THE NEW COGNITIVE NEUROSCIENCES

Second Edition

Michael S. Gazzaniga, *Editor-in-Chief*

0-262-07195-9

A Bradford Book
The MIT Press
Cambridge, Massachusetts
London, England

70 The Cognitive Neuroscience of Categorization

EDWARD E. SMITH AND JOHN JONIDES

ABSTRACT Behavioral work on categorization has proposed three different categorization procedures–classification by rule, by stored exemplars, and by prototype–and has often assumed that just one of these procedures underlies all acts of categorization. In this chapter, we use findings from cognitive neuroscience to argue that all three procedures are used in categorization and, further, that the three procedures have different neural bases. We discuss evidence from two kinds of studies. The first kind involves neuropsychological experiments with amnesic patients, and deals with the contrast between exemplars and prototypes. These studies indicate that amnesic patients can use implicit representations–presumably prototypes–in categorization tasks in which normal subjects use explicit exemplars. The second kind of study involves neuroimaging normal subjects while they perform categorization tasks that encourage the use of either rules or exemplars. These experiments indicate that rule use involves at least some neural circuits that are distinct from those involved in categorization by exemplars.

CATEGORIZATION PROCEDURES We routinely+ divide the world into categories–a particular wooden configuration is perceived as a *chair*, a particular assemblage of body parts as a *dog*, and so on. Categorization greatly reduces the amount of information we have to process, and allows us to draw inferences about imperceptible properties (e.g., "If it's a dog, it may bark when provoked"). Categorization may be what makes possible human perception, memory, communication, and thought as we know it.

To categorize some test object "x" is to come to treat it as an instance of some category (a nonarbitrary class of objects). How do people do this? Presumably, we have some mental representations of various categories and procedures for deciding which of these mental representations provides the best fit for object x (Smith, 1995). In this chapter, we focus on categorization procedures themselves (as opposed to mental representations of categories). In particular, we consider three categorization procedures that have been widely studied in cognitive science, in order to see what recent findings in cognitive neuroscience can tell us about these procedures.

EDWARD E. SMITH and JOHN JONIDES Department of Psychology, University of Michigan, Ann Arbor, Mich.

The three categorization procedures of interest can be roughly characterized in the following way (after Smith, Patalano, and Jonides, 1998). In deciding whether a test object belongs to a particular category, one may

1. determine whether the test object fits a rule that defines the category ("rule application");
2. determine the similarity of the test object to one or more remembered exemplars of the category ("exemplar similarity"); or
3. determine the similarity of the test object to a prototype of the category ("prototype similarity").

One issue arising from the extensive study of these procedures is central: Do people use a variety of categorization procedures? Or is all (or most) human categorization based on a single procedure? In its currently popular form the latter possibility involves the following reasoning: Categorization is always based on exemplar similarity; hence, what appear to be cases of rule application and prototype similarity can be shown, under more intensive analyses, to involve an exemplar-similarity procedure (Estes, 1994; Nosofsky, 1992a,b). We refer to this approach as the "unitary view." The obvious alternative is the "multiple view"–that all three procedures are widely used in categorization. Among those who argue for multiple procedures, the major issues concern the nature of these mechanisms: What component processes are involved and how do they differ for the different procedures?

CONSTRAINTS FROM COGNITIVE NEUROSCIENCE Purely behavioral studies may be limited in what they can tell us about the question of unitary versus multiple categorization mechanisms. Most behavioral experiments on this issue have produced results that fit better with exemplar- rather than prototype-based mathematical models; however, this may be because exemplar models contain more information (they contain representations of *all* relevant exemplars, whereas prototype models discard much of this information). Also, while quantitative modeling favors exemplar models, some qualitative considerations favor prototype models (e.g., Smith, 1995). Thus, there is a clear need to consider another kind of evidence about categorization; for this reason, we turn to more neurologically based research.

Findings from cognitive neuroscience can contribute both to the general issue of unitary versus multiple categorization procedures and to specific issues about component processes. In this chapter, we consider two cases of such contributions from cognitive neuroscience. In the first section, we focus on exemplar-similarity and prototype-similarity procedures, and we consider evidence from amnesic patients in categorization tasks. Such patients should have serious difficulty in learning exemplars of new categories, and consequently their performance in categorization tasks can be revealing about an exemplar mechanism. If categorization is based solely on similarity to remembered exemplars, then amnesics' performance should be very poor compared to that of normals; but if categorization can be accomplished by other means, amnesics' performance might be normal in some cases.

