Chapter 29

PERCEPTUAL AND SEMANTIC REORGANIZATION DURING CATEGORY LEARNING

ROBERT L. GOLDSTONE

BRIAN J. ROGOSKY

RACHEL PEVTZOW

MARK BLAIR

Indiana University

Contents

Handbook of Categorization in Cognitive Science, Edited by Henri Cohen and Claire Lefebvre

Abstract

Category learning not only depends upon perceptual and semantic representations; it also leads to the generation of these representations. We describe two series of experiments that demonstrate how categorization experience alters, rather than simply uses, descriptions of objects. In the first series, participants first learned to categorize objects on the basis of particular sets of line segments. Subsequently, they were given a perceptual part-whole judgment task. Categorization training influenced participants' part-whole judgments, indicating that whole objects were more likely to be broken down into parts that were relevant during categorization. In the second series, correlations were created or broken between semantic features of word concepts (e.g., ferocious vs. timid, and group-oriented vs. solitary animals). The best transfer was found between category learning tasks that shared the same semantic organization of concepts. Together, the experiments support models of category learning that simultaneously create the elements of categorized objects' descriptions and associate those elements with categories.

Human concept learning clearly depends upon the descriptions we give to the objects we categorize. Our concept of (DOG) is built out of features such as "furry," "barks," "four-legged," "domesticated," "friendly," and "loyal." However, recent research has found that the dependency works both ways. People's object representations not only influence, but are influenced by, the concepts that they learn. We have been exploring the psychological mechanisms by which concepts and descriptions mutually influence one another, and building computational models to show that the circle of influences is benign rather than vicious. Our efforts are not solitary. There is a growing body of behavioral [Shiffrin and Lightfoot (1997), Gauthier et al. (1998), Livingston, Andrews and Harnad (1998)], developmental [Needham (1999)], neural [Kaas (1991), Gauthier and Tarr (1997), Sigala, Gabbiani and Logothetis (2002) Gauthier et al. (2003)] and computational [Rumelhart and Zipser (1985), Hofstadter and Mitchell (1994), Harnad, Hanson and Lubin (1995), Behrmann, Zemel and Mozer (1998), Palmeri, Wong and Gauthier (2004)] evidence suggesting that it is necessary and desirable to develop categories and descriptions for objects simultaneously.

In the [DOG] example above, we purposefully merged what might be thought to be two different kinds of descriptions – perceptual and semantic. We aim to develop a unified account of perceptual and semantic reorganization that accompanies category learning. This is consistent with our larger effort to reunite perceptual and conceptual processes [Goldstone (1994), Goldstone and Barsalou (1998)]. In what follows, we describe two series of experiments implicating category learning in representational reorganization. The first series focuses on a case of perceptual reorganization, while the second focuses on semantic reorganization. However, similar mechanisms are likely to underlie both kinds of reorganization, encouraging the effort to unite perceptual and conceptual adaptation processes.

1. Concept learning and perception

Within traditional work on concept learning and categorization, there has been little suggestion that learned concepts influence perception. A working assumption made by many of the most influential theories of categorization [Bruner, Goodnow and Austin (1956), Medin and Schaffer (1978), Hintzman (1986)] is that objects to be categorized are described along a fixed set of features. The categorization procedure uses, but does not alter, the perceptual descriptions.

However, recently a number of researchers have argued that in many situations, the categorization process influences the featural descriptions that are used. Rather than viewing the "vocabulary" of primitives as fixed by low-level processes, this view maintains that the vocabulary is dependent on the higher-level processes that use the vocabulary. Some evidence for this comes from the study of expert/novice differences. Evidence suggests that experts perceive structures in X-rays [Lesgold et al. (1988), Norman et al. (1992), Sowden, Davies and Roling (2000)], beers [Peron and Allen (1988)], and infant chickens [Biederman and Shiffrar (1987)] that are missed by novices. Experts in these fields seem to acquire new ways of perceptually structuring objects as they learn new concepts.

1.1. Object segmentation

Objects often have more than one possible segmentation. The letter X can be viewed as composed of two crossing diagonal lines, or as a V and an upside-down V that just touch at their vertices. Segmenting objects into parts is an important part of the process of object recognition [Hoffman and Richards (1984), Hummel and Biederman (1992)]. Stephen Palmer (1977) argued that some segmentations of an object into parts are psychologically more natural than others. He developed a set of measures for determining the naturalness of a particular segmentation of an object. In one measure, Palmer assumed that the longer it took participants to verify whether a particular part was contained in an object, the less natural was the part. For example, in Figure 1, participants saw the whole object on the left and one of the four parts on the right. Participants would generally take longer to respond that the unnatural parts belonged to the whole than that the natural parts did. In general, Palmer's different measures of segmentation naturalness closely converged. Parts that were natural according to one measure were usually found to be natural according to other measures as well. Furthermore, the measures agreed well with a formal model of part naturalness that integrates several different sources of physical information. In this model, natural object parts tend to have components that are close to each other, have similar orientations, are connected to each other, and have similar lengths.

Our experiments used materials and tasks similar to those used by Palmer, and examined the possibility that information that is physically present in an object is not sufficient to determine its segmentation into parts. Rather, information about a person's categorization experience may also be necessary to determine the most natural segmentation.

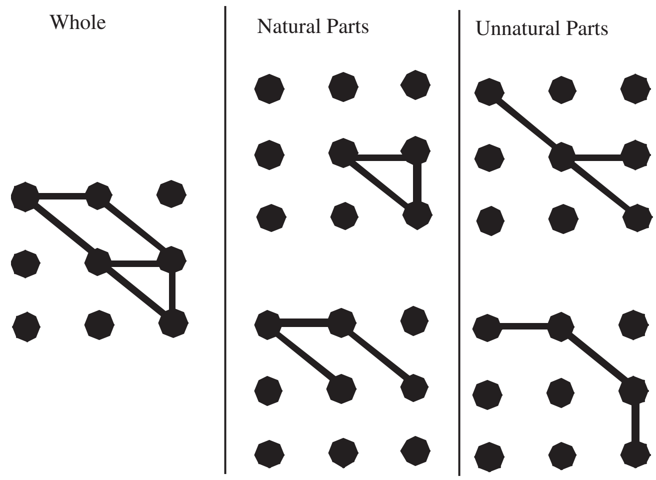

Fig. 1. The whole on the left can be segmented into either natural parts or unnatural parts.

1.2. Experiment 1

Experiment 1 tests whether categorization training can alter the naturalness of a part within a whole, as measured by part-whole response times. Participants' categorization experience is manipulated by giving them one of two different categories to learn. Both groups of participants are then given the same set of part-whole judgments.

The categorization conditions differ in the set of line segments that are diagnostic for categorization. The stimuli to be categorized are distorted versions of Objects A, B, C, and D in Figure 2. For one group of participants, A and B are placed in one category, and C and D are placed in another category. For this group of participants, the three line segments that comprise Part E and the three line segments that comprise Part F are diagnostic for categorization. Objects that belong in one category all have Part E, and objects that belong to the other category all have Part F. For the second group of participants, A and C are placed in one category, and B and D are placed in another category. For these participants, Parts G and H are diagnostic for categorization.

Categorization training could influence later part-whole judgments by highlighting segmentations of whole objects that involve diagnostic parts. For example, if Part F in Figure 2 was diagnostic during categorization training, then participants may be able to decide relatively quickly that Part F is contained in the whole object in Figure 1, even though it would be considered by Palmer's quantitative model of part goodness to be relatively unnatural. Experiment 1 tests for the influence of categorization training by comparing the part-whole judgments involving Parts E and F in Figure 2 to those involving G and H, as a function of the categorization training condition.

In Experiment 1, category parts and complements of those category parts are tested. A category part is defined as one of the sets of three line segments that were used to

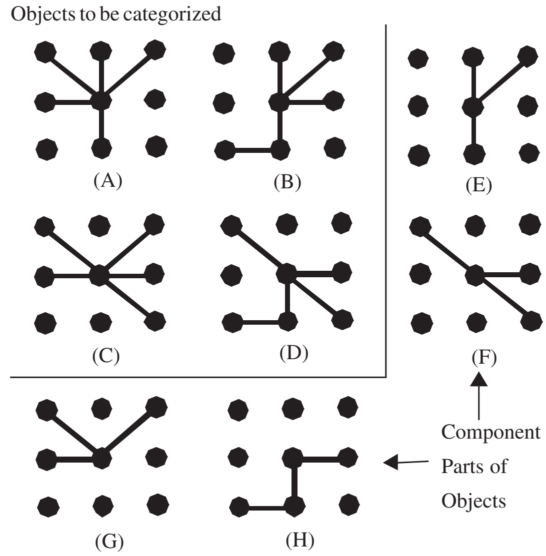

Fig. 2. Materials used in the categorization portion of Experiment 1. The four objects A, B, C, and D are categorized into two groups. Four other objects (not shown) are categorized into a third "junk" group. When A and B are placed in one group, and C and D are placed in the other, Parts E and F are diagnostic for the categorization. When A and C are placed in one group, and B and D are placed in the other, then Parts G and H are diagnostic.

construct the four objects to be categorized in Figure 2. Parts E, F, G, and H are all category parts. Category parts can either be diagnostic (if they are relevant for the categorization) or nondiagnostic. The complement of a part is defined as the line segments that remain after the category parts are removed from a whole. Figure 3 shows the four possible types of trials. On "Present Category Probe" trials, participants are probed with a category part that is present in the whole. On "Absent Category Probe" trials, participants are probed with a category part that is not present in the whole. On "Present Complement" trials participants are probed with a complement (all of the line segments except those belonging to the category part) that is present. On "Absent Complement" trials, a randomly chosen complement to a category part within another whole is used as a probe.

