The neurobiology of human category learning

F. Gregory Ashby and Shawn W. Ell

Categorization is among the most important skills that any organism can possess. Recent advances in cognitive neuroscience have led to new insights about the neural basis of category learning. Perhaps most important is the finding that many different, widely separated neural structures appear to participate in category learning, but to varying degrees that depend on category structure. In particular, different brain regions are implicated according to whether the category-learning task involves explicit rules, prototype distortion or information integration.

Is the plant edible or poisonous? Is the person friend or foe? Was the sound made by a predator or by the wind? All organisms assign objects and events in the environment to separate classes or categories. This allows them to respond differently, for example, to nutrients and poisons, and to predators and prey. Any species that lacked this ability would quickly become extinct[1].

Until just a few years ago, categorization was treated like the familiar black box – that is, little, if anything, was known about how the brain performs this vitally important skill. Instead, competing theories tried to account for observable categorization behavior by postulating hypothetical intervening processes. Because these were unobservable, they could not be used to test the competing theories. As a consequence, several theories were developed that assumed markedly different psychological processes but were about equally successful at accounting for categorization behavior[2]. Within the past decade, however, new discoveries in neuroscience and neuropsychology have greatly changed this outlook.

We are now beginning to learn a great deal about the neural structures and pathways that mediate category learning. This knowledge has yielded some surprising conclusions. Perhaps most important is the finding that many different, widely separated neural structures appear to participate in category learning, but to varying degrees that depend on category structure.

F. Gregory Ashby*
Shawn W. Ell
Department of Psychology, University of California, Santa Barbara, CA 93106, USA.
*e-mail: ashby@ psych.ucsb.edu

Category learning versus the representation of highly learned categories

This article focuses on category learning, and not on the categorization behavior of highly experienced experts. This distinction is important because there is good evidence that the neural mechanisms and pathways that mediate the learning of new categories are different from the neural structures that mediate the representation of highly learned categories. For example, neuropsychological groups that are impaired in category learning (e.g. frontal patients and individuals with Parkinson's disease) do not lose old, familiar categories (e.g. fruits and tools).

Interest in the representation of familiar categories has been sparked in recent years by reports of a variety of category-specific agnosias that result from damage to certain high-level visual cortical areas. Category-specific agnosia refers to the ability to perceive or categorize most visual stimuli normally, but a reduced ability to recognize exemplars from some specific category, such as inanimate objects (e.g. tools) or fruits. The most widely known of such deficits occur with human faces (i.e. prosopagnosia). The existence of category-specific agnosias raises two questions that are important for understanding the neurobiology of category learning. First, are familiar visual categories represented in visual cortex? Second, does category learning occur in visual cortex?

A complete answer to the first of these questions is beyond the scope of this article. Interested readers are referred to any of several excellent reviews of the category representation literature[3]. Suffice it to say that although category-specific agnosias are consistent with the hypothesis that category structure is represented in the visual cortex, they are also generally consistent with the hypothesis that visually similar objects are represented in nearby areas of the visual cortex. For example, it is well known that neighboring cells in this region tend to fire in response to similar stimuli. Thus, damage to some contiguous region of the visual cortex is likely to lead to perception deficits within a class of similar stimuli. The debate over these two hypotheses continues[4].

Even if the representation of familiar categories is in the visual cortex, this does not necessarily imply that the learning of these categories occurs in the same visual areas. For example, if this were true then category-learning deficits should co-occur with category-specific agnosias. As mentioned above, however, groups with known category-learning deficits do not exhibit category-specific agnosia.

Other evidence against the hypothesis that category learning occurs in the visual cortex has been obtained in single cell recording experiments with monkeys. For example, Rolls *et al.* recorded from cells in a high-level area of the visual (i.e. inferotemporal) cortex of monkeys[5]. In these experiments, one visual stimulus was associated with reward and one with a mildly aversive taste. After training, the rewards were switched. Thus, in effect, the animals were taught two simple categories (i.e. 'good' and 'bad') and then the category assignments were switched. If the categories were represented in visual cortex, then the firing properties of visual cortical cells should have changed when the category memberships were switched. However, Rolls *et al.* found no change in the

---

Box 1. Rule-based category-learning tasks

Figure I shows the stimuli and category structure of a recent rule-based task that used eight exemplars per category[a]. The categorization stimuli were colored geometric figures presented on a colored background. The stimuli varied on four binary-valued dimensions: background color (blue or yellow), embedded symbol color (red or green), symbol number (1 or 2) and symbol shape (square or circle). This yields a total of 16 possible stimuli. To create rule-based category structures, one dimension is selected arbitrarily to be relevant. The two values on that dimension are then assigned to the two contrasting categories.