In the second section, we focus on the contrast between exemplar similarity and rule application. Again, we consider some evidence from neuropsychology, noting how different kinds of patients perform in different kinds of tasks. Here, however, we concentrate on neuroimaging studies that attempt to show that different neural processes are involved when people categorize the same objects by different procedures.

Throughout, the studies that we review involve artificial categories (e.g., categories of dot patterns or imaginary animals) rather than natural categories (e.g., dogs or hammers). We emphasize artificial categories because of our concern with categorization procedures: Artificial categories maximize the chances that the categorization procedure of interest is employed. In seeking to determine if rule application is ever employed, one will do better to create, for example, a set of imaginary animals conforming to a novel rule instead of trusting that some particular natural category is, in fact, based on a rule.

Exemplars, prototypes, and categorization by amnesic patients

LOGIC OF THE RESEARCH It is useful to begin by fleshing out the two procedures of interest. To illustrate paradigmatic cases of exemplar- and prototype-similarity procedures, consider a situation in which a dermatologist must decide whether a particular skin lesion is an instance of disease Y.

Assuming the dermatologist has seen many patients, she will likely have stored in long-term memory numerous exemplars of various skin diseases. She may then note that the current lesion is very similar to stored exemplars of disease Y, and on this basis categorize the current lesion as an instance of Y. The sequence of processes presumably includes

1. retrieval of stored exemplars (of various disease categories) that are similar to the current lesion (the test object); and
2. selection of that category whose retrieved exemplars are (by some measure) most similar to the test object.

Note that if the exemplars retrieved in stage 1 all belong to the same category, then the selection process of stage 2 is trivial (or nonexistent). But if the exemplars retrieved in stage 1 belong to different categories, then stage 2 requires a systematic choice process (see, e.g., Estes, 1994; Nosofsky, 1992a,b).

Categorization based on prototype similarity is somewhat different. We assume that, as a consequence of seeing numerous patients, our dermatologist has abstracted from the individual cases a prototype of each relevant skin disease. This prototype is some measure of central tendency of lesions within a disease category—perhaps an average over the dimensions of shape, texture, and color, or the modal values on these dimensions. In any event, our dermatologist might note that the current lesion is more similar to the prototype of disease Y than it is to rival prototypes, and on this basis categorize the current lesion as an instance of Y. The sequence of processes presumably includes

1. retrieval of prototypes (of various disease categories) that are similar to the current lesion; and
2. selection of that category whose prototype is most similar to the test object.

In behavioral research we find numerous studies that directly contrast exemplar similarity and prototype similarity in specific categorization tasks (see Estes, 1994, for a partial review). Typically, the tasks require normal subjects to learn and subsequently use some artificial categories—say, two categories of dot patterns, or schematic faces. The researchers develop quantitative models of exemplar-similarity and prototype-similarity procedures, then determine which model provides a better fit to the data on category learning and use. As previously noted, many of these studies have favored exemplar-similarity models, so much so as to suggest that categorization with novel categories is always done by exemplar similarity (this is the unitary view). Importantly, these studies have frequently bolstered their case for exemplar-based categorization by showing that the representations assumed for categorization can also be invoked to explain performance on tests of recognition memory. That is, an experiment may require subjects to learn the instances of two novel categories of visual patterns, then test the subjects both on categorization with novel items and on recognition of the originally learned instances. One can then use modeling techniques to show that the same representations are involved both in categorization and long-term recognition memory (e.g., Nosofsky and Zaki, 1998; Shin and Nosofsky, 1992).

The link between categorization and recognition memory indicates that the memory representations involved are *explicit.* That is, the representations contain information about the context in which the item occurred, and may also permit conscious recollection of the item. These characteristics of explicit representations provide a basis for responding in tasks like recall and recognition. All of this is in contrast to *implicit* representations, which can provide a basis for performance changes but cannot be intentionally reinstated (see, e.g., Bower, 1998; Schacter, 1989).

This is where neuropsychology enters the picture. One of the best documented sets of findings in the field involves amnesic patients: Amnesics with damage either to medial-temporal–lobe or diencephalic structures have difficulty committing new information to explicit memory (see Squire, 1992, for a review). Amnesics should therefore have difficulty employing the exemplar-similarity procedure in learning and using novel categories, since the exemplars involved are presumably part of explicit memory. Under the unitary view that all categorization is based on exemplar similarity, amnesics should perform poorly in *all* tasks that require the learning and use of novel categories. Under the multiple-procedures view, however, amnesics should perform poorly in those categorization tasks that elicit an exemplar-similarity procedure, but they may perform normally on tasks that recruit other procedures, as long as these other procedures place a minimal load on explicit memory. Prototype similarity might be such an "other" procedure. These ideas were tested in the studies described in the following section.

