PEARSON NEW INTERNATIONAL EDITION

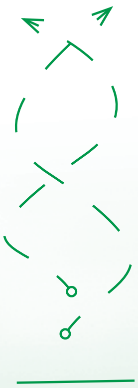

Cognitive Psychology: Mind and Brain

Edward E. Smith Stephen M. Kosslyn

First Edition

ALWAYS LEARNING

PEARSON

Pearson New International Edition

Cognitive Psychology: Mind and Brain
Edward E. Smith   Stephen M. Kosslyn
First Edition

PEARSON

processing. They reasoned that if simulations represent category knowledge, then switching costs analogous to those incurred while processing perceptual information should be incurred while processing information about categories. Participants in this study verified the properties of objects. On a given trial, the word for a category (for example, "cakes") was followed by a word for a possible property, both words presented visually. Half the time the property was true of the category ("frosting") and half the time it was false ("crust"). As in the earlier perception experiment, sometimes the properties referred to the same modality on two consecutive trials: a participant might verify that "rustles" is a property of "leaves" and on the next trial verify that "loud" is a property of "blenders." Most of the time, however, the properties across two consecutive trials referred to different modalities.

Pecher and colleagues (2003) found that switching modalities in this property verification task produced a switching cost, just as in the perception experiment by Spence and colleagues (2000). When participants had to switch modalities to verify a property, they took longer than when they did not have to switch modalities. This finding is consistent with the idea that perceptual mechanisms are used in the representation of category knowledge: to represent the properties of categories, participants appeared to simulate them in the respective modalities.

Many other behavioral findings similarly demonstrate that perceptual mechanisms play a role in the representation of category knowledge. The visual mechanisms that process occlusion, size, shape, orientation, and similarity have all been shown to affect category processing (e.g., Solomon & Barsalou, 2001, 2004; Stanfield & Zwaan, 2001; Wu & Barsalou, 2004; Zwaan et al., 2002). Motor mechanisms have also been shown to play central roles (e.g., Barsalou et al., 2003; Glenberg & Kaschak, 2002; Spivey et al., 2000). Across modalities, behavioral findings increasingly implicate modality-based representations in the storage and use of category knowledge.

3.4. Multimodal Mechanisms and Category Knowledge: Neural Evidence

When talking about modality-specific mechanisms, conclusions drawn from behavioral evidence, no matter how suggestive, have their limits: behavioral experiments don't measure brain mechanisms directly. But neuroimaging does, and much supportive evidence for the perceptual underpinnings of category knowledge comes from neuroimaging research. In these studies, participants lie in a PET or fMRI scanner while performing various category-related tasks, such as naming visually presented objects (for example, a dog), listening to the names of categories (for example, "hammer"), producing the properties of a category (for example, "yellow" for a lemon), or verifying the properties of a category (for example, answering the question "Does a horse run?").

For example, in a study by Chao and Martin (2000), participants were asked to observe pictures of manipulable objects, buildings, animals, and faces while their brains were scanned using fMRI. The investigators found that when participants viewed manipulable objects such as hammers, a circuit in the brain that underlies the