1.2.1. Method

There were two tasks in the experiment: categorization and whole-part decisions. In the categorization phase of the experiment, 49 participants were shown distortions of Objects A, B, C, and D as shown in Figure 2. Distortions of these objects were created by adding one line segment at a random location so that it was connected to at least one other line. Participants were asked to categorize an object into one of three groups.

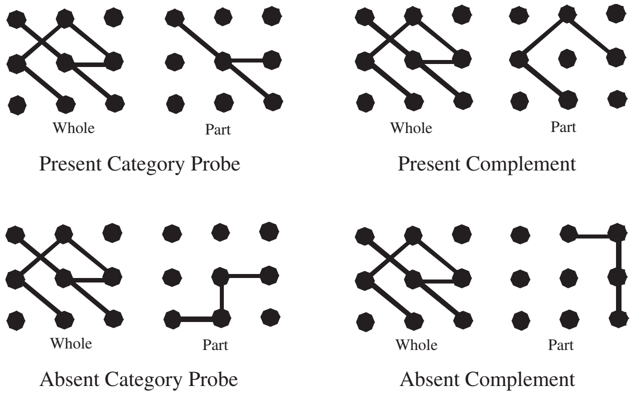

Fig. 3. The four types of possible trials in Experiment 1.

Following the response, a check mark was displayed if the participant was correct, or an X appeared if the participant was incorrect. In one categorization condition, Objects A and B belonged to one category and Objects C and D belonged to the other category. In the other categorization condition, Objects A and C belonged to one category, and Objects B and D belonged to the other category.

In the second phase of the experiment, trials consisted of displays with "wholes" and "probes." The participants' task was to decide whether the probed part was contained in the whole. The wholes consisted of one of the four category-defining parts (E, F, G, or H) from the categorization task, plus three connected line segments (complements), which were connected to the category part. The complements had no lines overlapping any of the category parts. The probes were either category parts (nondiagnostic or diagnostic) or complements.

There were four types of trials in the whole-part decomposition task: present category probe, absent category probe, present complement, and absent complement. For each of the trials shown in Figure 3, the object on the left is the whole, and the object on the right is the probe. In the first type of trial, the probe is a category part that is contained within the whole object. In the second, the probe is the complement to the category part. In the absent category probe trials, the probe is a category part, but is not contained within the whole object. For the last type of trial, absent complement, the probe is a randomly chosen complement from another object. Wholes were presented alone for 1000 ms, and then a probe was added to the display. The participants' task was to decide, as quickly and accurately as possible, whether or not the whole contained the part.

1.2.2. Results and discussion

Figure 4 shows the mean response times to decide whether or not the part was present in the whole, as a function of whether or not the whole contained a diagnostic category

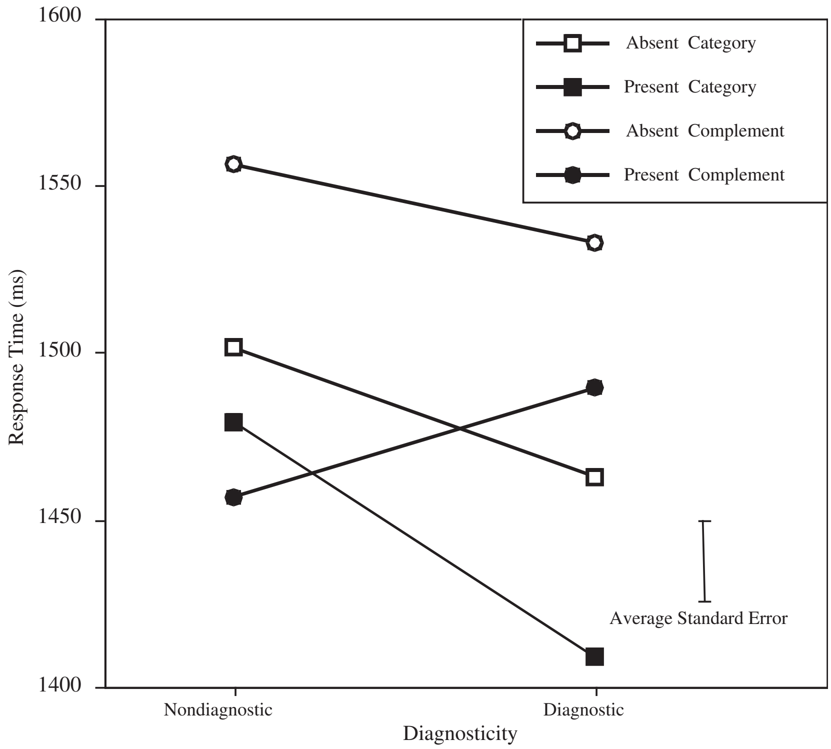

Fig. 4. Results of Experiment 1. Line segments were more readily identified as present in whole objects when they were diagnostic during categorization training than when they were nondiagnostic.

part. Response times to respond to category parts were faster for wholes that contained a diagnostic category part than for those that contained a nondiagnostic part. This diagnosticity advantage was significant only for present category parts. For complements, responses were faster for present than absent complements.

The results indicate that category learning influences perceptual sensitivity. Participants were more sensitive at responding to parts within whole objects when those parts were diagnostic. "Present" response times were significantly lower for diagnostic than nondiagnostic parts, and "absent" response times tended (nonsignificantly) to be lower as well.

1.3. Experiment 2

Experiment 2 further explores the hypothesis that category learning alters the subsequent segmentation of objects into parts. Experiment 2 introduces a new control for category parts: mirror-image reflections of category parts. During the part-whole judgment

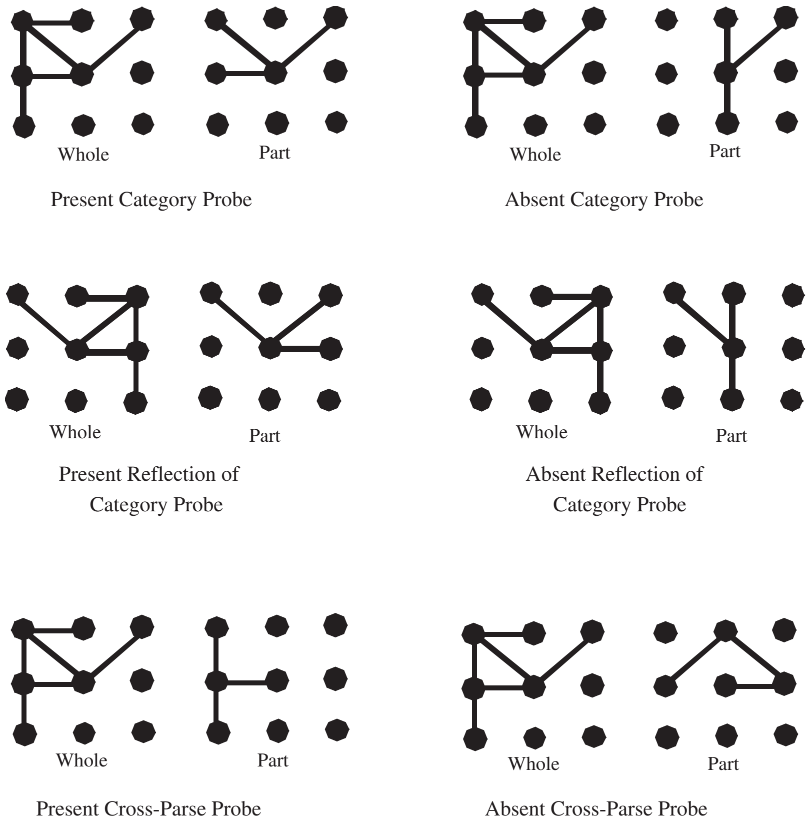

Fig. 5. Sample trials used in Experiment 2. When reflected parts were tested, the whole objects were also reflected. Cross-parse trials involve segmentations of a whole object that are incompatible with the segmentation suggested by the part presented during categorization.

task, participants were presented with category parts in some trials and reflections of category parts in other trials. Figure 5 shows six types of trials that were used. On "Present Category Part" trials, participants were presented with wholes that contained parts that were either diagnostic or nondiagnostic during categorization. On "Present Reflection of Category Part" trials, participants were presented with wholes and parts that were horizontal reflections (mirror images) of the category part trials. Finally, other parts were also tested that were neither category parts nor reflections of category parts.

