THE INFLUENCE OF CONCEPTUAL KNOWLEDGE ON VISUAL DISCRIMINATION

Isabel Gauthier, Thomas W. James, and Kim M. Curby
Vanderbilt University, Nashville, USA

Michael J. Tarr
Brown University, Providence, USA

Does conceptual knowledge about objects influence their perceptual processing? There is some evidence for interactions between semantic and visual knowledge in tasks requiring both long-term memory and lexical access. Here we assessed whether similar perceptual/semantic interactions arise during sequential visual matching, a task that does not require access to semantic information. Matching of two-dimensional or three-dimensional novel objects was facilitated when the objects were associated with arbitrarily assigned distinctive artificial semantic concepts as compared to similar semantic concepts. In contrast to prior demonstrations, this effect was obtained in a task that did not require naming objects, and was not affected by participants rehearsing consonant strings, suggesting a direct influence from semantic associations on visual object recognition.

INTRODUCTION

Successful object recognition requires that visual input interface with visual memory. There is considerable evidence that conceptual knowledge that is nonvisual interacts with visual input to facilitate or interfere with recognition. For instance, supplying participants with verbal information about faces improves subsequent face recognition (Kerr & Winograd, 1982; Klatzky, Martin, & Kane, 1982), while supplying them with verbal information unrelated to a scene can facilitate recognition of that scene for days (Wiseman, MacLeod, & Lootsteen, 1985). In addition, learning to categorise objects into arbitrarily assigned sets changes the perceived similarity of those objects. For example, Goldstone and his colleagues (Goldstone, Lippa, &

Shiffrin, 2001) found that faces were perceived as less similar when participants first learned to group them into different categories.

Bub and colleagues (Arguin, Bub, Dixon, Caille, & Fontaine, 1996a; Arguin, Bub, & Dudek, 1996b; Dixon, Bub, & Arguin, 1997, 1998) have investigated the interaction between visual input and conceptual knowledge by associating arbitrary names with either novel or familiar objects. They report on a patient (ELM) who presents with severe prosopagnosia, as well as a category-specific agnosia for living things. Although ELM's perceptual abilities remain intact (for instance, he is able tell whether two faces presented simultaneously are the same or different), he is unable to identify or recognise faces or other living things such as animals, fruits, or vegetables. Interestingly, when

---

such objects are arbitrarily paired with names that are semantically distinct, ELM can learn to identify a limited number of stimuli (objects or faces that would normally be difficult for him to identify). Thus, when ELM was asked to learn name–face pairings between a set of three unfamiliar faces and three famous names that were semantically unrelated (e.g., a famous actor, a pop singer, and a politician), he committed a similar number of naming errors to normal controls. When the faces were paired with three famous names that were semantically related (e.g., three famous figure skaters), however, ELM was unable to learn the pairings and performed far worse than controls (Dixon et al., 1998).

ELM also participated in similar experiments using name–object pairings with familiar or novel nonface objects, and living and nonliving object names (Arguin et al., 1996a, 1996b; Dixon et al., 1997). Again, ELM's performance was more similar to that of normal controls when the object names were semantically unrelated than when the names were related. These results have been replicated with a second patient who has a similar agnosia to ELM's (Schweizer, Dixon, Westwood, & Piskopos, 2001) and with a group of patients

with Alzheimer's dementia (Dixon, Bub, Chertkow, & Arguin, 1999). Such results suggest that associating dissimilar semantic concepts with objects makes said objects more perceptually discriminable and thus easier to name.

Two types of conceptual influence on visual perception

Conceptual influences on perception may be classified into two different experimental methodologies (Figure 1). First, the *effects of category learning on perception* (category learning—CL) have been studied by manipulating the categorisation of visual stimuli. For example, participants learn to categorise four objects, two into Group A and two into Group B. In such experiments, similarity ratings or psychophysical discriminations differ before and after categorisation training and this difference depends on whether the judged objects were categorised into the same or different groups. After training, objects in the same group are discriminated more slowly and judged as more similar than are objects in the different group (Goldstone, 1994; Goldstone et al., 2001; Sigala, Gabbiani, & Logothetis, 2002). In these experiments, the
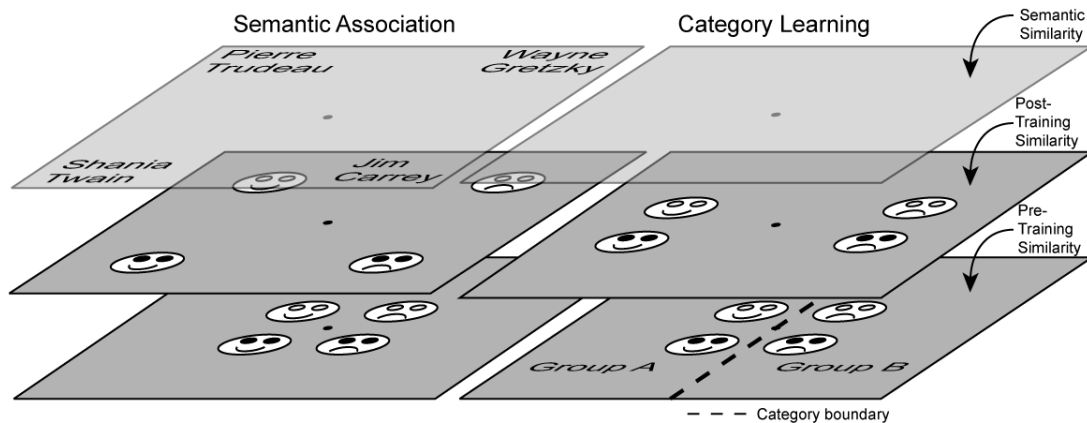
Figure 1. *Schematic representation of how perceived similarity of objects is affected by two different learning paradigms, semantic associations and category learning.* Semantic associations: *When names of famous people from dissimilar occupations are arbitrarily associated with face stimuli, those faces are perceived as more dissimilar after training. Post-training perceived similarity appears to depend on a combination of pre-training perceptual similarity and the similarity of the associated semantic information (Dixon et al., 1998).* Category learning: *When face stimuli are categorised into one of two groups, faces in different groups are perceived as more dissimilar after training. In this example, category membership was decided (arbitrarily) based on mouth expression, not on eye darkness. Post-training perceived similarity appears to depend mostly on the feature that was diagnostic for categorisation, in this case mouth expression (Goldstone et al., 2001).*

conceptual knowledge takes the form of the particular groups into which the objects were (arbitrarily) categorised. The speculation is that during category learning, perceptual dimensions that are diagnostic for a given category are given more weight than nondiagnostic dimensions, for instance, as described in the generalised context model (Nososfky, 1986).

Second, conceptual knowledge can influence perception *through the association of specific semantic features with objects* (semantic association—SA), as in the studies with patient ELM (Arguin et al., 1996a, 1996b; Dixon et al., 1997, 1998, 1999). In such experiments, the association of dissimilar semantic information with objects facilitates ELM's ability to name the objects.