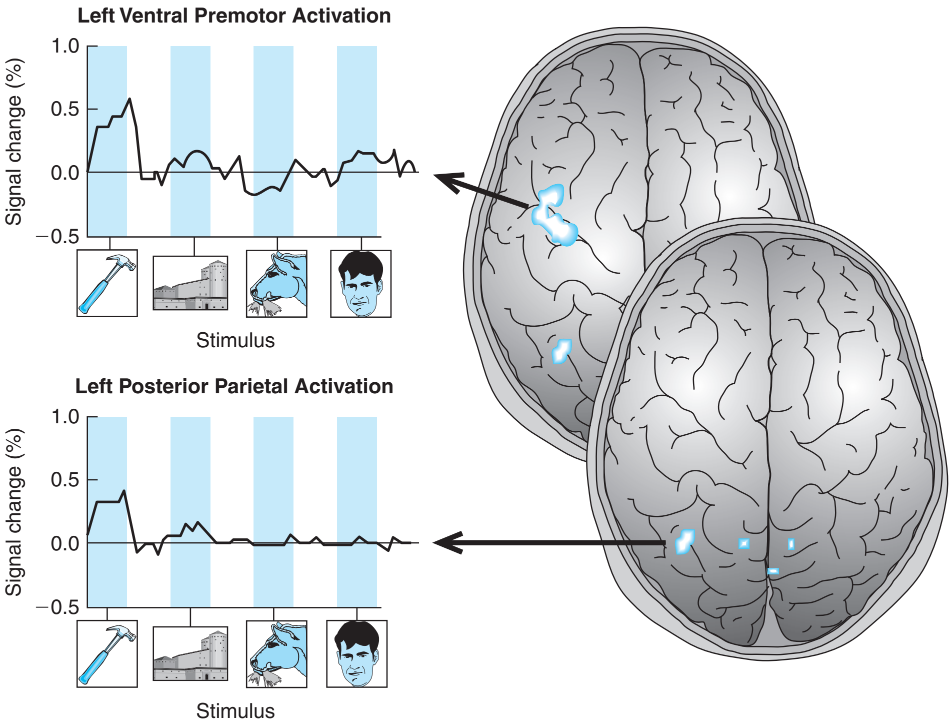

FIGURE 12 Neuroimaging support for category knowledge

The left-hemisphere grasping circuit (for right-handed participants) became active only while participants viewed pictures of tools, not while they viewed pictures of faces, animals, or buildings.

grasping of manipulable objects became active (Figure 12). This circuit did not become active when buildings, animals, or faces were observed. In much previous work, this grasping circuit has been found to become active while monkeys and humans perform actions with manipulable objects and while they watch others perform such actions (e.g., Rizzolatti et al., 2002). Even though Chao and Martin's participants were not allowed to move in the scanner, and even though they viewed no agents or actions, this grasping circuit nevertheless became active. From this result, the investigators concluded that activation of the grasping circuit constituted a motor inference about how to act on the perceived object. As participants viewed an object (for example, a hammer), they accessed category knowledge about it that included motor inferences (for example, "a hammer can be swung"). These inferences appear to be represented in the motor system, as we would expect if mental simulations are used to represent the objects and their categories.

Many further neuroimaging studies (reviewed by Martin, 2001; Martin & Chao, 2001; Martin et al., 2000) have shown that other modality-specific regions become active as other kinds of category knowledge are processed. In a striking correspondence, category knowledge about color, shape, and motion is processed near the respective brain areas that process this information in visual perception (Figure 13 on Color Insert B). When participants retrieve an object's shape properties, an area in the fusiform gyrus that overlaps visual shape processing areas becomes active during PET and fMRI scans. Similarly, when participants retrieve an object's color properties from category knowledge, an area in the occipital cortex that overlaps an area that processes color in perception (V4) becomes active. When participants think about performing actions on objects, motor areas become active. When participants retrieve an object's motion properties, regions in the posterior temporal gyrus that overlap motion processing areas in vision become active. When participants retrieve the sounds of objects, an auditory brain area becomes active (Kellenbach et al., 2001). And when they access knowledge of foods, gustatory areas in the brain become active that represent tastes (Simmons et al., 2005). Together these findings demonstrate that an object category's representation is distributed across the brain's perceptual and motor systems.

✓ Comprehension Check:

1. How might multimodal representations of a category's members become integrated in the brain to establish category knowledge?
2. What behavioral and neural evidence exists to support the hypothesis that the brain's modality-specific areas are involved in representing category knowledge?

4. STRUCTURES IN CATEGORY KNOWLEDGE

Category knowledge is not an undifferentiated mass of data; it contains many different structures, organized in many different ways. As we shall see in this section, exemplars, rules, prototypes, background knowledge, and schemata all play roles in creating the category knowledge that allows us to live lives cognizant of ourselves and the world around us. Furthermore, we possess powerful and dynamic abilities for using these structures.