Reflections of category parts are useful controls because the naturalness of a part within a whole remains invariant under reflection in Palmer's (1977) model of part goodness. For example, whatever the naturalness of Part P is in Whole W, Palmer's model predicts that the reflection of P should have the same naturalness in the reflection of W. Palmer's features for naturalness (e.g., cohesion, similarity, and proximity of line segments) remain unchanged if both the whole and the part are rotated or reflected in the same manner. In Figure 5, the "Present Category Probe" and "Present Reflection of Category Probe" conditions are predicted by Palmer's model to be equally difficult.

However, if category learning can alter the way in which an object is segmented, then it should be possible to change the quality of a part within a whole without much change to the quality of the part's reflection within the whole's reflection. If the top part of Figure 5 is diagnostic for categorization, then participants may be able to decide relatively quickly that the top whole contains this part.

1.3.1. Method

The procedure for Experiment 2 was similar to that of Experiment 1. Fifty-seven participants under the experimental conditions were given categorization training followed by a part-whole judgment task. The categorization training was identical to that of the first experiment, using the same stimuli (Figure 2). The two experimental conditions were identical to the two groups in the first experiment. A third group of 38 participants served as controls and received no categorization training.

The part-whole judgment task differed only slightly from that in Experiment 1. A new condition was added, in which the whole and the part were reflected. In addition, the parsing of an object was different: the probe was either a category part or a "cross-parse" part. Figure 5 shows examples of the different types of probes and trials. In the top two examples in Figure 5, the probe is a category part. These types of trials are identical to their comparable trials in Experiment 1. The middle two trials are similar to the top trials, in that the probes are category parts. However, unlike the top trials, the whole and the probe have been reflected (i.e., flipped horizontally). The last two examples of trials are present and absent cross-parse probes. For "present cross-parse" trials, the parsing of the whole into the cross-parse part is incompatible with the parsing required for "present category part" trials. The cross-parse cuts across the parsing needed to identify the category part, because the cross-parse part has an overlapping line segment in common with the category part. When a cross-parse probe is present, it shares a line with the category part contained within the whole. Absent cross-parse probes do not share any lines with the category part contained within the whole object; rather, they share a common line with one of the category parts that is not present within the whole. While complement parts (Experiment 1) were the remains of the whole after a category part was removed, the cross-parse parts used in Experiment 2 shared one line in common with the category part.

In this experiment, there were five factors of interest: type of probe (category or cross-parse), diagnosticity of the category part contained within the whole (diagnostic or nondiagnostic), diagnosticity of probe (diagnostic or nondiagnostic), trial type (present or absent), and reflection (normal or reflected stimuli). The two values along each of the factors occurred with equal frequency.

1.3.2. Results and discussion

Figure 6 shows the mean response times to decide whether or not the part was present in the whole, as a function of whether or not the whole contained a diagnostic category part. The baseline response times obtained from the control (no categorization) participants for the different types of probes were subtracted from the other conditions. By

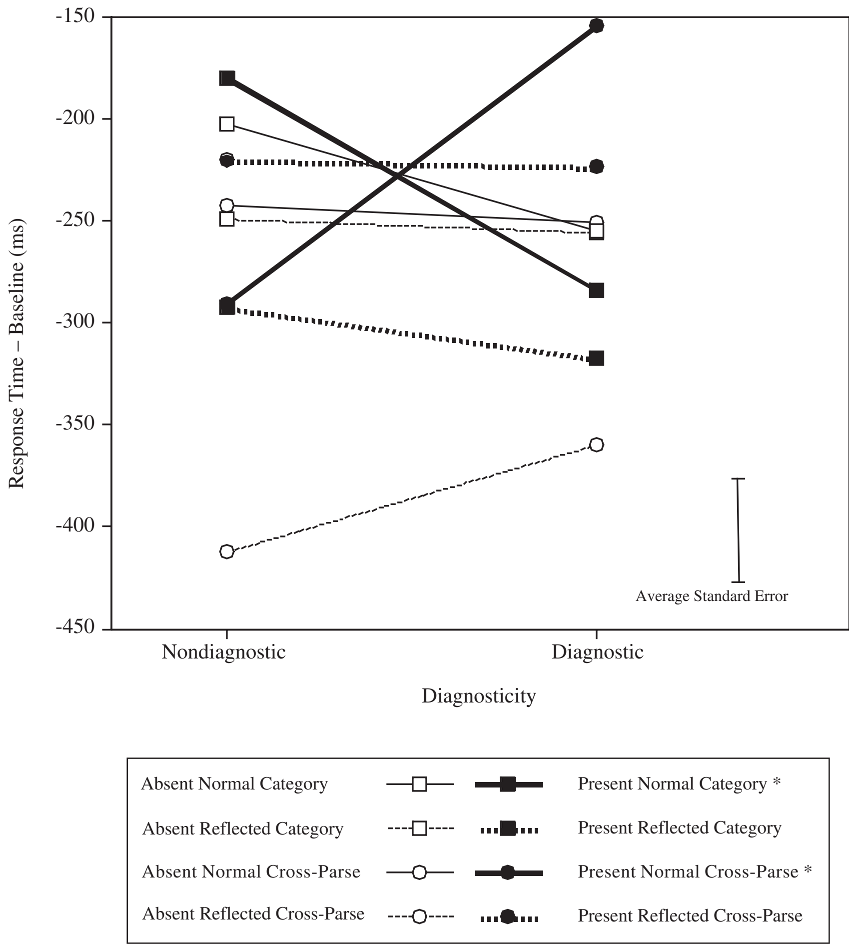

Fig. 6. Results from Experiment 2. Category parts were more quickly identified as present in whole objects when they were diagnostic during categorization training. Conversely, cross-parse parts were more quickly identified as present in whole objects when the whole objects contained a part that was nondiagnostic during categorization. Asterisks denote significant effects of diagnosticity.

subtracting out this baseline, we control for differences between the category parts and cross-parse parts in terms of intrinsic naturalness. The response times in Figure 6 are negative because the control group generally took longer to respond than the categorization groups. Thus, lower negative numbers are associated with greater advantages over the control group. Considering only trials in which the probe was a normal category part, participants were faster to respond "present" when the whole contained a diagnostic category part than when it contained a nondiagnostic part. This result replicates Experiment 1, in which participants were faster to respond to diagnostic than nondiagnostic category probes.

Diagnosticity had a significant effect on present, normal cross-parse probes. For this type of probe, response times were slower when the whole object contained a diagnostic category part than when it contained a nondiagnostic part. This is the opposite of the effect found for category probes, for which response times decreased for wholes containing diagnostic compared to nondiagnostic category parts.

Reflecting the stimuli had an effect on present, nondiagnostic category parts and cross-parse probes. When the whole contained a nondiagnostic category part, times to respond "present" were slower for normal category probes than for reflected ones. For both diagnostic and nondiagnostic absent cross-parse probes, response times were lower when the stimuli were reversed than when they were normal.

Categorization training had reliable effects on subsequent part-whole judgments, consistent with the position that participants tend to segment objects into parts that have been useful during categorization. The most straightforward effect of diagnosticity is on trials where a normal (not reflected) category part is present in the whole object, and participants are probed with this category part. On these trials, if the part was diagnostic during categorization, participants are faster to respond than if it was nondiagnostic.

The positive influence of diagnosticity of category parts was not found for horizontal reflections (mirror images) of the category parts. This result indicates a lack of transfer from learning about one part to other similar parts. A part and its reflection share commonly posited emergent features such as closure, angularity, length, height/width ratio, and density. The lack of transfer to reflected parts suggests that categorization learning sensitizes the particular three-line segments that are diagnostic rather than general stimulus properties of the diagnostic parts.

The second influence of categorization training was that, if a whole object contained a diagnostic part, then responses to present noncategory parts were slowed. In other words, on some trials, a whole object contained both a part that was relevant during categorization training and an additional part that had never been seen during categorization. If participants were probed with the never-before-seen part, they were relatively slow to respond "present." The critical aspect of the stimulus design that may explain this result is that category parts and cross-parse parts always shared one line segment. For example, in Figure 5, the category part in the top panel and the cross-parse part in the bottom panel have one line segment in common. Consequently, any segmentation that involved the category part was incompatible with the segmentation that involved the cross-parse part. If category learning biased participants to see the whole as containing the category part, then we would expect other, inconsistent segmentations of the object to be inhibited. Even though diagnosticity has a harmful influence on part-whole judgments involving new parts, this effect is consistent with the positive influence of diagnosticity. In short, object segmentations that are consistent with previously learned parts are facilitated, and those that are inconsistent with previously learned parts are inhibited.