Amnesic Performance in Categorization Tasks

Using the above logic, Kolodny (1994) compared normal controls and a group of medial-temporal–lobe and Korsakoff's amnesics on two different categorization tasks. One task involved paintings and presumably elicited an exemplar-similarity procedure, whereas the other task involved dot patterns and presumably triggered prototype similarity.

Consider first the dot-pattern task. In a *learning phase,* subjects were presented a sequence of dot patterns and were informed in which of three categories each pattern belonged. Every pattern, which contained nine dots, was generated by statistically distorting one of three prototype patterns. In a subsequent *test phase,* the subjects were presented the learned patterns plus novel ones, then asked to indicate the appropriate category for each one. This kind of task has been widely used in categorization studies with normal subjects, and performance on the task has typically been interpreted in terms of prototype similarity (e.g., Posner and Keele, 1968; Homa, Sterling, and Trepel, 1981). Following the categorization task, subjects were shown patterns they had categorized as well as completely novel ones, then were asked to decide which ones were "old" and which "new"; this is a recognition test of explicit memory.

The paintings task was similar in structure: a learning phase, followed by a test phase that involved learned and novel items, followed by a recognition task. But now the items were Renaissance paintings, and categorization required the subjects to learn which paintings were done by the same artist. Prior research indicated that categorization in this task was based on exemplar similarity; indeed, paintings done by the same painter were sufficiently dissimilar that it is difficult even to generate a plausible prototype for each artist (Hartley and Homa, 1981).

The results for the paintings task were exactly what one would expect if categorization were mediated by exemplar similarity. Amnesics performed far worse than normal controls during both the training and test phases of categorization. Indeed, amnesic categorization performance did not differ significantly from chance. As corroborating evidence that the control/amnesic difference in categorization was mediated by an underlying difference in retrieval from explicit long-term memory, amnesics performed more poorly than normals on the recognition memory test. There is, then, a clear connection between explicit memory and categorization of novel items, just as has been found in the studies with normals, and all of this is in keeping with an exemplar-similarity procedure.

The results are very different for the task involving dot patterns. Here, amnesics performed as well as normal controls during both the training and test phases of categorization. This suggests that categorization was based on a mechanism other than exemplar similarity. This suggestion is strengthened by the results on the recognition memory test, on which amnesics performed more poorly than normals. In the dot-pattern task, then, there is a dissociation between memory and categorization, with amnesics being impaired on the former but normal on the latter. The magnitude of this dissociation is striking, as amnesic patients achieved their normal categorization while performing at chance on the memory test. It therefore seems highly likely that a mechanism other than explicit exemplar similarity was involved in the categorization of dot patterns.[1]

A report by Squire and Knowlton (1995) contains even more dramatic evidence for a dissociation between categorization and explicit memory (see also Knowlton and Squire, 1993). These researchers worked with a severely amnesic patient E.P. who, according to standard tests, has virtually no capacity for explicit memory (unlike the patients in the preceding study, who scored above chance on recognition memory for paintings). But despite his complete loss of memory, E.P. is normal on the categorization of dot patterns. In one study, E.P. and a group of normal controls were presented a series of dot patterns during a training phase; all patterns were distortions of the same prototype. In a test phase, E.P. and the controls were presented novel patterns, which included the prototype itself, distortions relatively similar to the prototype ("low distortions"), distortions relatively dissimilar to the prototype ("high distortions"), and random dot patterns. The subjects' task was to decide which of these test patterns belonged to the same category as that exemplified in the training phase. As in prior studies with normal subjects, the controls gave their highest ratings ("Yes, it's a member of the category") to the prototype, next highest to the low distortions, and next to the high distortions, giving their lowest rating to the random patterns. E.P. did the same, and to the same degree. There were no significant differences between E.P.'s gradient of categorization and that of the normals.

This is only half the dissociation between categorization and memory. E.P. and the normal controls were also given an extremely simple recognition-memory test. A single dot pattern was presented 40 times in succession, and 5 minutes later subjects were given a recognition test in which they had to decide whether each pattern presented was the memory pattern or not. Unsurprisingly, control subjects were almost perfect on this test. But E.P. performed at chance. Thus E.P. completely failed the simplest test of explicit memory, but was perfectly normal on a categorization test with the same kind of stimuli. Clearly, E.P.'s categorization performance cannot be based on explicit memory, hence not on an (explicit) exemplar-similarity procedure. There must be some other procedure in use.