Although both CL and SA are examples of a conceptual influence on perception, there are important differences between the two paradigms, in particular, how objects are categorised in the two cases as well as the relationship between perceptual and conceptual information. In SA experiments, objects are identified at the individual level during training (i.e., no two objects are given the same label), so that there is no between-versus within-category comparisons, the crucial manipulation in CL experiments. In contrast to CL experiments, where objects are explicitly put into categories, in SA manipulations any category would be implicitly formed (e.g., all faces given individual politicians' names can be grouped in a "politician" category). More importantly, however, in CL experiments the categories can be learned by attending selectively to one of the perceptual dimensions of the stimulus set (e.g., Goldstone et al., 2001), whereas in SA studies, there is typically no systematic relationship between perceptual features of the objects and the category (e.g., there is no perceptual dimension that separates faces of politicians from non-politicians). Thus, in an SA experiment, a novel stimulus (e.g., a new face) could not be categorised solely on the basis of perceptual information, whereas it could in a CL experiment.

One property common to both CL and SA experiments has been the use of tasks that involve recognition memory or naming. Because naming involves both semantic and lexical access, studies using naming tasks to investigate the interaction between visual processing and cognitive processes may not actually reflect an influence of semantic knowledge on early vision (Pylyshyn, 1999). Rather, such studies may be measuring a conceptual influence on the naming process. For instance, Humphreys, Lloyd-Jones, and Fias (1995) hypothesised that the influence of semantics on visual judgements occurs at the mapping between semantics and phonology. In other words, they claim that conceptual knowledge influences access to an object's name but not the perceptual processes that are performed as a precursor to naming.

Humphreys et al.'s argument was based on a "post-cue" procedure in which the presentation of a pair of coloured objects that are semantically related or unrelated is *followed by* a cue (the colour of one object) indicating which of the two objects is to be named. Semantic interference on naming is inferred when naming is slower for objects paired with semantically-related distractors. Using a modified version of the same post-cue procedure, Dean, Bub, and Masson (2001) obtained results suggesting that semantic interference affects more than object naming. As in the post-cue procedure, they presented a pair of objects that were either semantically related or unrelated followed by a cue. The cue, however, was an achromatic object (one of the pair) and the participants were to respond with the *colour* of the cued object. Crucially, in this case, retrieving the name of the object was not required to perform the task. Dean et al. still obtained evidence for semantic interference, suggesting that conceptual knowledge influences a visual task that does not require a naming response and, in fact, requires participants to remember only a single visual feature.

In summary, there appear to be at least two different ways in which conceptual knowledge can influence perceptual judgments. First, category learning affects perceptual judgments, presumably because perceptual information diagnostic for the learned categorisation is weighted more heavily in visual memory. Second, effects of semantic associations may reflect a more direct link between perceptual and conceptual representations. Such a link is

at least hinted at in studies demonstrating that associating nonperceptual information with particular objects improves naming performance in patients with category-specific visual agnosia, especially when the associated concepts are relatively distinctive. In the present study, we consider whether such semantic associations affect perceptual judgements in normal individuals. Furthermore, we explore whether these effects necessarily require mediation through naming, that is, can they be obtained in the context of visual matching judgements that neither require nor encourage naming.

Overview of experiments

In the following experiments, we examined the influence of semantic information on perceptual decisions in normal individuals. To circumvent the fact that semantics almost certainly influence visual processing through access to an object's name, we used novel objects (which have no names) and a visual sequential-matching task (in which naming is not required to generate a correct response). Critically, in the visual recognition literature sequential matching is often treated as cognitively impenetrable (e.g., Biederman & Gerhardstein, 1993). For instance, it is not necessary to endow an artificial system with semantic memories for it to be able to match two images of an object successfully (whether simultaneously or sequentially presented), even when those images present very different views of the object (Riesenhuber & Poggio, 1999).

This property of the sequential-matching task makes it an attractive tool for studying visual representation independently of other knowledge (Biederman & Gerhardstein, 1993; Ellis & Allport, 1986; Hayward & Williams, 2000; Lawson & Humphreys, 1996; Tarr, Williams, Hayward, & Gauthier, 1998) and to distinguish between perceptual and semantic impairments in neuropsychology. For instance, one of the factors used to distinguish patients with apperceptive and associative visual agnosia is their ability to perceptually match or copy objects (Farah, 1990; Humphreys & Riddoch, 1987; Kolb & Whishaw, 1996). It remains an open question, however,

whether the information used to make a sequential-matching decision is exclusively visual. Here, rather than address the important but thorny question of whether there is any stage of visual processing that is cognitively impenetrable (Pylyshyn, 1999), we ask a more practical question: can participants' performance in matching judgements be assumed to reflect only perceptual knowledge?

One concern of particular importance in the study of conceptual influences on visual perception is the common confound between the semantic similarity and the structural similarity in visual stimuli. That is, objects that are semantically related (e.g., different vehicles) also tend to be visually similar. To address this issue, some studies have used novel stimuli and/or arbitrarily assigned semantic descriptions to the stimuli (for examples, see Dixon et al., 1997, 1998). In the experiments presented here, we addressed this problem in a similar way; by arbitrarily associating semantic information with novel shapes and objects. In Experiment 1, we used two-dimensional shapes rotated in the picture-plane and arbitrarily associated them with conceptual labels from a single basic-level category (to create similar concepts associated with each shape; Rosch, Mervis, Gray, Johnson, & Boyes-Braem, 1976) or from multiple basic-level categories (to create dissimilar concepts associated with each shape). In Experiments 2 and 3, we used three-dimensional objects rotated in depth and associated them with "artificial concepts" that were comprised of sets of three semantic features. The amount of overlap of these features between sets determined the semantic similarity. For both experiments, viewpoint manipulations were included to reduce idiosyncratic strategies based on salient local features (Boucart & Humphreys, 1997). Moreover, how observers generalise across changes in viewpoint is consider one of the most critical aspects of visual recognition (Biederman & Gerhardstein, 1993; Tarr et al., 1998). Thus, an effect of conceptual knowledge in this task would illustrate the importance of considering such information in any theory of object recognition.

Our overarching hypothesis was that conceptual information would facilitate the discrimination of

visually similar objects when this information was semantically dissimilar. We expected a conceptual influence with only visually similar objects for one of the two following reasons: (1) visually similar objects are more confusable than visually dissimilar objects and thus would benefit from the addition of any type of information (including semantic) that would help distinguish them; and (2) visually similar objects take longer to differentiate than visually dissimilar objects, which in turn may allow more time for semantic associations to influence the discrimination process. Semantic knowledge, however, will only increase discriminability when the information associated with each object is relatively disparate, and therefore provides additional (useful) evidence for discriminating among them (Dixon et al., 1997, 1998). At the same time, it remains to be shown whether similar semantic associations actually interfere with object discriminations or not. The present study is not designed to address this question, in that we tested only the *relative* effect of similar and dissimilar semantic associations on visual recognition; whether any observed differences are due to facilitation or interference remains an open question.

EXPERIMENT 1

Participants learned to associate four words with four shapes. The shapes were visually similar or dissimilar, and the words belonged either to different or to a common basic-level category (e.g., four species of fish). Following training, we tested whether, for only visually similar shapes (the more difficult discrimination), responses in a sequential-matching task were facilitated for shapes associated with words from *different* categories.

Methods

Participants

Sixty-four participants who reported normal or corrected to normal vision participated in this experiment for payment or course credit (with informed written consent).