1.4. Conclusions on perceptual reorganization

These two experiments are generally consistent in indicating that concept learning influences later perceptual part-whole judgments. Participants were more quickly able to identify parts as being present in an object when they were relevant, rather than irrelevant, for an earlier categorization task. This effect could not be explained by a bias to respond "present" because "absent" responses were never slowed, and were sometimes facilitated, for diagnostic parts. The pattern of results in general suggests that the manner in which an object is segmented into parts depends on the learned informativeness of the parts.

On the basis of our results, we can ask whether categorization training improves the response to previously relevant parts, or impedes the processing of irrelevant parts. Evidence in favor of both processes was found in the experiments. In favor of training having a positive effect on relevant parts, it was found in Experiment 1 that relevant parts were identified as present or absent more quickly than either irrelevant parts, or complements of relevant parts.

There is also convincing evidence that training causes irrelevant parts to be ignored or rejected. In Experiment 1, participants were quicker to respond "absent" when a nondiagnostic feature was present in the whole object than when a complement was present. A similar bias to respond "absent" quickly was found in Experiment 2 when comparing nondiagnostic normal category parts to reflections of these same parts. Even more persuasive evidence that irrelevant features are processed less effectively comes from the comparison of normal parts and their reflections in Experiment 2. In Experiment 2, both "present" and "absent" judgments were slow for nondiagnostic parts relative to reflections of those parts. "Present" and "absent" judgments for diagnostic parts were roughly equal in speed to judgments about reflections of diagnostic parts. Thus, by comparing judgments to their reflected controls, it becomes clear that one influence of categorization training is to desensitize irrelevant parts.

This desensitization of irrelevant parts is particularly surprising because it requires that the items should not be simply interpreted in terms of their diagnostic parts. Rather, the nondiagnostic parts must also be registered at some level in order for it to be inhibited. Although parsings of items into nondiagnostic and diagnostic parts are mutually inconsistent because they involve overlapping line segments, participants seem to generate both parsings. Rather than simply being ignored, nondiagnostic information seems to be actively suppressed. This conclusion is consistent with recent results showing that alternative figure-ground interpretations of a display compete against one another [Peterson and Lampignano (2003)].

Our current results complement other related studies showing the influences of category learning on the segmentation of objects. Hock, Webb and Cavedo (1987) showed that category learning increased the likelihood of segmenting a pattern into parts that were similar for patterns that were members of the same category. Finally, researchers have shown that participants' ability to perform a figure-ground segmentation depends on their familiarity with the stimuli [Peterson and Gibson (1994), Vecera and O'Reilly (1998), Peterson and Lampignano (2003), Vecera et al. (2004)]. People's lifelong familiarity with objects facilitates their ability to extract these objects from surrounding context and treat them as figures [Schyns and Murphy (1994)].

If our results are best explained by postulating that people create perceptual units for often-repeated patterns that are useful for categorization, one question that remains is, "How are these new units acquired?" Some researchers [Shiffrin and Lightfoot (1997), Goldstone (2000)] refer to a process of perceptual unitization by which conjunctions of stimulus features are "chunked" together so that they become perceived as a single whole unit. Simple co-occurrence of line segments is not sufficient for their unitization; nondiagnostic and diagnostic parts occur equally often during categorization. Within this framework, the sensitization of diagnostic over nondiagnostic features must be due to a unitization process that depends on categorical relevance as well as co-occurrence of features.

Mozer et al. (1992) have developed a connectionist model that learns how to segment objects. Their MAGIC system learns how to group features based on a set of presegmented examples. Object parts that belong to the same segment are represented in MAGIC by units that have the same phase of activation (they are firing in synchrony). Our experiments provide support for MAGIC's flexible, rather than fixed, segmentation procedure. Mozer (1994) added a learning principle to MAGIC that does not require explicit feedback to be provided about part segmentations. In this new model, objects tend to be segmented into parts that are uniform across instances. According to his regularity principle, features within a natural part tend to have higher correlations in their structures than do features from different parts [for a similar principle, see Schyns and Murphy (1994)]. This newer approach shows even more promise of being able to account for our results because our categorization training does not provide explicit feedback about what segments should be bound together, but it does provide information about co-occurrence relations between line segments. Again, in order to account for our experiments, this model would have to incorporate information about the categorization of objects, and not just relations between features within an object.

Goldstone (2003) presents a model of unitization, and the complementary process of differentiation, that does take into account the categorization of objects as well as unsupervised statistics across the entire set of objects. It possesses units that intervene between inputs and category outputs and can be interpreted as learned feature detectors. The CPLUS model is given a set of pictures as inputs, and produces a categorization of each picture as output. Along the way to this categorization, the model comes up with a description of how the picture is segmented into pieces. The segmentation that CPLUS creates will tend to involve parts that (1) obey the Gestalt laws of perceptual organization by connecting object parts that have similar locations and orientations, (2) occur frequently in the set of presented pictures, and (3) are diagnostic for the categorization. The network builds detectors at the same time as it builds connections between the detectors and categories. The psychological implication is that our perceptual systems do not have to be set in place before we start to use them. The concepts we need can and should influence the perceptual units we create.

2. Semantic reorganization during category learning

Several models of object perception have assumed that we recognize objects by compounding primitive elements such as features [Treisman and Gelade (1980)] or shapes [Biederman (1987)]. Likewise, many theories of conceptual representation have also been based on a fixed set of primitive semantic concepts [Schank (1972), Wierzbicka (1992)]. Just as we have favored approaches with adaptive perceptual elements, we have been led by our research to conclude that conceptual elements are similarly adaptive.

2.1. Integral versus separable dimensions

Our second line of research explores the flexibility of conceptual dimensions as they apply to classification. There has been a long history of research into how pairs of dimensions are processed, starting with Garner (1974, 1976) and Monahan and Lockhead (1977). Garner made the distinction between separable dimensions, for which one dimension can be attended to while the other is ignored, and integral dimensions, for which such selective attention is impossible. This distinction was based on patterns of results in classification tasks developed by Garner (1974). In the "correlated" task, values on both dimensions were varied together to form the stimulus set. For example, if the dimensions were size and shape of figures, then the correlated task would consist of large squares in one category and small circles in the other category. In the orthogonal ("filter") task, the categorization rule depends on only one dimension and the other, irrelevant, dimension must be ignored. For example, figures might be categorized based on size (large vs. small) regardless of their shape (square vs. circle). Performance on these tasks was compared to a univariate ("control") task in which the stimuli were categorized on a single dimension with no variation in the irrelevant dimension.

In these tasks, one of two patterns often emerged for a given pair of dimensions. For integral dimensions (e.g., saturation and brightness), the correlated task was performed better than the control task, and the filter task was performed worse than the control task. For separable dimensions (e.g., size and brightness), the correlated and filter task performances were approximately equal to the performance on the control task. The degree of integrality of the stimuli was judged according to the amount of facilitation in the correlated task and the amount of interference by the irrelevant dimension in the filter task, as compared to the control task. The interference of the irrelevant dimension can be understood as the result of an inability to selectively attend to the relevant dimension.

Likewise, the benefit of the redundant information in the correlated task could be due to both dimensions being used to perform the task, even though only one dimension is logically necessary. Monahan and Lockhead (1977) proposed that stimuli consisting of integral dimensions are initially processed in terms of overall similarity and then in terms of individual aspects. The reverse may be true for separable dimensions.

King, Gruenwald and Lockhead (1978) studied performance on the Garner classification tasks for animal terms based on the dimensions of size and ferocity. They found that the correlated task was performed better than the control task, which was performed better than the filter task. They interpreted the pattern of results as an indication of integral dimensions.

2.2. Experiment 3

To investigate the effects of category training on the integrality of semantic dimensions such as those used by King et al. (1978), we used a training-transfer paradigm using the correlated, filter, and conjunctive classification tasks. As Figure 7 shows, in the correlated task, either dimension can be used to perform the classification, or both can. In the filter task, only one dimension is relevant and the other dimension is irrelevant. In the conjunctive task, both dimensions are necessary. We hypothesized that the correlated task would induce more integral processing of the semantic dimensions since a conjunction of values indicates the category membership and this should facilitate the use of both dimensions as a unified single dimension. The conjunctive task should also induce a more fused representation of the two dimensions since both dimension values must be attended to in order to make a category choice. In the filter task, only one dimension is relevant to the categorization, so we hypothesized that this should induce a more separate use of the two dimensions. We measured these effects by training participants on one task (correlated, filter, or conjunctive) and then transferring them to a different task (correlated, filter, or conjunctive). If category training can affect the integrality of semantic dimensions, then positive transfer should occur if participants are trained on an integrating task (correlated or conjunctive) and then transferred to the other integrating task. Negative transfer should occur if participants are trained on an integrating task (correlated or conjunctive) and then transferred to the separating task (filter) or vice versa.