4.1. Exemplars and Rules

The simplest structures that category knowledge contains are memories of individual category members; these are known as **exemplars**. The first time you see an unfamiliar type of dog and are told its breed, a memory of that dog is stored along with the name of the breed. As you see more of these dogs, a memory for each one similarly becomes associated with the breed name, and thereby with other memories of that breed. Over time, a collection of memories results for these category exemplars, all integrated in the appropriate memory store (as illustrated earlier in Figure 11a). This sort of content is relatively simple because each type of memory is stored independently of the others.

Much research has shown that exemplar memories are common in our category knowledge (e.g., Brooks, 1978; Lamberts, 1998; Medin & Schaffer, 1978; Nosofksy, 1984), and that they play a powerful role. For example, participants in a study by Allen and Brooks (1991) were told about two categories of imaginary animals; *builders* and *diggers*. Individual animals might have long or short legs, an angular or curved body, and spots or no spots. A **rule**—that is, a precise definition of the criteria for a category—determined whether a particular animal was a *builder* or *digger:*

> An animal is a *builder* if it has *two* or *three* of the following properties: *long legs, angular body, spots;* otherwise it is a *digger.*

Some participants were told the two-out-of-three rule. These participants were then shown pictures of the imaginary animals sequentially and instructed to indicate which were *builders* and which were *diggers*. Presumably they used the rule to do this, counting the number of critical properties for each animal. If they made an error, the experimenter told them the correct category. Once the participants demonstrated that they could apply the rule for the categories effectively, they received a surprise test. On each trial they saw an animal that they hadn't seen earlier. Again they had to say whether the animal was a *builder* or *digger,* but this time the experimenter didn't say whether their categorizations were correct or incorrect.

Allen and Brooks (1991) suspected that even though participants knew a rule for the categories, they might nevertheless be storing exemplar memories and using them in categorization. From earlier research, the investigators believed that the human brain automatically stores and uses exemplar memories, even when doing so is not necessary. But how to determine this? Figure 14 illustrates the clever technique that the investigators used. In the test phase of the experiment, participants were shown some of the animals they had seen before and some new ones. Two of the new ones were *builders*. One of these differed from a builder seen during training in only one characteristic; this type of correspondence, between two entities of the same category, is referred to as a *positive match*. The other new builder, while fulfilling the rule, differed in only one characteristic from a *digger* seen previously; this kind of correspondence, between two entities of different categories, is a *negative match*.

Here's the key prediction. If participants do not store exemplar memories and use only the rule, the positive- and negative-match animals should be equally easy to categorize: both fulfill the rule for builders. If, however, participants stored exemplar memories—even though they did not have to in order to make the call—then the negative-match animal, which shared more characteristics with the *digger,* should be harder to categorize correctly than the positive-match one.

Why? Think about what happens when participants encounter the negative-match animal. If an exemplar memory of its counterpart from training exists, it is likely to become active. If it does, then because the two animals are so similar, sharing as they do two characteristics, the negative-match animal is a reminder of the counterpart seen earlier. *But the counterpart was in the other category!* So if the exemplar memory is active, the temptation to miscategorize is strong; rule and exemplar memory conflict.

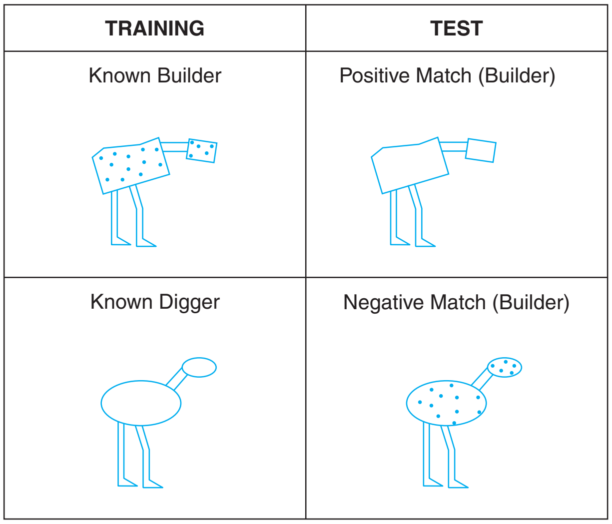

FIGURE 14 The original builders and diggers

Left column: A *builder* and a *digger* that participants studied while they learned the rule for *builders*. Right Column: Positive and negative test matches. A positive match was a *builder* that differed only by one property from a *builder* studied earlier; a negative match was a *builder* that differed only by one property from a *digger* studied earlier. If participants use only rules to categorize *builders,* the positive and negative matches should be equally easy to categorize, given that both have two of the three *builder* properties. Alternatively, if participants also use exemplars to categorize *builders,* the negative match should be harder to categorize because it is so similar to a member of the wrong category.