PROTOTYPES AND EXPLICIT VERSUS IMPLICIT MEMORY On the face of it, that other procedure would likely be prototype similarity. For one thing, prior work suggests that the categorization of dot patterns is accomplished by a prototype-similarity mechanism (e.g., Posner and Keele, 1968; Homa, Sterling, and Trepel, 1981; but see Shin and Nosofsky, 1992). Another point is that prototype similarity makes minimal demands on long-term explicit memory. Because the dot patterns are presented every few seconds during the training phase, subjects can extract a prototype simply by holding in working memory (which is spared in amnesia) the measure of central tendency they have extracted thus far, combining it with the next item to form a new prototype. But there is a problem with the claim that prototype similarity is the mechanism operative in the dot-pattern categorization tasks of the preceding studies. Subjects would have to store the prototype in long-term memory between the training and test phases (usually, a matter of minutes), as well as during the test phase (again, a matter of minutes). However, patient E.P., who cannot store a single item for minutes, performs normally on dot-pattern categorization.

There are a couple of possible solutions to this dilemma. One is that the categorization procedure that E.P. and other amnesic patients use is not prototype similarity, but some more primitive procedure that has been spared in amnesia. The only such primitive procedure in sight is a simple mechanism that forms associations between stimulus cues and categories in a manner analogous to classical conditioning. Amnesic subjects might, for example, learn to associate a set of dots in a certain region with category A. This kind of associative mechanism has been used to explain performance in probabilistic categorization tasks, tasks in which the features of objects are only probabilistically related to categorization. In such tasks, amnesics again sometimes perform as well as normals (e.g., Knowlton, Squire, and Gluck, 1994). However, it seems unlikely that the associative mechanism that underlies probabilistic categorization is involved in dot-pattern categorization, since there is a dissociation between the two tasks. Parkinson's patients are impaired in the probabilistic categorization task, yet perform normally on dot-pattern categorization. Hence, two different categorization mechanisms may be involved in the two tasks (Reber and Squire, 1997).[2]

So what procedure is used in dot-pattern categorization? Perhaps the abstraction and use of prototypes requires only implicit, not explicit, memory; consequently, amnesic patients may be able to use an implicit prototype for categorization purposes (which, of course, would not work for recognition). What would an implicit prototype contain? Likely, it would represent the central tendencies of the items presented (like an explicit prototype), but not associations with context (unlike an explicit prototype) (e.g., Bower, 1998). The establishment of such contextual associations is at least part of what distinguishes explicit from implicit representations, and this process requires an intact medial-temporal–diencephalic system.[3]

CONCLUSIONS What can we conclude about categorization procedures from this line of work? First, at least in some tasks, categorization need not be based on explicit memory. Because exemplar similarity has been tied to explicit memory in experiments with normal subjects, it follows that categorization need not be based on exemplar similarity. In contrast, there *are* tasks in which categorization performance is tied to explicit memory (e.g., the paintings task used by Kolodny, 1994), and in such cases exemplar similarity seems the most plausible account of categorization. Thus, in contrast to the unitary view, there appears to be more than one mechanism of categorization, and at least one of these mechanisms is spared in amnesia. The best guess about the spared mechanism is that it is prototype similarity, with the prototypes involved being implicit representations. These conclusions are in rough agreement with conclusions drawn by Knowlton, Squire, and their co-workers.[4]

Exemplars, rules, and categorization by normal subjects

LOGIC OF THE RESEARCH We turn now to a second line of research on categorization procedures, one that focuses on the contrast between exemplar similarity and rule application, and involves the neuroimaging of normal subjects as well as the behavioral study of neurological patients. We have already characterized the exemplar-similarity procedure, and it is useful to do the same for rule application. Returning to our dermatologist, suppose she knows the additive rule: If the lesion has a sufficient number of the following features–elliptic shape, bumpy texture, reddish-brown coloring, etc.–then disease Y is indicated. If the dermatologist applies this rule in making her diagnosis (categorization), presumably she will engage in the following sequence of processes (after Smith, Patalano, and Jonides, 1998):

1. selectively attend to each critical attribute of the test object (e.g., the shape, texture, and color of the lesion);
2. for each attended-to attribute, determine whether the perceptual information instantiates the value specified in the rule (e.g., "Is this color reddish-brown"?); and
3. amalgamate the outcomes of stage 2 so as to determine the final categorization.

The first stage involves selective attention, the second involves the perceptual instantiation of abstract conditions, and the third requires the working-memory operations of storing and combining information.