Materials

Sets of four 2-dimensional shape stimuli were selected from a larger set of 16 shapes (Figure 2) for the study and test phases of the experiment. Shapes were chosen to be either visually similar (V+) or visually dissimilar (V−). Shapes within a V+ set shared a common structure and could not be distinguished based on any one line-segment, whereas shapes within a V− set did not share a common structure. Thus, processing of the spatial relations of parts within a shape was necessary for V+ shape sets. In addition to the set of four shapes that was used during the study and test phases, two other shapes were also used for the test phase only (Figure 2), for a total of six test shapes for any given participant.

During the study phase, words were associated with the shape stimuli. Sets of 4 words were chosen from a larger set of 16 words that belonged to four categories: *birds* (crow, jay, pigeon, owl), *trees* (pine, palm, cedar, cypress), *flowers* (iris, violet, orchid, tulip), and *fish* (guppy, tuna, trout, carp). For the semantically similar condition (S+), words were chosen from within one category (e.g., crow, jay, pigeon, owl), whereas for the semantically dissimilar condition (S−), words were chosen across categories (e.g., crow, pine, iris, guppy).
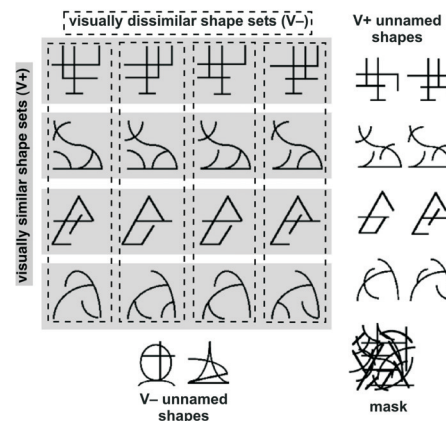

Figure 2. *Stimuli used in Experiment 1. Each participant learned four shapes, from either one of the four visually similar sets (V+, shaded areas) or from one of the four visually dissimilar sets (V−, dashed areas). The unnamed shapes used during sequential-matching test tasks in V+ and V− conditions are shown (two unnamed shapes per condition) as well as the pattern mask.*

All stimuli were presented to participants on a CRT monitor via an Apple Macintosh G3 computer running RSVP software (http://www.cog.brown.edu/~tarr/rsvp.html).

Design and procedures

Each participant was randomly assigned to one of four groups and learned four shapes and four words. The four groups were arrived at by crossing the V+ and V− conditions with the S+ and S− conditions, yielding four groups (V+S+, V+S−, V−S+, V−S−) with 16 participants each. Shape–word combinations were counterbalanced so that no two participants associated the same words with the same shapes. Participants were told that they were learning about four shapes on planet Zol, "where things look very different."

The experiment was divided into a study phase and a test phase. During the first part of the study phase, each shape was shown eight times, simultaneously with its word label, for 5 s. During the second part of the study phase, participants matched the shapes to these same-word labels. On each matching trial, a word was presented above the set of four shapes. Each shape was identified by a number that was presented below it on the screen. Participants responded by pressing the number key on the keyboard that corresponded to the shape that matched with the word. To prevent associations between the shapes and locations on the screen, and to encourage learning of shape–word pairings, the positions of the four shapes on the screen varied on each trial. Matching trials were presented in blocks of eight trials. Participants completed a minimum of three blocks and continued until they were able to complete two consecutive blocks with no errors.

After the shape–label pairings were learned to criterion, participants performed a speeded sequential-matching task. Six shapes were used in this test phase, the four shapes that were learned during the study phase, plus two more (Figure 2). Each trial consisted of the following events in order: a fixation cross for 200 ms, the first shape (S1) for 500 ms, a pattern mask for 1000 ms, and the second shape (S2). S2 was displayed until a "same" or "different" response was made (participants were asked to judge whether the two shapes were identical, although they could be rotated). S1 was always one of the four shapes from the study phase and was always shown at its studied orientation, denoted as zero degrees (0°). S2 was one of the four studied shapes or one of two unstudied shapes. S2 was rotated in the picture plane 0°, 55°, 110°, or 165° from the studied orientation. For each orientation, two shapes were rotated clockwise and the other two shapes were rotated counterclockwise. In total, there were 160 trials of the following types: 80 same trials, 48 different trials with a different studied shape, and 32 different trials with a different unstudied shape.

Results and discussion

During the study phase, the number of blocks of matching trials that were required to learn the shape–word pairings was primarily determined by the visual similarity of the shapes. A 2 × 2 ANOVA with visual similarity and semantic similarity revealed a significant effect of visual similarity, $F(1, 60) = 7.9$, $MS_E = 5.957$, $p \leq .01$, but no effect of semantic similarity, $F < 1.0$, n.s., and no interaction between visual similarity and semantic similarity, $F < 1.0$, n.s. The visually similar shapes ($M_{(S−V+)} = 4.40$, $M_{(S+V+)} = 5.30$) required longer to learn than the visually dissimilar shapes ($M_{(S−V−)} = 3.06$, $M_{(S+V−)} = 3.13$). Importantly, there was no evidence that the level of semantic similarity influenced how quickly the participants learned the shape–word pairings.

For the post-training sequential-matching test task, both mean sensitivity (d′)[1] and mean response time for correct responses (RT)[2] were analysed. A 2 × 2 × 4 split-plot ANOVA with visual similarity and semantic similarity as between-subjects factors and orientation as a within-subjects factor was

---

[1] Mean hits and correct rejections were also analysed and showed the same pattern of results as d′.

[2] Mean RT was calculated as a geometric mean (Alf & Grossberg, 1979) for correct responses only. This estimate of central tendency is less susceptible to outliers than the arithmetic mean. See also Gauthier, Behrmann, and Tarr (1999) and Gauthier, Williams, Tarr, and Tanaka (1998).

performed on d′ and RT. As expected (Tarr & Pinker, 1989; Gauthier & Tarr, 1997), there was a significant main effect of orientation for both d′, $F(3, 180) = 4.14$, $MS_E = 1.30$, $p < .01^3$, and RT, $F(3, 180) = 40.2$, $MS_E = 3755.1$, $p < .001$. Also expected, due to the obvious difference in difficulty between the visually similar and visually dissimilar conditions, was a significant main effect of visual similarity for both d′, $F(1, 60) = 17.2$, $MS_E = 27.5$, $p < .001$, and RT, $F(1, 60) = 41.2$, $MS_E = 111861.6$, $p < .001$. There was also a significant interaction for RT between orientation and visual similarity, $F(3, 180) = 8.96$, $MS_E = 3755.1$, $p < .001$, due to a less pronounced effect of orientation in the V− condition compared to the V+ condition.

The interaction between visual similarity and semantic similarity that was predicted based on our hypothesis was not significant for either d′, $F(1, 60) = 2.63$, n.s., or RT, $F(1, 60) = 2.30$, n.s.; see Figure 3.[4]

To address our specific a priori hypothesis that performance for the V+S− condition would be better than for the V+S+ condition, we conducted a focused t-test. This comparison revealed no significant difference in d′ between these conditions, $t(30) = 1.61$, $p = .075$, with greater sensitivity for the V+S− condition. There was no significant difference in RT between the conditions, $t(30) = 1.11$, $p = .28$, however, the V+S− condition was faster than the V+S+ condition, ruling out a possible speed accuracy tradeoff. Thus, although our results were not conclusive, they do suggest that a more powerful experimental design might reveal effects consistent with our predictions.