(Adapted from Allen, S. W., & Brooks, L. R. (1991). Specializing the operation of an explicit rule. *Journal of Experimental Psychology: General, 120,* pp. 3–19, Fig. 1, p. 4. Copyright © 1991 American Psychological Association. Adapted-with permission.)

What happens when participants encounter the positive-match animal and exemplar memory is active? Again, they will be reminded of the similar animal they saw in training, which in this case is in the correct category. This time both the exemplar memory and the rule point to the right answer.

Numbers told the story: the results demonstrated clearly not only that exemplar memories had been stored but also that they had a profound impact on categorization. Participants correctly categorized the positive-match exemplars 81 percent of the time, but correctly categorized the negative-match exemplars only 56 percent of the time. Even though participants knew a good rule for categorizing all the test animals, their memories of earlier exemplars intruded on categorization of negative-match animals, causing 25 percent more errors than in the other condition. If exemplar memories hadn't been stored, there shouldn't have been any difference in categorizing positive- and negative-match animals, given that both satisfied the rule equally well. Many similar findings in the literature demonstrate that exemplars are ubiquitous structures in category knowledge.

Does this finding suggest that we store only exemplars and not rules? Before we can answer, we need to look at another side of the coin. A second group of participants received the same training, learning by feedback from the experimenter whether their categorizations were correct. The difference? This second group was *not* told the rule for categorizing *builders* and *diggers*. These "no-rule" participants then were shown the same series of positive- and negative-match animals at test as the "rule" participants in the first group.

Two findings are of interest. Like the rule participants, the no-rule participants were more accurate on the positive-match animals (75 percent correct) than on the negative-match animals (15 percent correct). For these participants, similarity to exemplar memories played the central role in categorization. The effect of exemplar memories was significantly larger in the no-rule condition (incorrect negative-match categorization, 85 percent) than in the rule condition (incorrect negative-match categorization, 44 percent). Rule participants had stored a rule in their category knowledge that made them less vulnerable to exemplar memories than were no-rule participants. By applying the rule on some occasions, rule participants were more likely to categorize negative-match animals correctly. These and other results demonstrate that we can store rules for categories, not just exemplars (e.g., Ashby & Maddox, 1992; Blok et al., 2005; Nosofsky et al., 1994).

Thus the Allen and Brooks (1991) study established that, depending on the training conditions, we acquire exemplar memories, rules, *or both* for the categories we learn. To corroborate these behavioral findings with neural evidence, neuroimaging was conducted while two groups performed the task, either in the rule condition or the no-rule condition (Patalano et al., 2001; see also E. Smith et al., 1998). The investigators made the following predictions. First, in the no-rule condition, the brain areas used should be those that store exemplar memories (because the exemplars were experienced only visually, the primary sites of brain activation should be in the visual system). Second, in the rule condition, the primary brain areas used should be those that represent rules. (Because people rehearse a rule to themselves while they assess its fit to exemplars, motor areas that implement the implicit speech actions for rehearsal should become active.)

The results of the brain scans bear out the predictions. In the no-rule condition, most of the active sites were in occipital areas where vision is processed. As predicted, when participants did not know a rule, they primarily used visual memories of exemplars to categorize. In the rule condition, there were active sites in frontal motor areas. Again as predicted, when participants knew a rule, they rehearsed it silently to themselves, and the actions of internal rehearsal engaged the motor system.

The conclusion? Different brain systems become active to represent exemplars and to represent rules. Furthermore, the particular systems that become active support the notion that category knowledge is represented in modality-specific areas: visual areas represent the content of exemplars, motor areas implement the process of rehearsing rules. (For other work that also localizes various category representations in the brain, see Ashby & Ell, 2001.)