Given this characterization of rule application and our previous one of exemplar similarity, we can now consider how data from cognitive neuroscience can be used to determine if the putative mechanisms are distinct, as well as the nature of their component processes. One approach is to look for dissociations between the processes of interest in different patient populations; for example, brain damage in frontal regions may lead to a deficit in tasks requiring rule application, but not in tasks requiring exemplar similarity. We will briefly consider such neuropsychological evidence. A second approach is to use neuroimaging techniques to compare categorization based on rule application versus that based on exemplar similarity. This approach allows us to determine the neural regions activated in a particular categorization procedure, and then to use what is known about the functionality of the activated regions to infer the processes involved in the categorization procedure. We will review a PET study that embodies this approach.

NEUROLOGICAL PATIENT PERFORMANCE ON CATEGORIZATION TASKS What neural regions are involved in the application of rules? For years, the best guess has been frontal regions, particularly the dorsolateral prefrontal cortex (DLPFC). One source of suggestive evidence comes from clinical observations of patients with selective frontal lesions. Such patients seem to be particularly deficient in complex tasks, like planning and decision making, and some of this deficiency is thought to arise from the patients' inability to follow explicit rules (e.g., Luria, 1969). A second line of evidence comes from experiments that demonstrate that frontal-lobe patients, particularly those with DLPFC damage, have difficulty in a categorization task that requires the use of explicit rules, the Wisconsin Card Sort task. In this task, on each trial a card is presented that contains geometric forms, the forms varying from card to card with respect to number, shape, size, and background shading. The subjects must first learn which of the four attributes to use as a basis for sorting the cards; and, once they have learned this, the experimenter switches the relevant attribute, so subjects must now discover the new critical attribute. The basic finding is that frontal-lobe patients are relatively normal in learning the initial rule, but are impaired in shifting to a new rule when the experimenter switches relevant attributes (e.g., Milner, 1964).

This finding has frequently been interpreted in terms of deficient use of rules, but a related explanation may be more plausible. The frontal-lobe patients' major deficit may be in *switching* between rules rather than in *applying* them. Indeed, frontal-lobe patients may be deficient in switching their attention between *any* two mental processes. Support for this interpretation comes from studies showing that patients with DLPFC lesions are selectively impaired in switching between two simple tasks (e.g., sorting by color versus sorting by shape), even when compared to neurological patients who have lesions in other parts of frontal cortex (Rubenstein, Evans, and