EXPERIMENT 2

In Experiment 2, we again tested whether arbitrary associations between semantic information and novel objects would affect subsequent performance in a perceptual task. To increase the influence of semantics on perceived similarity, we changed the


Figure 3. *Mean sensitivity and response times for correct responses in Experiment 1. Error bars show the standard error of the mean.*

way that semantic associations were generated. Participants learned to associate "artificial concepts" with three-dimensional novel objects (Figure 3). The objects were all visually similar and thus there was no visually dissimilar condition in Experiment 2. The artificial concepts were triads of nonvisual semantic features, which is an important difference between Experiment 1 and Experiment

---

[3] F statistics, MS_E, and p values have been corrected using Greenhouse-Geisser epsilon when necessary.

[4] Tests of significance for the visual similarity and semantic similarity factors were calculated in two ways, collapsing across orientation and by isolating only the 0 orientation. These analyses always produced the same effects.

2. With the shape–word pairings used in Experiment 1, activation of the word concept activated not only nonvisual semantic features, but also visual semantic features (e.g., a "crow" <is black>, <has wings>, etc.). The features that were used to create the artificial concepts were all nonvisual semantic features (Table 1). Also, sets of four artificial concepts were either made semantically similar (S+) or semantically dissimilar (S−) based on the overlap between features (Table 1), which was manipulated experimentally. Thus, the use of artificial concepts instead of single count noun labels may provide both a stronger and more direct test of semantic–perceptual interactions.

We again used a sequential-matching task to measure post-training performance with the objects. As mentioned earlier, it was unlikely that naming influences a sequential-matching task and participants did report that they were not naming the shapes in Experiment 1. Nevertheless, in an effort to further discourage participants from covertly naming the objects we introduced a subvocal verbal rehearsal task (Baddeley, 1986) for half of the sequential-matching trials. Furthermore, in Experiment 1, participants might have been encouraged to name the objects because their training task was to associate a single word with

each shape and that word was always a concrete noun. In contrast, the objects in Experiment 2 were associated with triads of nonvisual semantic features.

Methods

Participants
Thirty-two participants who reported normal or corrected to normal vision participated for payment or course credit (with informed written consent).

Materials
Four novel three-dimensional objects (YUFOs; Figure 4) were used. The four YUFOs were highly visually similar and could only be discriminated using subtle differences in shape. They were created using FormZ (Autodessys Inc., Columbus, OH) and rendered with a blue texture in Lightscape (Lightscape Technologies, Inc., San Jose, CA) to create highly realistic images. The four objects were rendered at an arbitrary canonical pose (referred to as 0°) as well as four other viewpoints generated by progressive 30° rotations in depth around the vertical axis. Artificial concepts were triads of features selected from a larger pool of 16 features. These features were "fast," "flexible," "friendly," "cold,"

Table 1. *Examples of similar and dissimilar nonvisual artificial concepts used in Experiment 2*

	Object 1	Object 2	Object 3	Object 4
Semantically dissimilar				
fast	✓			
flexible	✓			
friendly	✓			
cold		✓		
rare		✓		
sweet		✓		
fragile			✓	
hollow			✓	
nervous			✓	
sticky				✓
soft				✓
wet				✓
Semantically similar				
loud	✓	✓	✓	
heavy		✓	✓	✓
nocturnal	✓		✓	✓
strong	✓	✓		✓
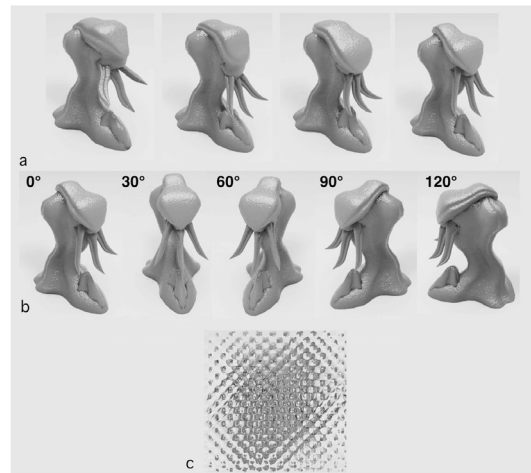
Figure 4. *Stimuli used in Experiment 2 (YUFOs). (a) The four YUFOs learned by each participant; (b) the five orientations in depth; and (c) the pattern mask used in the sequential-matching test task.*

"rare," "sweet," "fragile," "hollow," "nervous," "sticky," "soft," "wet," "loud," "heavy," "nocturnal," and "strong". In the semantically dissimilar (S−) condition, four artificial concepts were created by selecting three features per concept that did not overlap between concepts. Thus, for the S− condition, 12 features were used to create four concepts (Table 1). In the semantically similar (S+) condition, four concepts were created from just four features. Thus, each feature was used in three of the four concepts, leading to considerably feature overlap between the concepts (Table 1). The object–concept combinations were counterbalanced as in Experiment 1.

All stimuli were presented to participants on a CRT monitor via a MacIntosh G3 computer running RSVP software. Six sheets of paper with the pictures of the four YUFOs on them were also given to the participants.

Design and procedures

Participants were randomly assigned to either the S+ or S− conditions. The experiment was divided into a study phase and a test phase. For the study phase, a different training procedure was used for Experiment 2 than was used for Experiment 1, due to the assumption that learning to associate semantic feature triads with objects would be more difficult than associating single words.

During the first part of the study phase, each object was shown four times, simultaneously with two features from its three-feature concept, for 5 s. During the second part of the study phase, participants answered 16 questions about the association between pairs of features and the objects (e.g., "Is this one *cold* and *flexible*?"). During the third part of the study phase, two objects were shown six times, simultaneously with all three features from their associated concept, for 5 s. During the fourth part of the study phase, participants answered 24 questions about the association between single features and these two objects (e.g., "Is this one *sticky*?"). During the fifth part of the study phase, the remaining two objects were shown six times, simultaneously with all three features from the associated concept, for 5 s. During the sixth part of the study phase, participants then answered 50 questions about associations between single features and all four objects. These six parts of the study phase (110 trials) were then repeated, for a total of 220 study trials. To facilitate learning, at six points during the study phase, participants were asked to write down all of the adjectives (features) that they could remember about each object, on sheets of paper depicting pictures of the four YUFOs. This occurred after Part 2, Part 5, and Part 6 during each of the two repetitions.

In the final part of the study phase, matched triads of features with the correct object until they reached a criterion. On each matching trial, the three features of a single artificial concept were presented above three of the four YUFOs. Each YUFO was identified by a number that was presented below it on the screen. Participants responded by pressing the number key on the keyboard that corresponded to the YUFO that matched with the concept. To prevent associations between the YUFOs and locations on the screen, and to encourage learning of object–concept pairings, the positions of each YUFO on the screen varied on each trial. Matching trials were presented in blocks of 12 trials. Participants completed a minimum of two blocks and continued until they were able to complete a block with only two errors.