2.2.1. Method

Three word sets of 40 words each were designed from the categories of animals, vehicles, and clothing. For each word set, two dimensions were used. Two values were designated for each dimension, and the two dimensions were crossed, resulting in four cells with ten words in each cell (see Table 1 for the vehicles example). The dimensions were ferocity and sociability, capacity and speed, and warmth and casualness for the animal, vehicle, and clothing word sets, respectively.

One hundred and sixteen participants performed a training task followed by a testing task for each of the three word sets. Each of the three tasks in training was paired

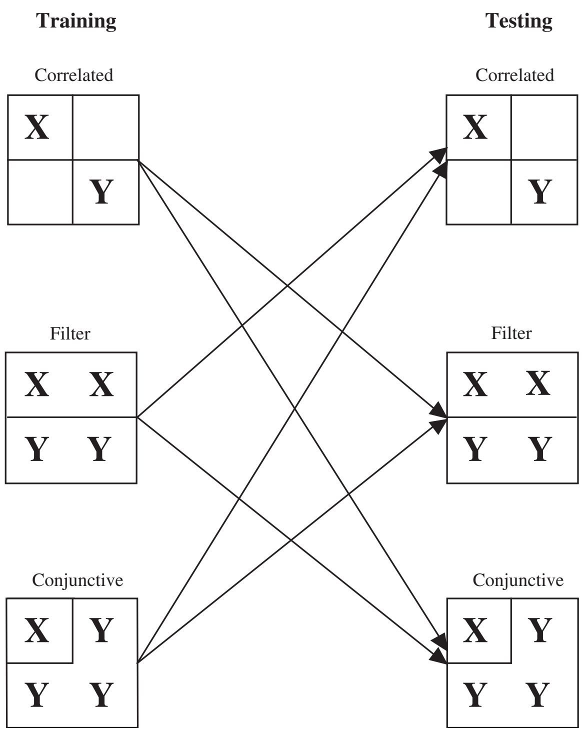

Fig. 7. The design of the training (left column) and testing tasks (right column) used in Experiments 3 and 4. Participants were transferred to a task that differed from the training task. This results in six possible train-test combinations. Each participant was given a different combination for each of the three word sets.

Table 1
'Vehicle' word set stimuli for each dimension-value combination in Experiments 3 and 4

	Only a few passengers		Many passengers	
	bicycle	cart	sailboat	ferry
	carriage	rowboat	trailer	escalator
	raft	canoe	yacht	gondola
Slow	tractor	dogsled	riverboat	balloon
	wagon	skateboard	elevator	barge
	pickup	biplane	bus	subway
	car	tank	streetcar	submarine
Fast	taxi	speedboat	van	transport
	jeep	helicopter	train	trolley
	motorcycle	rocket ship	airline	blimp

Table 2
Category descriptions for each task in Experiment 3

Task	Category X	Category Y
Correlated	Vehicles that are capable of having only a few passengers and slow	Vehicles that are capable of having many passengers and fast
Filter	Vehicles that are slow	Vehicles that are fast
Conjunctive	Vehicles that are capable of having only a few passengers and slow	Vehicles that are not both capable of having only a few passengers and slow

with one of the two different tasks in testing, resulting in six training-testing conditions (see Figure 7). In the correlated task, the categorization rule was based on the combination of two values on the dimensions that varied together. Words from two diagonally positioned cells were shown, but not from the other two cells along the reverse diagonal. In the filter task, the categorization rule was based on a single dimension that divided the set into two categories with two cells in each category. In the conjunctive task, the categorization rule was based on the combination of two dimensions for Category X (one cell), while the remaining three cells formed Category Y. Table 2 shows the categorization rules for the set of vehicle words. Before each task, participants were told the general category (e.g., vehicles), the rule for both categories (e.g., Table 2), and the list of words for each category listed in columns (e.g., Table 1).

The participants were given 160 trials in each of the three training tasks and 120 trials in each testing task. The training tasks were divided into four blocks of 40 trials each. The testing tasks were divided into three blocks of 40 trials each. For each block, the words were selected with equal frequency from each cell and presented in a randomized order.

On each trial, the word was presented on the computer screen with the first letter at the center of the screen. Participants made their category choice using the number keys. They were given feedback on their choice using a check mark for correct answers and an X for incorrect answers.

2.2.2. *Results and discussion*

The average response time results for correct trials are shown in Figure 8. During training, the correlated and conjunctive tasks were both performed significantly faster than the filter task. The correlated testing task was not performed significantly differently based on the training task that preceded it. Performance on the conjunctive testing task was significantly more accurate when it was preceded by the correlated training as compared to the filter training.

The performance during training provides a baseline to which we can compare the relative effects of training on that task. The correlated training had no significant effects on the conjunctive testing task compared to initial conjunctive training performance. The filter training had a significant negative effect on the conjunctive testing task compared

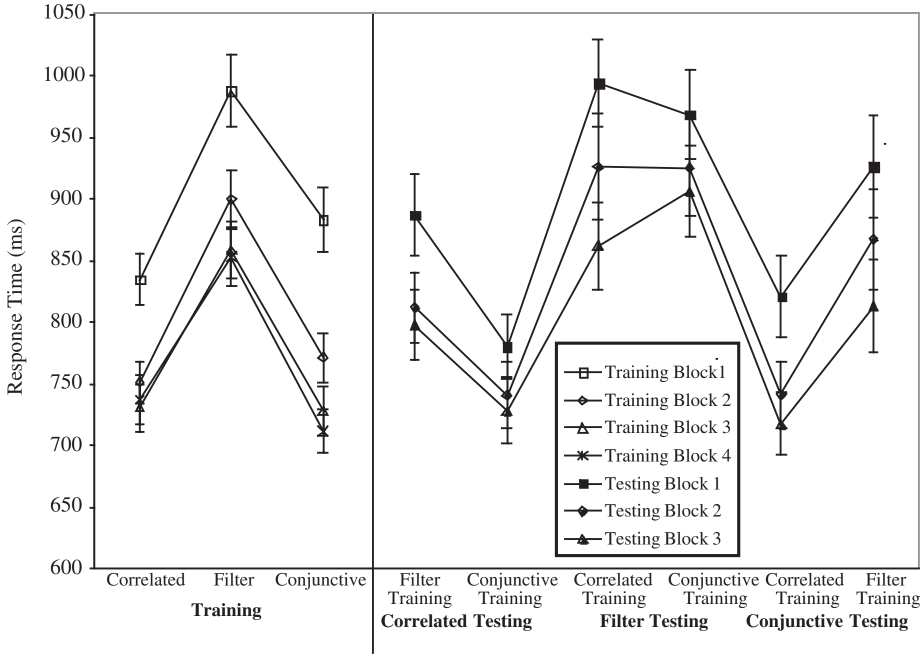

Fig. 8. Response time data from Experiment 3 for the initial training tasks and the testing tasks (error bars ±1 S.E.).

to the initial conjunctive training performance for accuracy and response time. The filter training also had a negative effect on the correlated task compared to the correlated training task performance that was significant in terms of response time. The correlated training had a significant negative effect on accuracy of the filter testing task compared to filter training accuracy. The conjunctive training also had a significant negative effect on the accuracy of the filter testing task compared to filter training accuracy.

The filter task training resulted in negative transfer to the conjunctive task and the correlated task. The correlated task training did not have any effect on transfer to the conjunctive task. Relative to initial performance, we have evidence of negative transfer of training on the filter task for both the conjunctive and correlated tasks and negative transfer of training on the correlated and conjunctive tasks for the filter task. This matches the prediction that there would be negative transfer effects between tasks inducing separation of dimensions and those inducing fusion of dimensions.

2.3. Experiment 4

Experiment 3 showed that the effects of classification task training on subsequent testing tasks are consistent with an adaptation of the conceptual dimensions. However, two possible types of adaptation could be taking place: a change in the representation of dimensions for individual word exemplars of the category, or a change in the representation of the dimension at the category level. In Experiment 4, this question is explored using a design in which new words are introduced during the test tasks.

The design of the tasks was the same as in Experiment 3, except that new exemplars were presented in the testing phase that had not been presented in the training phase. We hypothesized that training in a task that induces processing of two semantic dimensions in an integral manner (correlated and conjunctive) will result in positive transfer to the other fusion-inducing task. Conversely, negative transfer is expected from training in the task that is thought to induce separate processing of dimensions (filter) to the fusion-inducing tasks (correlated or conjunctive) and vice versa. We also hypothesized that the same pattern of results in the testing tasks would be found for the both the novel words and the words previously seen during training although the negative transfer effects may be more pronounced for words that had previously been seen.

2.3.1. Method

The materials were similar to those used in Experiment 3, except that the number of words in each domain was doubled to 80. The three domains were animals, activities, and things. The animal words were again placed in a 2 × 2 table along the dimensions of ferocity and sociability. The activities word set consisted of sports and hobbies that were classified according to how physical the activity is ('strenuous' vs. 'light') and the riskiness of the activity ('risky' vs. 'safe'). The things word set consisted of various objects and materials that were classified according to their naturalness ('natural' vs. 'artificial') and their fluidity ('solid' vs. 'fluid'). All of the categories were pretested by having participants rate words on two dimensions and using these ratings to select words that fell most clearly into one category or the other.