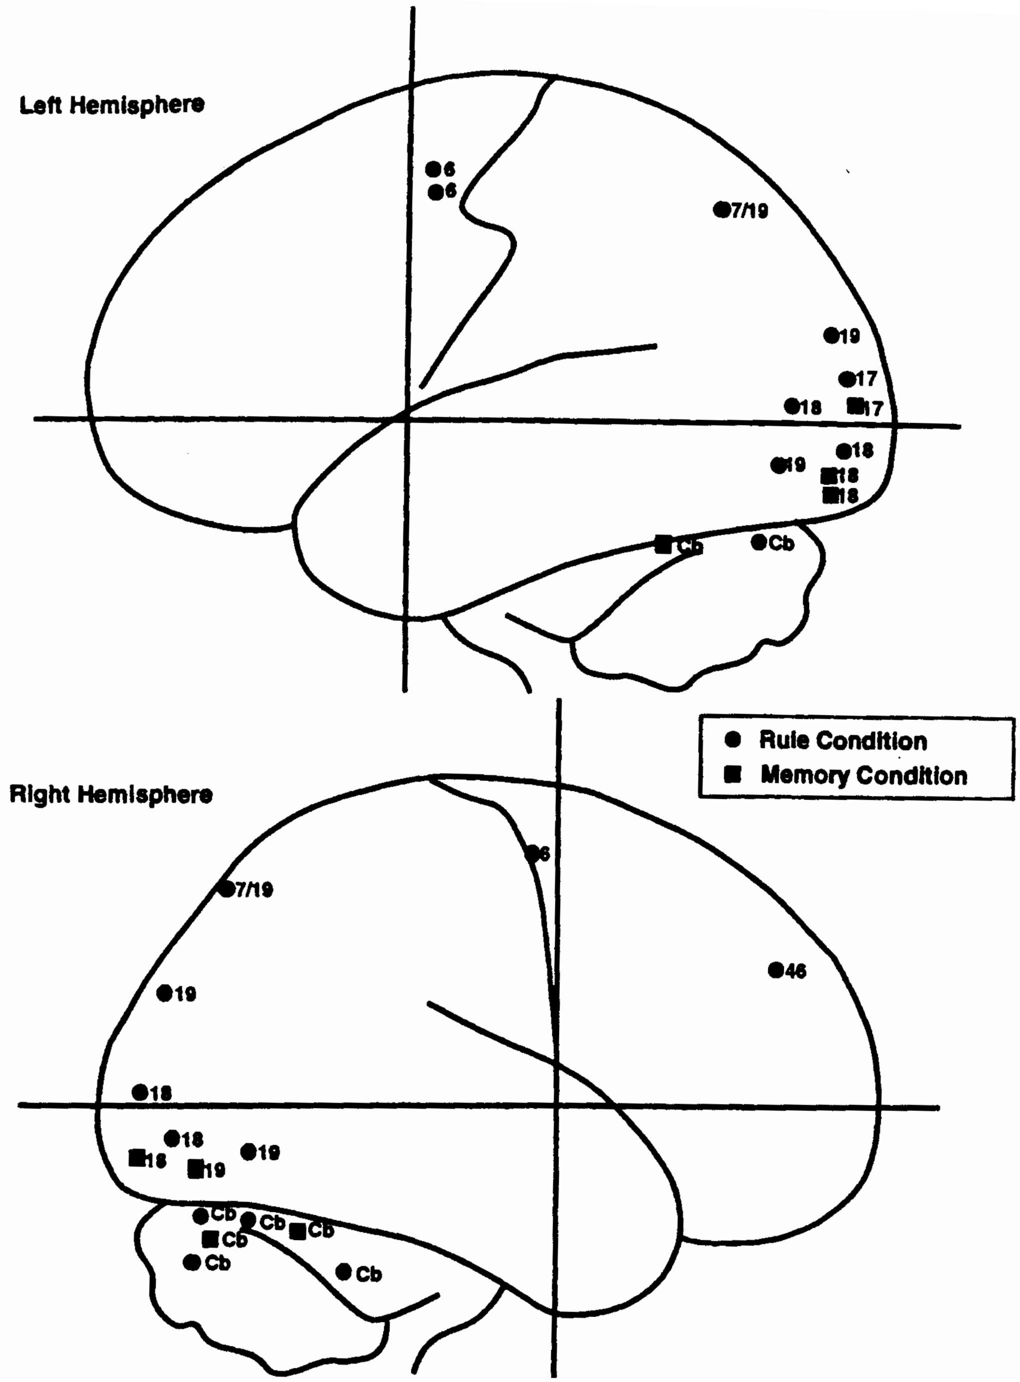

FIGURE 70.2 The marked areas on the left and right lateral surfaces of the cortex denote regions of significant activation in the rule- and memory-condition subtraction images. Circles designate the rule condition, squares the memory condition; numbers designate Brodmann areas, and Cb designates cerebellum. The only active regions that are not visible from this perspective are left BA 32 (memory condition only) and right thalamus (rule condition only). (Adapted from Smith, Patalano, and Jonides, 1998.)

This distinction between procedures is also supported by the neuropsychological evidence that we reviewed, though that evidence is considerably less diagnostic than the PET findings. All things considered, the evidence from cognitive neuroscience supports the cognitive distinction between rule application and exemplar similarity as categorization mechanisms. Earlier, we showed comparable evidence for a distinction between exemplar similarity and prototype similarity. Putting this together, we have support for three different categorization procedures, which argues against the unitary view of categorization and for the multiple-procedures view.

In addition, the various studies tell us something about the processes involved in the three categorization procedures. The poor performance of amnesic patients on a task that clearly requires exemplar retrieval (the paintings task) indicates that the heart of this procedure is retrieval from explicit, long-term memory. This proposal is consistent with the PET findings showing that most of the activation during memory-based categorization is in the posterior part of the brain, where visual memories are presumably stored. The good performance of amnesics on tasks that seem to require prototype similarity (dot-pattern categorization) indicates that, whatever the component processes involved in this procedure, the representations are implicit. And, finally, the PET findings obtained during rule-based categorization indicate that the processes involved in rule application include selective attention and working memory.

Finally, we note that while we have focused on procedures of categorization, much of the neural activation mediating any act of categorization is determined by the contents of the categories, i.e., by *what* is been categorized, not just *how* it is categorized. For example, in our PET study, we might have used verbal descriptions of the imaginary animals instead of pictures–and likely would have found less activation in occipital regions and more in language areas in both the rule and memory conditions. Similarly, neuroimaging studies of dot-pattern categorization show activation in occipital areas, areas that are known to be involved in the initial encoding of visual-spatial information (Reber, Stark, and Squire, 1998).