After the object–concept pairings were learned to criterion, participants performed a speeded sequential-matching task in the test phase of the experiment. Each trial consisted of the following events in order: a fixation cross for 1000 ms, the first YUFO (O1) for 1500 ms, a mask for 500 ms, and the second YUFO (O2). O2 was displayed until a "same" or "different" response was made. O1 was always shown at its studied orientation, denoted 0° here. O2 was about the vertical axis 0°, 30°, 60°, 90°, or 120° from the studied orientation. O1 and O2 were the same YUFO for half of the trials. It was made clear to the participants that they were to make this decision based on the images that were presented on the screen. No mention was made of the adjectives that had been associated with the objects. Debriefing revealed that a several participants believed that the matching part of the experiment was in fact a separate experiment that

was not related to the association phase. Thus, it seemed unlikely that participants were explicitly recalling the associated information to aid them in matching task. To further study this possibility, though, participants performed two blocks of 120 trials and during one block of trials, they performed a subvocal verbal-rehearsal task. During the verbal rehearsal block (the order of which was counterbalanced), participants were presented with a string of seven consonants (e.g., H K T V X C F) and then rehearsed them while performing the sequential-matching trials. After each set of 20 trials, participants were prompted to recall the string, after which they were presented with a new string for the next 20 trials.

Results and discussion

The number of blocks required to reach criterion during the final part of the study phase was small and not significantly different between the S– condition (M = 3.06, SE = 0.25) and the S+ condition (M = 4.31, SE = 3.11); $t(30)$ = 1.46, n.s. The number of blocks required for the S+ condition was more variable due to two participants who took longer to reach criterion (10 and 13 blocks). Apart from these two participants, all others reached criterion in four blocks or less. The mean number of blocks without the largest two values in each condition was: S+, 3.21, SE = 0.11, and S–, 3.00, SE = 0.0; $t(26)$ = 1.88, n.s. Thus, there was little evidence that the level of semantic similarity influenced how quickly the participants learned the object–concept pairings. Although it may seem paradoxical that similar concepts were learned as quickly as dissimilar concepts, it should be noted that highly similar objects had to be discriminated simultaneously with learning the semantic associations. Because of the high degree of visual similarity of the novel objects, it is likely that the factor that constrained the number of trials to reach criterion was the visual similarity (which was the same) as opposed to the semantic similarity (which was varied).

For the post-training sequential-matching test task, both mean sensitivity (d′) and mean response time for correct responses (RT) were analysed. A 2 × 2 × 4 split-plot ANOVA with semantic similarity

as a between-subjects factor and verbal rehearsal and orientation as within-subjects factors was performed on d′ and RT. Again, there was a significant main effect of orientation for both d′, $F(4, 120)$ = 41.6, MS_E = 15.1, $p < .001$, and RT, $F(4, 120)$ = 19.5, MS_E = 192683.3, $p < .001$, but orientation did not interact with the other factors.

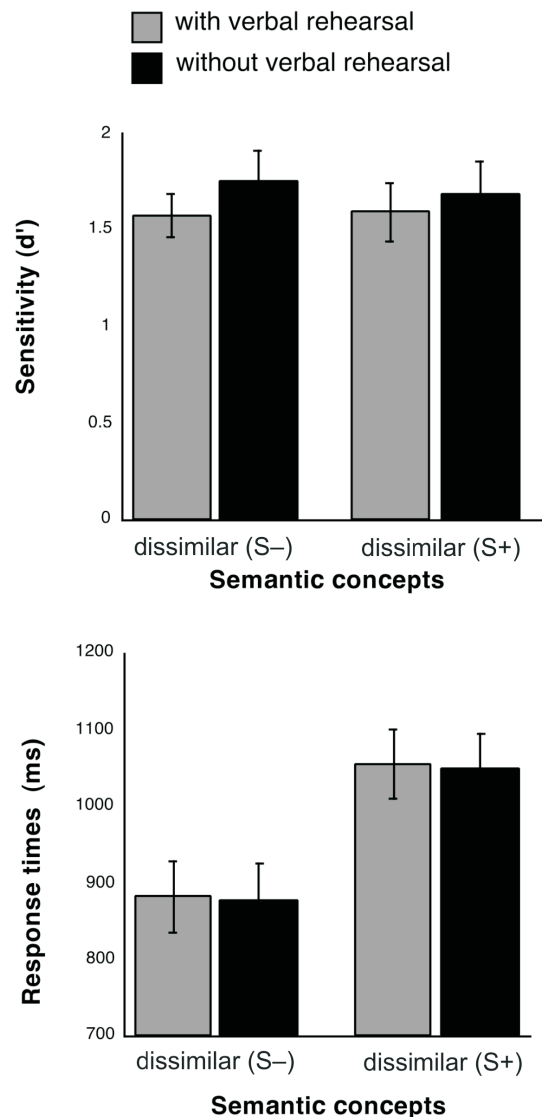As illustrated in Figure 5, the effect of semantic similarity was obtained in RT, $F(1, 30)$ = 8.02, MS_E

Figure 5. *Mean sensitivity and response times for correct responses in Experiment 2. Error bars show the standard error of the mean.*

= 2347032.7, p < .01, but not in d′, $F(1, 30)$ < 1.0. There was no significant effect of verbal rehearsal on d′, $F(1, 30)$ = 3.92, MS_E = 0.382, n.s., nor RT $F(1, 30)$ < 1.0. Importantly, there was no interaction of verbal rehearsal with semantic similarity for either d′ or RT $F(1, 30)$ < 1.0, suggesting that performing a sequential-matching task is more difficult during verbally rehearsal but that verbal rehearsal did not disrupt the effect of semantic similarity on the ability to perceptually match the objects.

Relatively dissimilar semantic associations facilitated the speed of matching over similar semantic associations three-dimensional objects. This result was in the same direction as the nonsignificant result described in Experiment 1. In addition, manipulating verbal rehearsal in Experiment 2 demonstrated that the effects of semantic similarity were not due to covert object naming.

By associating objects with artificial concepts comprised of triads of nonvisual semantic features instead of using single concrete nouns, we removed the possibility that the effects of semantic similarity were due to semantic information regarding visual properties. For instance, associating the label "crow" with an object would activate semantic information about visual properties (e.g., <has wings>, <is black>) in addition to activating semantic information about nonvisual properties (e.g., <eats carrion>, <is light>). Although dividing semantic information into only visual and nonvisual features may be overly simplistic in some contexts (McRae & Cree, 2002), because we are dealing with the visual perception of objects, this division suffices for the point we wish to make. Specifically, any influence obtained using the procedure in Experiment 1 could be interpreted as arising due to associations developed between *visual* semantic information and the stimulus shapes, with no contribution from the *nonvisual* semantic information. Because only nonvisual semantic features were used in Experiment 2, these results provide evidence that semantic information influences visual perception more generally. That is, although semantic information can be divided into different types, a correspondence between the perceptual sensory modality and the type of semantic information is not a necessary condition for perception and semantics to interact.

EXPERIMENT 3

In Experiment 2, we obtained an effect of semantic associations on response times. To investigate whether this effect could also be observed in sensitivity, we made the sequential-matching task from Experiment 2 more difficult by restricting the presentation time of the second stimulus.