Each of the 248 participants was given only one of the combinations of training and testing tasks. They repeated the particular train-test condition for each of the three word sets. All other aspects of the task design were the same as in Experiment 3 except for the use of new words in the testing phase. In the training task, only half of the available words in each cell were presented (10 words). In the testing task, all the available words for the cells used in the task were presented. The category frequency was balanced in each task and the order of the word sets and word presentations was randomly selected for each participant.

2.3.2. Results and discussion

The response times for the experiment are shown in Figure 9. The correlated testing task was performed significantly slower when preceded by the filter training than by the conjunctive training, but this effect was limited to previously trained words only. In the filter testing task, the previously trained words were performed more accurately when preceded by the correlated training than the conjunctive training. In the conjunctive

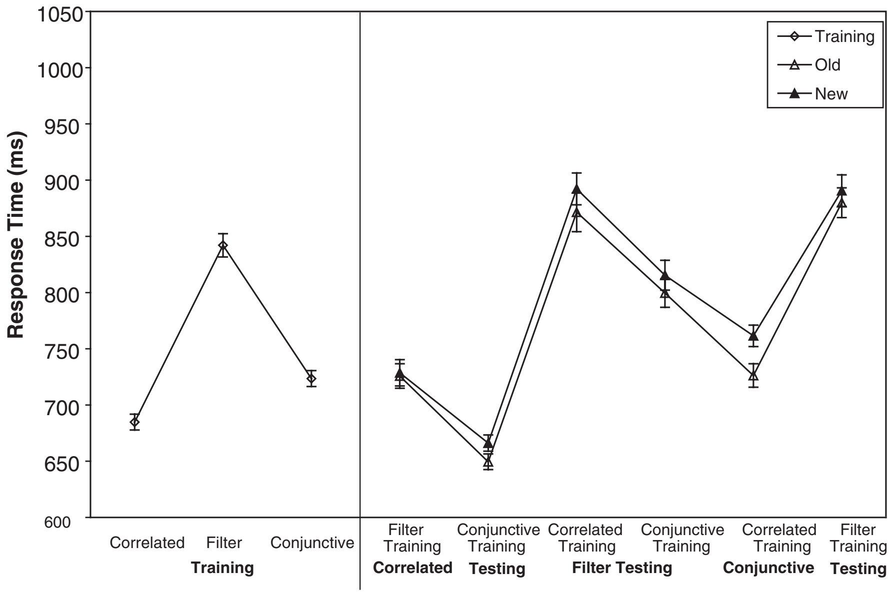

Fig. 9. Response time data from Experiment 4 for the initial training tasks and the testing tasks (error bars ±1 S.E.).

testing task, the novel words were judged faster when preceded by correlated training than filter training. Likewise, old words were categorized more accurately and faster when preceded by correlated training than filter training.

As in Experiment 3, the training task performances were compared using the average performance over each of the three blocks. And again, the correlated task was performed more accurately and faster than the conjunctive task. In turn, the conjunctive task was performed more accurately and faster than the filter task. Thus, for the initial training task performance, the same pattern of results was found as in Experiment 3: the correlated task elicited the best performance followed by the conjunctive task, followed by the filter task.

Novelty of the words during the testing task was the crucial factor tested in Experiment 4. Overall, words previously seen in training were responded to more quickly than new words and there was a larger range of improvement for the speed of response to new words than old words. These effects are not surprising since the old words presumably were made easier to classify in the testing task by the prior exposure in the training task. Therefore, more improvement is expected in the new words.

In the interactions between novelty and condition, there was little difference between the old and new words. For novel words, the effect of transfer is based on a shift in the integrality or separability of the semantic dimensions alone, and not due to a direct change in the specific item representation since the words were not present in training. The degree to which the training induced a change in the dimension representation over and above the changes to individual item representations can be measured by the degree to which the same pattern of results is found for both old and new words. The correlated testing task revealed a benefit in terms of accuracy from the conjunctive training task compared to filter training for old words, but not for new words. The conjunctive task exhibited a positive transfer effect on old words from the correlated training compared to filter training, in terms of both accuracy and response time, and on new words in terms of response time. These results suggest that changes may occur in both the item representations and the semantic dimensions.

3. Conclusions on semantic reorganization

Both of the experiments in this series obtained the same result for the initial task performances. The correlated task was performed the best, followed by the conjunctive task, and the filter task was performed least well. The fact that both the correlated and conjunctive tasks were performed better than the filter task is likely the result of the dimensions' initially being integrally processed such that they can be easily processed together but not so easily processed separately. These results echo the findings of King et al. (1978).

Negative transfer effects were obtained in the filter task following conjunctive or correlated training. Likewise, negative transfer effects were obtained in the correlated and conjunctive tasks following filter training. These effects support the hypothesis that training may induce a change in the integrality of the semantic dimensions.

Experiment 4 tested whether the adaptation occurs on an individual linguistic concept level, whereby the features of a particular item become integrated, or on a semantic dimension level, whereby changes generalize to other concepts defined by the altered dimensions. While some effects did not generalize to concepts not seen during training, correlated training had a positive effect on the conjunctive testing task relative to the filter training effects for both old and new words, suggesting that changes took place at the level of the semantic dimensions.

Experience using semantic dimensions in classification tasks can alter the processing of those dimensions. There were shifts in the apparent integrality of the dimensions such that tasks that incorporate both dimensions together may create a more fused representation of the dimensions. Other tasks that require the use of a single dimension and the discounting of an irrelevant dimension tend to cause a separated representation of the dimensions. More generally, our studies provide a methodological tool for examining how any number of semantic dimensions across domains are processed and adapted during classification tasks.

4. Integrating perceptual and semantic reorganization

Together, the four reported experiments suggest an alternative approach to theories that have posited fixed sets of perceptual [Treisman and Gelade (1980), Biederman (1987)] or semantic [Schank (1972), Wierzbicka (1992)] features. According to this alternative, category learning not only uses existing object descriptions, but also alters object descriptions to facilitate the learning. Understandably, the claim that new perceptual or semantic features are created during category learning has been controversial [Schyns, Goldstone and Thibaut (1998)], and we would like to dispel the notion that feature creation is a magical process, or that once feature creation is allowed, then "anything goes."

4.1. Characterizing psychological features

To understand what we mean by feature creation, it is helpful to first analyze what we mean by "feature." By "feature" we mean a psychological unit of perception or thought. "Dimensions" are also psychological entities, but refer to a set of values that can be ordinally positioned. Brightness, then, is a psychological dimension only because it is processed as a unit. If luminance energy were not psychologically isolated, then there would not be a (psychological) dimension of brightness reflecting this physical quantity.

If features and dimensions are units of perception and thought, then we can ask what physical aspects are bundled together into these psychological units. Features can be interpreted as packages of stimulus elements that are separated from other sets of elements and that reflect the subjective organization of the whole stimulus as a set of components. Features can be revealed using several experimental operationalizations. If two pieces of physical information, X and Y, are packaged together in the same psychological feature and Z is not, then several empirical predictions follow. We predict that searching for X and Y simultaneously should be easier than simultaneously searching for X and Z [Treisman and Gelade (1980)]. We predict that searching for X should be affected by contextual variation to Y more than to Z [Gauthier and Tarr (2002)]. And we predict that categorization based on X should be slowed more by irrelevant variation to Y than to Z [Garner (1974, 1976)]. It should be easier for people to attend to X and Y simultaneously than X and Z. All of these operationalizations tie into the notion that X and Y are being processed together.

It is also noteworthy that all of these operationalizations imply a continuum of featurehood. There will be various degrees to which stimulus aspect Y intrudes upon or facilitates processing of X. Although we conceive of features as packages of stimulus components, we are not proposing that packages are completely discrete or mutually exclusive. Rather, they are packages in the same way that cliques can be circled in social networks or regions can be identified in brain neural networks. In all three domains, a unit (feature, clique, or region) is characterized by relatively dense within-unit connectivity among elements and relatively sparse connectivity between elements within the unit and external elements. Features are useful idealizations because they capture the notion of elements that are densely interconnected, but it is important to recognize that (1) features (e.g., densely interconnected clusters) may exist at multiple levels of resolution, (2) elements processed as one feature may not have uniform interconnectivity to other elements of the same feature, and (3) the internal integrity of different features may vary.

4.2. Characterizing featural change

Having characterized psychological features, we can now turn to the meaning of feature creation. By this account, feature creation simply involves alterations to the organization of stimulus elements into features. Figure 10 shows two ways that this can happen. By unitization, stimulus elements (circles) that were originally processed as three different features (ovals) come, with practice, to be processed as only two features. Elements that were originally processed separately are now processed together [Shiffrin and Lightfoot (1997), Goldstone (2000)]. By differentiation, the same three-element object comes to be processed as four features. Elements that were originally psychologically fused together become isolated [Smith and Kemler (1978), Smith, Gasser and Sandhofer (1997), Goldstone and Steyvers (2001)].