Content effects have also been obtained in neuroimaging studies of natural categories. In some experiments, subjects are first told the name of a target category (e.g., *vegetable*), and then decide whether named test objects (e.g., lettuce, apple) are instances of the target category. Activation is routinely observed in the left-hemisphere angular gyrus, known to be involved in the processing of linguistic information (e.g., Grossman, Robinson, and Jaggi, 1996). In other studies, the same kind of categorization task is used, but the test items are presented pictorially rather than verbally (e.g., Kosslyn, Alpert, and Thompson, 1995). Now activation is observed in occipital cortex, not in the angular gyrus. Thus, while the categorization procedure is presumably the same in these verbal and visual tasks–prototype or exemplar similarity (e.g., Smith, 1995)–the activation patterns vary in predictable fashion with the kind of information being categorized. A true understanding of the neural bases of categorization, then, must consider both the content of the representations and the procedures that operate on them.

ACKNOWLEDGMENTS Preparation of this paper, and some of the research reported here, was supported by grants from the National Institute on Aging and the Office of Naval Research.

NOTES

1. A dissociation between performance in categorization and recognition memory tasks has often been taken as strong evidence that different processes are involved in the two tasks. But insightful analyses by Nosofsky and Zaki (1998) challenge the evidentiary status of such dissociations. Essentially, these authors demonstrate that such a dissociation could arise if: (1) categorization is based on exemplar similarity, and hence both tasks involve the same memory representations; (2) the categorization task requires a less demanding memory discrimination than does the memory task; and (3) the amnesics have diminished memory capacity (which they surely do). Kolodny's (1994) specific dissociation between dot-pattern categorization and memory might not survive Nosofsky and Zaki's critique. However, the study of amnesic patient E.P. (Squire and Knowlton, 1995; Knowlton and Squire, 1993) would survive this critique (as noted by Nosofsky and Zaki themselves), presumably because the recognition memory task involved was extremely easy.
2. Nosofsky and Zaki's (1998) distrust of dissociations (see note 1) applies here as well. However, the dissociation at issue–that between probabilistic categorization and dot-pattern categorization–is unlikely to hinge on the two tasks' having differential memory requirements. But the dissociation could arise if (1) categorization is based on prototype similarity in both tasks; (2) the extent to which test objects were more similar to one category prototype than the other was greater in the dot-pattern than the probabilistic-categorization task, and hence the decision was more difficult in the probabilistic-categorization task; and (3) Parkinson's patients are impaired in making similarity-based decisions, and hence performed less well in probabilistic categorization.
3. Since we have raised the specter of an implicit prototype, what about the possibility that amnesics rely on implicit exemplars when categorizing dot patterns? At least two reasons argue against the latter possibility. First, if categorization can be based on exemplars, why didn't the amnesic subjects in Kolodny's (1994) experiment use such exemplars to perform normally on the paintings task? This task almost certainly requires exemplar-based categorization, yet amnesic performance on it was at chance, suggesting that the exemplars involved had to be explicitly represented. Second, as noted earlier, there is some independent evidence that the categorization of dot patterns is done by a prototype (e.g., Homa, Sterling, and Trepel, 1981).
4. Ashby, Alfonso-Reese, and Turken (1998) have also proposed distinct neural systems for explicit and implicit categorization mechanisms, but they identify the explicit system only with rule application.
5. Each area of activation contains a peak–a point of greatest change in activation within the area–that can be specified in

an x,y,z-coordinate system. An area of activation found in the rule condition is considered "common" to an area found in the memory condition if the peaks of the two areas differ by less than 10 mm–the approximate spatial resolution of PET–on each coordinate.

REFERENCES

ALLEN, S. W., and L. R. BROOKS, 1991. Specializing the operation of an explicit rule. *J. Exp. Psychol.: Gen.* 120:3–19.

ASHBY, F.G., L. A. ALFONSO-REESE, A. U. TURKEN, and E. M. WALDRON, 1998. A neuropsychological theory of multiple systems in category learning. *Psychol. Rev.* 105:442–481.

BOWER, G. H., 1998. An associative theory of implicit and explicit memory. In *Theories of Memory II*, M. A. Conway, S. E. Gathercole, and C. Cornoldi, eds. East Sussex: Psychology Press, pp. 25–60.

D'ESPOSITO, M., J. DETRE, D. C. ALSOP, R. K. SHIN, S. ATLAS, and M. GROSSMAN, 1995. The neural basis of the central executive system of working memory. *Nature* 378:279–281.

ESTES, W. K., 1994. *Classification and Cognition.* New York: Oxford University Press.

GROSSMAN, M., K. ROBINSON, and J. JAGGI, 1996. The neural basis for semantic memory: Converging evidence from Alzheimer's disease. *Brain Lang.* 55:96–98.

HARTLEY, J., and D. HOMA, 1981. Abstraction of stylistic concepts. *J. Exp. Psychol.: Hum. Learn. Mem.* 7:33–46.