Methods

Participants
Thirty-two participants who reported normal or corrected to normal vision participated for payment or course credit (with informed written consent).

Materials
The same stimuli as in Experiment 2 were used.

Design and procedures
Procedures were identical to those in Experiment 2, with the following differences in the sequential matching task that followed the study phase. Each trial consisted of the following events in order: a fixation cross for 1000 ms, the first YUFO presented centrally for 1500 ms, a 500 ms blank followed by the second YUFO for 175 ms, presented centred up to 70 pixels away from the centre of the screen (the position was randomly determined so that the YUFOs, each 283 pixels high × 203 pixels wide, would be randomly displayed within a window about 1.5 times their height and 1.7 times their width) and finally a square pattern mask (450 pixels wide) for 200 ms.[5] There was no verbal distractor task during the test phase.

---

[5] Details of the sequential matching task in Experiment 3 were set to provide a reasonable match to similar studies using lateralised presentation, albeit with central presentation (Curby, Hayward, & Gauthier, 2003). In particular, the second stimulus was masked (as it would be when presented to only one hemisphere) and it appeared with some position uncertainty.

Results and discussion

The number of blocks required to reach criterion during the final part of the study phase was not significantly different between the S− ($M = 3.06$, SE = 0.06) and S+ conditions ($M = 5.50$, SE = 1.61); $t(30) = 1.51$, n.s. The number of blocks to reach criterion was highly variable in the S+ condition due to two participants who took 20 blocks to reach criterion. Apart from these two participants, all others reached criterion in four blocks or less. Without these two participants, the mean number of blocks in the S+ condition was 3.14, SE = 0.10, and that in the S− condition was 3.00, SE = 0.0; $t(26) = 1.47$, n.s.

Sensitivity and mean RT in the sequential matching judgements are shown in Figure 6. As expected, the effect of the semantic association on sequential matching was obtained in sensitivity, $F(1, 120) = 4.12$, $MS_E = 3.170$, $p = .05$, which also showed an effect of orientation, $F(4, 120) = 24.62$, $MS_E = 5.912$, $p < .0001$, but no interaction of orientation with condition ($F < 1.0$). An ANOVA on mean response time for correct responses only revealed a significant effect of orientation, $F(4, 120) = 8.00$, $MS_E = 36277.3$, $p < .0001$. Response times in Experiment 3 were considerably faster than in Experiment 2, as can be appreciated by comparing Figures 5 and 6. This is undoubtedly due to the fast presentation of the second stimulus, preventing extended inspection times as permitted in Experiment 2.

Thus, depending on the allotted time course of visual processing, effects of nonvisual knowledge may manifest themselves in either the speed of processing or the accuracy of the discrimination. Here our results indicate that when the visual task is more difficult due to shortened processing time, relatively dissimilar semantic information associated with novel objects improves sensitivity on subsequent visual matching judgements.

GENERAL DISCUSSION

We hypothesised that arbitrarily associating semantic information with objects would influence subsequent perceptual judgements of those objects. In addition, we argued that the influence should be independent of object naming. These hypotheses were derived from previous research on categorisation (Goldstone, 1994; Goldstone et al., 2001) and category-specific agnosia (Arguin et al., 1996a, 1996b; Dixon et al., 1997, 1998). Our results support these hypotheses and extend prior work in that we combined a manipulation consisting of the distinctiveness of nonperceptual information


Figure 6. *Mean sensitivity and response times for correct responses in Experiment 3. Error bars show the standard error of the mean.*

associated with novel objects (as in studies with patients with category-specific agnosia) with a task that could be performed on the basis of perceptual information alone (as in studies of the effects of category learning), rather than using naming judgements. Associating novel objects with relatively semantically dissimilar concepts produced relatively better performance in sequential-matching judgements (as expressed either in response times or sensitivity, depending on task difficulty) than associating objects with semantically-similar concepts.

The influence of semantics on perception did not reach significance when the semantic information was a single concrete noun and when the objects were two-dimensional shapes rather than three-dimensional shapes. However, when two-dimensional shapes were visually similar the effect was in the same direction as for three-dimensional objects associated with "artificial concepts" (triads of semantic features), so it would seem that the same processing principles apply rather broadly across stimulus categories. We also found that the effect of semantic similarity can be obtained either in sensitivity or in response times depending on the specifics of the matching task: One possibility is that neither dependent variable reached significance in Experiment 1 because the effect was split between domains (visually similar shapes were matched both more accurately and faster when associated with semantically dissimilar concepts).

One somewhat surprising result is that we found no evidence that the associated conceptual information interacted with the viewpoint effects obtained when objects were matched across different orientations. One interpretation of this finding is that the "additive" nature of viewpoint and semantic similarity implies separate processing stages (Sternberg, 1966)—the former perceptual and the latter conceptual. However, given that we still do not have a clear model of how observers compensate for viewpoint (Tarr & Bülthoff, 1998), it is difficult to attribute viewpoint normalisation to a perceptual system and the effects of semantic similarity to a conceptual system. In particular, more recent models of viewpoint normalisation rely on an "accumulation of evidence" approach (see Perrett,

Oram, & Ashbridge, 1998) that, plausibly, is occurring at an earlier point in object processing than once thought. Consistent with this hypothesis, a recent fMRI study comparing viewpoint effects for object recognition judgements to those obtained in mental rotation judgements found that activity related to viewpoint specifically during object recognition occurred at several points in the occipito-temporal pathway, including as early as BA 19 (Gauthier, Hayward, Tarr, Anderson, Skudlarski, & Gore, 2002). Thus, a great deal of *visual* processing may occur post-normalisation—for instance, visual processing to bind features together or to extract object structure. In short, there is a great deal of room for perceptually mediated semantic effects. Thus, the most that can be said is that the absence of an interaction between viewpoint and conceptual similarity suggests that these two factors recruit different stages of perceptual processing.

There is another interesting inference that may be based on the lack of an interaction between semantic similarity and viewpoint, namely that the effect of semantic similarity was not dependent on the time required to discriminate the objects. This supposition is further supported by the results of Experiment 3, in which participants responded almost twice as quickly as in Experiment 2, yet showed the same effect of semantic similarity. Taken together, these results suggest that semantic similarity interacted with visual similarity and not with the time taken to discriminate the objects; however, this issue deserves further investigation.

Semantic, perceptual, and phonological interactions

We obtained an effect of semantic similarity on visual matching regardless of whether the matching task was performed while participants engaged in verbal rehearsal or not, that is, semantic similarity produced an effect under conditions that would generally not be expected to be affected by nonperceptual manipulations. Based on this result, we propose that semantic knowledge influences perceptual decisions without mediation by other

nonvisual processes, such as phonological lookup or lexical access.