From Figure 10 it may appear like there are two separate, perhaps contradictory, tracks for featural change. In fact, not only are unitization and differentiation compatible with each other, but they often occur simultaneously. They are compatible because both processes created appropriate-sized units for a task. If elements covary together and their co-occurrence predicts an important categorization, then the elements will tend to be unitized. If elements vary independently of one another and they are differentially relevant for categorizations, then the elements will tend to be differentiated. Experiments 1 and 2 are good examples of simultaneous unitization and differentiation. During category learning, the three line segments that jointly indicate a category are unitized together, and are isolated from other line segments in the objects to be categorized. Accordingly, we do not support theories that propose monolithic developmental trends toward either increasingly unitized [Gauthier and Tarr (2002)] or differentiated [Kemler and Smith 1978)] representations. We believe that both occur, and furthermore, that the same learning algorithm can do both simultaneously [Goldstone (2003)].

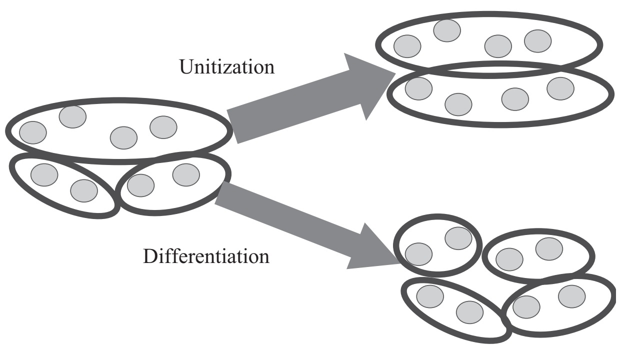

Fig. 10. Two varieties of featural reorganization. Stimulus elements are shown by circles and psychological packages of those elements – in features – are shown by ovals. By unitization, stimulus elements that were once processed as different features come to be processed as a single feature. By differentiation, stimulus elements that were once processed as the same feature come to be processed by multiple features.

Features are not created "out of nothing." They are reorganizations of stimulus elements. A critic might respond, "Then how is your account any different from the standard fixed-features approach in which primitive elements are combined in new arrangements to create object representations?" For now, we will give three replies [see Schyns et al. (1998) for others]. First, by our account, features are not (always) created from a set of psychological primitives. Often they are created from stimulus elements that originally have no parsing in terms of psychological primitives. For example, people can create a "saturation" detector that is relatively uninfluenced by brightness even if there was originally no detector that had this response profile [Burns and Shepp (1988)]. To be sure, if brightness and saturation affected a brain identically, then there would be no way to develop a detector that responded to only one of these properties. However, as long as two properties have some differential effects, then increasingly differentiated detectors can emerge, if the training encourages their isolation. The critic might counter, "But dimensions that are fused together at some point in perceptual processing can never be split later." By analogy, once red ink has been poured into blue ink, there is no simple procedure for later isolating the blue ink. Fortunately, this analogy is misleading, and there are several computational models that can differentiate fused dimensions [Smith et al. (1997), Edelman (1999), Goldstone (2003)]. For example, competitive learning networks differentiate inputs into categories by developing specialized detectors for classes of stimuli [Rumelhart and Zipser (1985)]. Random detectors that are slightly more similar to an input than other detectors will learn to adapt themselves toward the input and will inhibit other detectors from doing so. The end result is that originally homogeneous detectors become differentiated and heterogeneous over the course of training.

Second, feature creation often involves delineating spatial regions rather than composing elements. For example, a bounded segment of a curve can be extracted by identifying its end points by rapid changes in curvature [Hoffman and Richards (1984)]. This extraction does not require piecing together elements. What would these putative elements be – line segments or pixels? There is good evidence that neither small line segments nor pixels are functionally useful features for object recognition. They are too low-level to provide diagnostic evidence for actual objects. Moreover, pixels cannot be true features because they are not identified by intrinsic attributes like 'red' or '4 cm.' Their essential nature depends upon their location in spatial media. Much of feature creation involves forming bounded regions in a spatial medium rather than symbolically composing atomic elements.

Third, there are clear-cut cases, where something like new perceptual devices are created. By becoming physically modified, systems can learn to represent properties that they were unable to represent originally. In evolutionary time, organisms developed ears sensitive to acoustic properties that no early organisms (e.g., bacteria) could detect. This is also possible within a system's own lifetime. The cybernetician Gordon Pask built a device that could create its own primitive feature detectors. It consisted of an array of electrodes partially immersed in an aqueous solution of metallic salts. Passing a current through the electrodes caused dendritic metallic threads to grow. Eventually the threads created bridges between the electrodes, which subsequently changed the behavioral repertoire of the device. Cariani (1993) reports that within a half a day, the system could grow to be sensitive to a sound or magnetic field. With more time, the device could discriminate between two musical pitches. Similarly, there is good neurophysiological evidence that training can produce changes in the early somatosensory, visual, and auditory cortex [see Goldstone (1998) for a review]. While these changes are not as radical as sprouting a new ear, they are existence proofs that early perceptual devices can be systematically and physically altered by the environment to change their representational capacities.

4.3. Prospects for synthesizing perceptual and semantic reorganization

We have juxtaposed two series of experiments with the intention of highlighting similarities and differences between the perceptual and semantic reorganization that accompanies concept learning. In the first series of experiments, people apparently create shape complexes during category learning, and use those shape complexes as building blocks for describing subsequently presented objects. In the second series, people create either fused or separated semantic descriptions that subsequently affect their later categorizations. Is the process of creating a three-line-segment complex similar to creating an integrated representation of the timidity and sociability of animals, or the speed and capacity of vehicles?

One apparent discrepancy between perceptual and semantic unit construction is that there are strong visuospatial constraints on perceptual unit creation. People have a strong bias to create units that obey Gestalt laws of proximity, similarity and good continuation. These biases are needed for computational models that aim to create psychologically plausible perceptual units [Goldstone (2003)], and are useful in limiting the combinatorial explosion of potential units that could be built. At first sight, semantic units do not have any corresponding constraints on their construction.

Upon further reflection, we believe that there are biases affecting semantic unit construction, and that these biases play a loosely analogous role to the Gestalt laws of perceptual organization. Informal interviews with some of the participants in Experiments 3 and 4 suggest that in the conjunctive and correlated conditions, participants often created conceptions that fused the two component dimensions into a semantic Gestalt. For example, for the animals category, people often created a schema for social, timid animals that consisted of groups of small animals huddled together for protection. For the vehicles category, participants sometimes created a fused notion of fast, high-capacity vehicles by imagining mass transportation systems. By this account, just as it would be difficult to create a unit for two line segments that are far apart, of different thicknesses, and not part of a continuous path, it should be difficult to create semantic units for hard-to-relate semantic dimensions such as Jorge Luis Borges' (1966) dimensions of animals: "those that tremble as if they were mad" and "those that have just broken a flower vase." Furthermore, we believe that these constraints on semantic unit construction are important for creating nontrivial units. There is a trivial sense in which any features, such 'square' and 'blue,' can be combined to create a conjunctive unit 'square and blue.' However, these conjunctions are inert, being no more than the Boolean concatenation of their elements. For these conjunctions, the standard compositional account of unit construction is perfectly adequate. However, semantic reorganization often differs from logical combination, and the elements interact to create complexes with emergent properties. Much of the recent work on knowledge-based categorization provides insight into the development of semantic complexes [Murphy (2002)].

Much of the most important work in characterizing representational reorganization will involve specifying mechanisms that are tightly tied to particular classes of materials. Still, we are sanguine about the heuristic utility of attempting to unify perceptual and semantic reorganization processes. Complementary mechanisms of differentiation and unitization are found for both. Both are guided by unsupervised statistics and supervised feedback provided by categorizations. Moreover, it may prove difficult to draw a clean dividing line between perceptual and conceptual processing [Goldstone and Barsalou (1998)], not just because we lack precise enough empirical diagnostics, but because they emanate from a shared substratum.

Acknowledgments

Many useful comments and suggestions were provided by John Kruschke, Robert Nosofsky, Linda Smith, and Richard Shiffrin. This research was funded by NIH Grant MH56871, and NSF Grant 0125287.

References

Behrmann, M., R.S. Zemel and M.C. Mozer (1998), "Object-based attention and occlusion: Evidence from normal participants and a computational model", Journal of Experimental Psychology: Human Perception and Performance 24:1011–1036.

Biederman, I. (1987), "Recognition-by-components: a theory of human image understanding", Psychological Review 94:115–147.

Biederman, I., and M.M. Shiffrar (1987), "Sexing day-old chicks: a case study and expert systems analysis of a difficult perceptual-learning task", Journal of Experimental Psychology: Learning, Memory, and Cognition 13:640–645.