HOMA, D., S. STERLING, and L. TREPEL, 1981. Limitations of exemplar-based generalization and the abstraction of categorical information. *J. Exp. Psychol.: Hum. Learn. Mem.* 7:418–439.

KNOWLTON, B. J., and L. R. SQUIRE, 1993. The learning of categories: Parallel brain systems for item memory and category knowledge. *Science* 262:1747–1749.

KNOWLTON, B. J., L. R. SQUIRE, and M. A. GLUCK, 1994. Probabilistic classification learning in amnesia. *Learn. Mem.* 1:106–120.

KOLODNY, J. A., 1994. Memory processes in classification learning: An investigation of amnesic performance in categorization of dot patterns and artistic styles. *Psychol. Sci.* 5:164–169.

KOSSLYN, S. M., H. M. ALPERT, and W. L. THOMPSON, 1995. Identifying objects at different levels of hierarchy: A position emission tomography study. *Hum. Brain Map.* 3:107–132.

LURIA, A. R., 1969. Frontal lobe syndromes. In *Handbook of Clinical Neuropsychology*, Vol. 2, P. J. Vinken and G. W. Bruyn, eds. Amsterdam: North Holland.

MILNER, B., 1964. Some effects of frontal lobectomy in man. In *The Frontal Granular Cortex and Behavior*, J. M. Warren and K. Akert, eds. New York: McGraw-Hill.

NOSOFSKY, R. M., 1992a. Exemplar-based approach to relating categorization, identification, and recognition. In *Multidimensional Models of Perception and Cognition*, F. G. Ashby, ed. Hillsdale, N.J.: Lawrence Erlbaum, pp. 363–393.

NOSOFSKY, R. M., 1992b. Exemplars, prototypes, and similarity rules. In *From Learning Theory to Connectionist Theory: Essays in Honor of William K. Estes*, Vol. 1, A. F. Healy, S. M. Kosslyn, and R. M. Shiffrin, eds. Hillsdale, N.J.: Lawrence Erlbaum, pp. 149–168.

NOSOFSKY, R. M., and S. R. ZAKI, 1998. Dissociations between categorization and recognition. *Psychol. Sci.*, 9:247–255.

POSNER, M. I., and S. DEHAENE, 1994. Attentional networks. *Trends Neurosci.* 17:75–79.

POSNER, M. I., and S. W. KEELE, 1968. On the genesis of abstract ideas. *J. Exp. Psychol.* 77:353–363.

POSNER, M. I., S. E. PETERSEN, P. T. FOX, and M. E. RAICHLE, 1988. Localization of cognitive functions in the human brain. *Science* 240:1627–1631.

REBER, P. J., and L. R. SQUIRE, 1997. Implicit learning in patients with Parkinson's disease. Poster presented at the 38th Annual Meeting of the Psychonomic Society. Philadelphia, Pa.

REBER, P. J., C. E. L. STARK, and L. R. SQUIRE, 1998. Cortical areas supporting category learning identified using functional MRI. *Proc. Natl. Acad. Sci. U.S.A.* 95:747–750.

RUBENSTEIN, J., J. E. EVANS, and D. E. MEYER, 1994. Task switching in patients with prefrontal cortex damage. Paper presented at the annual meeting of the Cognitive Neuroscience Society, San Francisco, Calif.

SCHACTER, D. L., 1989. Memory. In *Foundations of Cognitive Science*, M. I. Posner, ed. Cambridge, Mass.: MIT Press, pp. 683–726.

SHIN, H. J., and R. M. NOSOFSKY, 1992. Similarity-scaling studies of "dot-pattern" classification and recognition. *J. Exp. Psychol.: Gen.* 121:278–304.

SMITH, E. E., 1995. Concepts and categorization. In *Invitation to Cognitive Science*, Vol. 3: *Thinking* (2d Ed.), E. E. Smith and D. Osherson, eds. Cambridge, Mass.: MIT Press.

SMITH, E. E., and J. JONIDES, 1997. Working memory: A view from neuroimaging. *Cogn. Psychol.* 33:5–42.

SMITH, E. E., A. PATALANO, and J. JONIDES, 1998. Alternative strategies of categorization. *Cognition* 65:167–196.

SQUIRE, L. R., 1992. Memory and the hippocampus: A synthesis from findings with rats, monkeys, and humans. *Psychol. Rev.* 99:195–231.

SQUIRE, L. R., and B. J. KNOWLTON, 1995. Learning about categories in the absence of memory. *Proc. Natl. Acad. Sci. U.S.A.* 92:12470–12474.