This claim stands in contrast to one by Humphreys and colleagues, who argued that semantic influences on perception are attributable to interference between semantic and phonological representations (Humphreys et al., 1995) or interference between two semantic representations (Boucart & Humphreys, 1992, 1994; Humphreys & Boucart, 1997). Beyond the post-cue procedure described earlier, Humphreys and colleagues also found that visual matching judgements that required attending to global shape were subject to interference from semantic information. As in the post-cue procedure, semantic interference was attributed to the semantic similarity between two matched images (e.g., more interference for matching the orientation of lines overlapping a helicopter and a truck than lines overlapping a helicopter and a rabbit). However, in this case the associations between semantic and visual features are not completely arbitrary, as they were in our design, and at least part of their effect may be attributed to greater visual similarity within than between classes (something that is accepted for visual categories). In contrast, our results indicate that matching two images of the *same object* can be influenced by an experimental manipulation of the semantic similarity of concepts associated with each of the objects in a stimulus set.

Contextual effects of semantic associations

Having hypothesised that semantics does influence perception, that is, without mediation from other processes, we would caution that the posited interaction between perceptual and semantic representations need not have occurred during the test phase of our experiment, as we originally assumed. Our semantic manipulation may have exerted its influence during the study phase, during the test phase, or it may have influenced both phases. Our working assumption has been that associations are formed during the study phase between the perceptual representation of a given object and the semantic representation of the specified semantic information. During the test phase, these associations caused the

automatic activation of associated semantic knowledge that then influences performance. An alternative explanation is that studying the novel objects in the presence of conceptual information changed the way that the perceptual representations were encoded. That is, the semantic information acted as a context within which the perceptual representations were learned. During the test phase, these different perceptual representations would produce differences in performance. Some categorisation research has favoured the latter account, in which categorisation influences the nature of the perceptual representation (Goldstone, 1994; Goldstone et al., 2001; Schyns & Rodet, 1997). A variant of this idea is that the learning phases of the two semantic conditions differed in difficulty and that this difference caused differences in the way that the perceptual representations of the objects were created. Specifically, if the S+ condition was more difficult, then the objects that were learned in the S+ condition would develop poorer representations, which might produce poorer performance during the test phase.

Our experiments were not designed to address this question directly; however, we did not find evidence of differential learning in the two semantic conditions in any of the experiments. This was perhaps because study phase duration was limited primarily by the difficulty of learning to discriminate the visually similar objects and not by learning of the artificial concepts. Interestingly, a small number of participants took a very long time to reach criterion in the study phase of Experiments 2 and 3—and these were all in the S+ condition. Although this suggests that the S+ condition may indeed have been slightly more difficult, it also suggests that participants in the S+ condition received more exposure to the objects than participants in the S− condition. Such additional experience should have, if anything, led to *better* performance during the test phase for the S+ condition, contrary to the results obtained. Thus, it is difficult to make any definitive claims about the precise locus of the effect of any semantic associations learned simultaneously with visual information. Whether different perceptual representations were built at study or associations automatically engaged at test is impor-

tant to the question of the modularity of visual processes (Fodor, 1983; Pylyshyn, 1999), but in practice, both situations still suggest that it is difficult to study visual processes independently of semantic influences. As such, it may be unwarranted to assume that any visual task is protected from such influences or that novel visual stimuli would guard against this possibility.

Implications of semantic–perceptual interactions

One implication of our findings concerns the interpretation of performance in tasks that can in principle be performed solely on the basis of perceptual information. We obtained an influence of semantic similarity in such a task and our effects arose from semantic–perceptual associations learned over a short period of time (much less than 1 hour). Furthermore, our criterion for having learned the associations successfully was not speeded and required only that participants remember them correctly. Given these parameters for learning arbitrary associations, our effects could be comparable to those of similar nonvisual associations spontaneously generated by participants as they become familiar with novel objects learned in many studies of visual cognition. Moreover, similar effects would be expected to be even more pronounced for familiar objects that have a rich pool of already-established semantic associations.

What is made clear by these results is that the field should be cautious in invoking the oft-made distinction between low-level perceptual and high-level visual recognition tasks. Such a distinction is typical in neuropsychology, where researchers often distinguish between brain-injured participants who can perform perceptual matching tasks and those who cannot (Farah, 1990; Humphreys & Riddoch, 1987; Kolb & Whishaw, 1996). Our present findings suggest that the low-level/high-level dichotomy may not be as strong as once supposed. For instance, brain-injured participants may perform more poorly than normal controls on perceptual tasks with novel objects because of perceptual impairments (the standard explanation) or, as implied by our results, due to an impaired ability to

generate semantic associations with those objects. The extent to which participants automatically generate semantic attributes for novel objects they learn in the laboratory is often ignored and its impact may need to be reconsidered.

Semantic–perceptual interactions investigated with neuroimaging

We close with the suggestion that one means for addressing some of the open questions in our work is through neuroimaging. Such methods might prove useful in differentiating between two possible mechanisms that may underlie our effect. Previous neuroimaging studies have determined that the posterior cortex is functionally heterogeneous. For example, regions have been found that respond preferentially to different perceptual attributes of objects, such as their colour, form, and movements (Corbetta, Miezin, Dobmeyer, Shulman, & Petersen, 1991). More recently, regions have been identified that respond selectively during semantic retrieval of object-related knowledge. Interestingly, these regions also appear to be somewhat heterogeneous. Martin and his colleagues (Martin, Ungerleider, & Haxby, 1999) have investigated the neural substrates that underlie the representation of both perceptual and semantic features or attributes. Their work suggests semantic and perceptual information related to a particular attribute (i.e., colour) is stored in neighbouring regions of cortex. For instance, perceptual access to an attribute activates an area of cortex that neighbours the area of cortex that is activated during semantic access to the same attribute. In particular, the regions that respond to semantic access appear to be just anterior to the related regions that respond to perceptual access. Thus, combining neuroimaging techniques with our training and testing procedure may allow us to identify some of the neural substrates through which an influence of semantics on perceptual judgements occurs.

Conclusion

The association of nonvisual semantic features with novel objects can influence object discrimination

judgements that may be made on the basis of visual information alone. Beyond complicating the task of studying perceptual processes in isolation from other higher-level processes (contrary to Fodor, 1983), it is our belief that future models of visual processing should more carefully consider the role of top-down knowledge (e.g., Mumford, 1992) and, in particular, whether or not perceptual systems can be engaged independently of non-perceptual systems.

REFERENCES

Alf, E. F., & Grossberg, J. M. (1979). The geometric mean: Confidence limits and significance tests. *Perception and Psychophysics*, *26*, 419–421.

Arguin, M., Bub, D., Dixon, M., Caille, S., & Fontaine, S. (1996a). 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. (1996b). Shape integration for visual object recognition and its implication in category-specific visual agnosia. *Visual Cognition*, *3*, 221–275.

Baddeley, A. D. (1986). *Working memory*, London: Oxford University Press.

Biederman, I., & Gerhardstein, P. C. (1993). Recognizing depth-rotated objects: Evidence and conditions for three-dimensional viewpoint invariance. *Journal of Experimental Psychology: Human Perception and Performance*, *19*, 1162–1182.

Boucart, M., & Humphreys, G. W. (1992). Global shape cannot be attended without object identification. *Journal of Experimental Psychology: Human Perception and Performance*, *18*, 785–806.

Boucart, M., & Humphreys, G. W. (1994). Attention to orientation, size, luminance, and color: Attentional failure within the form domain. *Journal of Experimental Psychology: Human Perception and Performance*, *20*, 61–80.