Borges, J.L. (1966), Other Inquisitions (Washington Square Press, New York).

Bruner, J.S., J.J. Goodnow and G.A. Austin (1956), A Study of Thinking (Wiley, New York).

Burns, B., and B.E. Shepp (1988), "Dimensional interactions and the structure of psychological space: the representation of hue, saturation, and brightness", Perception and Psychophysics 43:494–507.

Cariani, P. (1993), "To evolve an ear", Systems Research 10:19–33.

Edelman, S. (1999), Representation and Recognition in Vision (MIT Press, Cambridge, MA).

Garner, W.R. (1974), The Processing of Information and Structure (Lawrence Erlbaum, Potomac, MD).

Garner, W.R. (1976), "Interaction of stimulus dimensions in concept and choice processes", Cognitive Psychology 8:98–123.

Gauthier, I., T.W. James, K.M. Curby and M.J. Tarr (2003), "The influence of conceptual knowledge on visual discrimination", Cognitive Neuropsychology 3/4/5/6:507–523.

Gauthier, I., and M.J. Tarr (1997), “Becoming a “Greeble” expert: exploring mechanisms for face recognition”, Vision Research 37:1673–1682.

Gauthier, I., and M.J. Tarr (2002), “Unraveling mechanisms for expert object recognition: bridging brain activity and behavior”, Journal of Experimental Psychology: Human Perception & Performance 28:431–446.

Gauthier, I., P. Williams, M.J. Tarr and J. Tanaka (1998), “Training “greeble” experts: a framework for studying expert object recognition processes”, Vision Research 38:2401–2428.

Goldstone, R.L. (1994), “The role of similarity in categorization: providing a groundwork”, Cognition 52:125–157.

Goldstone, R.L. (1998), “Perceptual learning”, Annual Review of Psychology 49:585–612.

Goldstone, R.L. (2000), “Unitization during category learning”, Journal of Experimental Psychology: Human Perception and Performance 26:86–112.

Goldstone, R.L. (2003), “Learning to perceive while perceiving to learn”, in: R. Kimchi, M. Behrmann and C. Olson, eds., Perceptual Organization in Vision: Behavioral and Neural Perspectives (Lawrence Erlbaum Associates, Mahwah, NJ) 233–278.

Goldstone, R.L., and L. Barsalou (1998), “Reuniting perception and conception”, Cognition 65:231–262.

Goldstone, R.L., and M. Steyvers (2001), “The sensitization and differentiation of dimensions during category learning”, Journal of Experimental Psychology: General 130:116–139.

Harnad, S., S.J. Hanson, and J. Lubin (1995), “Learned categorical perception in neural nets: Implications for symbol grounding”, in: V. Honavar and L. Uhr, eds., Symbolic Processors and Connectionist Network Models in Artificial Intelligence and Cognitive Modeling: Steps Toward Principled Integration (Academic Press, Boston) 191–206.

Hintzman, D.L. (1986), ““Schema abstraction” in a multiple-trace memory model”, Psychological Review 93:411–429.

Hock, H.S., E. Webb and L.C. Cavedo (1987), “Perceptual learning in visual category acquisition”, Memory & Cognition 15:544–556.

Hoffman, D.D., and W.A. Richards (1984), “Parts of recognition”, Cognition 18:65–96.

Hofstadter, D.R., and M. Mitchell (1994), “The Copycat project: a model of mental fluidity and analogy-making”, in: K.J. Holyoak and J.A. Barnden, eds., Advances in connectionist and neural computation theory, vol. 2. (Ablex, Norwood, NJ) 31–112.

Hummel, J.E., I. Biederman (1992), “Dynamic binding in a neural network for shape recognition”, Psychological Review 99:480–517.

Kaas, J.H. (1991), “Plasticity of sensory and motor maps in adult mammals”, Annual Review of Neuroscience 14:137–167.

Kemler, D.G., and L.B. Smith (1978), “Is there a developmental trend from integrality to separability in perception?” Journal of Experimental Child Psychology 26:498–507.

King, M.C., P. Gruenwald and G.R. Lockhead (1978), “Classifying related stimuli”, Journal of Experimental Psychology: Human Learning & Memory 4:417–427.

Lesgold, A., R. Glaser, H. Rubinson, D. Klopfer, P. Feltovich and Y. Wang (1988), “Expertise in a complex skill: Diagnosing X-ray pictures”, in: M.T.H. Chi, R. Glaser and M.J. Farr, eds., The Nature of Expertise (Laurence Erlbaum Associates, Hillsdale, NJ) 315–335.

Livingston, K.R., J.K. Andrews and S. Harnad (1998), “Categorical perception effects induced by category learning”, Journal of Experimental Psychology: Learning, Memory, and Cognition 24:732–753.

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

Monahan, J.S., and G.R. Lockhead (1977), “Identification of integral stimuli”, Journal of Experimental Psychology: General 106:94–110.

Mozer, M.C. (1994), “Computational approaches to functional feature learning”, Proceedings of the Sixteenth Annual Conference of the Cognitive Science Society (Lawrence Erlbaum Associates, Hillsdale, NJ) 975–976.

Mozer, M.C., R.S. Zemel, M. Behrmann and C.K.I. Williams (1992), “Learning to segment images using dynamic feature binding”, Neural Computation 4:650–665.

Murphy, G.L. (2002), The Big Book of Concepts (MIT Press, Cambridge, MA).

Needham, A. (1999), "The role of shape in 4-month-old infants' segregation of adjacent objects", Infant Behavior and Development 22:161–178.

Norman, G.R., L.R. Brooks, C.L. Coblentz and C.J. Babcock (1992), "The correlation of feature identification and category judgments in diagnostic radiology", Special issue: Memory and cognition applied. Memory & Cognition 20:344–355.

Palmer, S.E. (1977), "Hierarchical structure in perceptual representation", Cognitive Psychology 9:441–474.

Palmeri, T.J., A. C.-N. Wong and I. Gauthier (2004), "Computational approaches to the development of perceptual expertise", Trends in Cognitive Sciences 8:378–386.

Peron, R.M., and G.L. Allen (1988), "Attempts to train novices for beer flavor discrimination: a matter of taste", The Journal of General Psychology 115:403–418.

Peterson, M.A., and B.S. Gibson (1994), "Must figure-ground organization precede object recognition? An assumption in peril", Psychological Science 5:253–259.

Peterson, M.A., and D.W. Lampignano (2003), "Implicit memory of novel figure-ground displays includes a history of cross-border competition", Journal of Experimental Psychology: Human Perception and Performance 29:808–822.

Rumelhart, D.E., and D. Zipser (1985), "Feature discovery by competitive learning", Cognitive Science 9:75–112.

Schank, R. (1972), "Conceptual dependency: a theory of natural language understanding", Cognitive Psychology 3:552–631.

Schyns, P.G., R.L. Goldstone and J. Thibaut (1998), "Development of features in object concepts", Behavioral and Brain Sciences 21:1–54.

Schyns, P.G., and G.L. Murphy (1994), "The ontogeny of part representation in object concepts", in: D.L. Medin, ed., The Psychology of Learning and Motivation, vol. 31 (Academic Press, San Diego) 305–354.

Shiffrin, R.M., and N. Lightfoot (1997), "Perceptual learning of alphanumeric-like characters", in R.L. Goldstone, P.G. Schyns and D.L. Medin, eds., The Psychology of Learning and Motivation, vol. 36 (Academic Press, San Diego) 45–82.

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

Smith, L.B., M. Gasser and C. Sandhofer (1997), "Learning to talk about the properties of objects: a network model of the development of dimensions", in: R.L. Goldstone, P.G. Schyns and D.L. Medin, eds., The Psychology of Learning and Motivation, vol. 36(Academic Press, San Diego) 219–255.

Smith, L.B., and D.G. Kemler (1978), "Levels of experienced dimensionality in children and adults", Cognitive Psychology 10:502–532.

Sowden, P.T., I.R.L. Davies and P. Roling (2000), "Perceptual learning of the detection of features in X-ray images: a functional role for improvements in adults' visual sensitivity?", Journal of Experimental Psychology: Human Perception and Performance 26:379–390.

Treisman, A., and G. Gelade (1980), 'A feature-integration theory of attention", Cognitive Psychology 12:97–136.

Vecera, S.P., A.V. Flevaris, C.K. Morewedge and T.D. Wilson (2004), "Exogenous spatial attention influences figure-ground assignment", Psychological Science 15:20–26.

Vecera, S.P., and R.C. O'Reilly (1998), "Figure-ground organization and object recognition processes: an interactive account", Journal of Experimental Psychology: Human Perception and Performance 24:441–462.

Wierzbicka, A. (1992), "Semantic primitives and semantic fields", in: A. Lehrer and E.F. Kittay, eds., Frames, fields, and contrasts: New essays in semantic and lexical organization (Lawrence Erlbaum Associates, Hillsdale, NJ) 209–228.