Boucart, M., & Humphreys, G. W. (1997). Integration of physical and semantic information in object processing. *Perception*, *26*, 1197–1209.

Corbetta, M., Miezin, F. M., Dobmeyer, S., Shulman, G. L., & Petersen, S. E. (1991). Selective and divided attention during visual discriminations of shape, color, and speed: Functional anatomy by positron emission tomography. *The Journal of Neuroscience*, *11*, 2383–2402.

Curby, K. M., Hayward, W. G., & Gauthier, I. (2003). *Laterality effects in the recognition of depth rotated objects.* Manuscript submitted for publication.

Dean, M. P., Bub, D. N., & Masson, M. E. J. (2001). Interference from related items in object identification. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *27*, 733–743

Dixon, M. J., Bub, D. N., & Arguin, M. (1997). The interaction of object form and object meaning in the identification performance of a patient with category-specific visual agnosia. *Cognitive Neuropsychology*, *14*, 1085–1130.

Dixon, M. J., Bub, D. N., & Arguin, M. (1998). Semantic and visual determinants of face recognition in a prosopagnosic patient. *Journal of Cognitive Neuroscience*, *10*, 362–376.

Dixon, M. J., Bub, D. N., Chertkow, H., & Arguin, M. (1999). Object identification deficits in dementia of the Alzheimer type: Combined effects of semantic and visual proximity. *Journal of the International Neuropsychological Society*, *5*, 330–345.

Ellis, R., & Allport, D. A. (1986). Multiple levels of representation for visual objects: A behavioural study. In A. G. Cohn & J. R. Thomas (Eds.), *Artificial intelligence and its applications* (pp. 245–247). New York: Wiley.

Farah, M. J. (1990). *Visual agnosia: Disorders of object recognition and what they tell us about normal vision.* Cambridge, MA: MIT Press.

Fodor, J. A. (1983). *Modularity of mind.* Cambridge, MA: The MIT Press.

Gauthier, I., Behrmann, M., & Tarr, M. J. (1999). Can face recognition really be dissociated from object recognition? *Journal of Cognitive Neuroscience*, *11*, 349–370.

Gauthier, I., Hayward, W. G., Tarr, M. J., Anderson, A. W., Skudlarski, P., & Gore, J. C. (2002). Bold activity during mental rotation and viewpoint-dependent object recognition. *Neuron*, *34*, 161–171.

Gauthier, I., & Tarr, M. J. (1997). Orientation priming of novel shapes in the context of viewpoint-dependent recognition. *Perception*, *26*, 51–73.

Gauthier, I., Williams, P., Tarr, M. J., & Tanaka, J. W. (1998). Training "Greeble" experts: A framework for studying expert object recognition processes. *Vision Research*, *38*, 2401–2428.

Goldstone, R. (1994). Influences of categorization on perceptual discrimination. *Journal of Experimental Psychology: General*, *123*, 178–200.

Goldstone, R., Lippa, Y., & Shiffrin, R. M. (2001). Altering object representations through category learning. *Cognition*, *78*, 27–43.

Hayward, W. G., & Williams, P. (2000). Viewpoint dependence and object discriminability. *Psychological Science*, *11*, 7–12.

Humphreys, G. W., & Boucart, M. (1997). Selection by color and form in vision. *Journal of Experimental Psychology: Human Perception and Performance*, *23*, 136–153.

Humphreys, G. W., Lloyd-Jones, T. J., & Fias, W. (1995). Semantic interference effects on naming using a postcue procedure: Tapping the links between semantics and phonology with pictures and words. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *21*, 961–980.

Humphreys, G. W., & Riddoch, M. J. (1987). The fractionation of visual agnosia. In G. W. Humphreys & M. J. Riddoch (Eds.), *Visual object processing: A cognitive neuropsychological approach*. Hove, UK: Lawrence Erlbaum Associates Ltd.

Kerr, N. H., & Winograd, E. (1982). Effects of contextual elaboration on face recognition. *Memory and Cognition*, *10*, 603–609.

Klatzky, R. L., Martin, G. L., & Kane, R. A. (1982). Semantic interpretation effects on memory for faces, *Memory and Cognition*, *10*, 195–206.

Kolb, B., & Whishaw, I. Q. (1996). *Fundamentals of human neuropsychology*. New York: W. H. Freeman.

Lawson, R., & Humphreys, G. W. (1996). View specificity in object processing: Evidence from picture matching. *Journal of Experimental Psychology: Human Perception and Performance*, *22*, 395–416.

Martin, A., Ungerleider, L. G., & Haxby, J. V. (1999). Category specificity and the brain: The sensory/motor model of semantic representations of objects. In M. Gazzaniga (Ed.), *The new cognitive neuroscience* (2nd ed.). Cambridge, MA: MIT Press.

McRae, K., & Cree, G. (2002). Factors underlying category-specific semantics deficits. In E. M. E. Forde & G. Humphreys (Eds.), *Category-specificity in mind and brain*. Hove, UK: Psychology Press.

Mumford, D. (1992). On the computational architecture of the neo-cortex: II. The role of the cortico-cortical loop. *Biological Cybernetics*, *66*, 241–251.

Nososfky, R. M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General*, *115*, 39–61.

Perrett, D. I., Oram, M. W., & Ashbridge, E. (1998). Evidence accumulation in cell populations responsive to faces: An account of generalisation of recognition without mental transformations. *Cognition*, *67*, 111–145.

Pylyshyn, Z. (1999). Is vision continuous with cognition? The case of impenetrability of visual perception. *Behavioral and Brain Sciences*, *22*, 341–423.

Riesenhuber, M., & Poggio, T. (1999). Hierarchical models of object recognition in cortex. *Nature Neuroscience*, *2*, 1019–1025.

Rosch, E., Mervis, C. B., Gray, W. D., Johnson, D. M., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology*, *8*, 382–439.

Schweizer, T. A., Dixon, M. J., Westwood, D., & Piskopos, M. (2001). Contribution of visual and semantic proximity to identification performance in a viral encephalitis patient. *Brain and Cognition*, *46*, 260–264.

Schyns, P. G., & Rodet, L. (1997) Categorization creates functional features. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *23*, 681–696.

Sigala, N., Gabbiani, F., & Logothetis, N. K. (2002). Visual categorization and object representation in monkeys and humans. *Journal of Cognitive Neuroscience*, *14*, 187–198.

Sternberg, S. (1966). High-speed scanning in human memory. *Science*, *153*, 652–654.

Tarr, M. J., & Bülthoff, H. H. (1998). Image-based object recognition in man, monkey, and machine. *Cognition*, *67*, 1–20.

Tarr, M. J., & Pinker, S. (1989). Mental rotation and orientation-dependence in shape recognition. *Cognitive Psychology*, *21*, 233–282.

Tarr, M. J., Williams, P., Hayward, W. G., & Gauthier, I. (1998). *Nature Neuroscience*, *1*, 275–277.

Wiseman, S., MacLeod, C. M., & Lootsteen, P. J. (1985). Picture recognition improves with subsequent verbal information. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *11*(3), 588–595.

Wiseman, S., & Neisser, U. (1974). Perceptual organization as a determinant of visual recognition memory. *American Journal of Psychology*, *87*, 675–681.