An important property of rule-based category-learning tasks is that the optimal rule is often easy to describe verbally[b]. As a result, subjects can learn the category structures via an explicit process of hypothesis testing[c] or theory construction and testing[d]. Unlike most information-integration tasks (see Box 2), at the end of training subjects are usually able to describe quite accurately the rule they used in rule--based tasks. Virtually all categorization tasks used in neuropsychological assessment are rule based, including the well-known Wisconsin Card Sorting Test[e].

Fig. I. Category structure of a rule-based category-learning task. The optimal rule is: respond A if the background color is blue, and respond B if the background color is yellow.

References
a Waldron, E.M. and Ashby, F.G. The effects of concurrent task interference on category learning: evidence for multiple category learning systems. *Psychol. Bull Rev.* (in press)
b Ashby, F.G. *et al.* (1998) A neuropsychological theory of multiple systems in category learning. *Psychol. Rev.* 105, 442–481
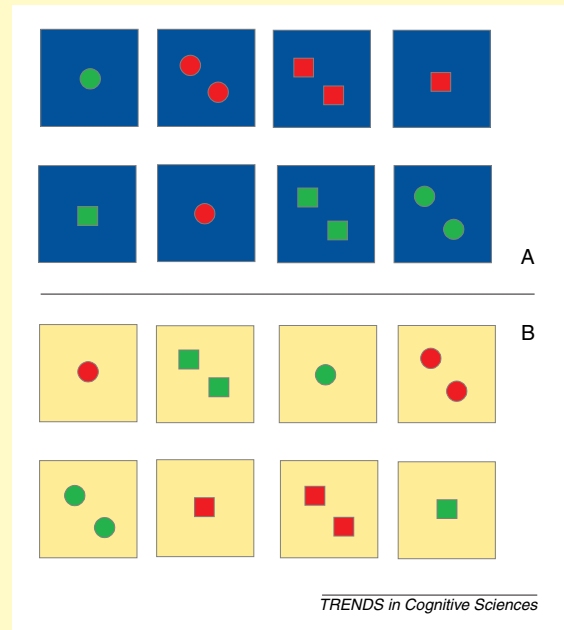c Bruner, J.S. *et al.* (1956) *A Study of Thinking*, John Wiley & Sons
d Murphy, G.L. and Medin, D.L. (1985) The role of theories in conceptual coherence. *Psychol. Rev.* 92, 289–316
e Kolb, B. and Whishaw, I.Q. (1990) *Fundamentals of Human Neuropsychology*, 3rd edn, Freeman

---

response of any of these cortical cells, although other similar studies have found changes in the responses of cells in other brain areas (e.g. orbitofrontal cortex).

Category-learning tasks

If the goal is to study category learning rather than category representation, then it is necessary to present subjects with unfamiliar categories and observe their behavior during the period when their ability to assign stimuli to these categories rises from chance to some stable level. In experiments with adults, the prevailing method of ensuring unfamiliarity is for the experimenter to create new, arbitrary categories of objects (so-called 'artificial categories'). In the past, little attention has been paid to the manner in which these arbitrary categories have been created. However, much recent evidence suggests there might be at least some differences in the neural circuitry that mediate category learning, depending on exactly how the categories are constructed. In fact, the available evidence identifies at least three different kinds of category-learning tasks.

Rule-based tasks are those in which subjects can learn the category structures via some explicit reasoning process. Frequently, the rule that maximizes accuracy (i.e. the optimal rule) is easy to describe verbally[6]. In the most common applications,

only one stimulus dimension is relevant, and the subject's task is to discover this relevant dimension and then to map the different dimensional values to the relevant categories. An example is shown in Box 1. Virtually all standard neuropsychological categorization tasks are of this type – including the well known Wisconsin Card Sorting Test[7]. Rule-based tasks, which have a long history in cognitive psychology, have been favored by proponents of the so-called classical theory of categorization, which assumes category learning is the process of discovering the set of necessary and sufficient conditions that determine category membership[8].

Information-integration tasks are those in which accuracy is maximized only if information from two or more stimulus components (or dimensions) must be integrated at some pre-decisional stage[9]. A conjunction rule (e.g. respond A if the stimulus is small on dimension *x* and small on dimension *y*) is a rule-based task, rather than an information-integration task, because separate decisions are first made about each dimension (e.g. small or large) and then the outcome of these decisions is combined (integration is not pre-decisional). In many cases, the optimal rule in information-integration tasks is difficult or impossible to describe verbally[6]. The neuropsychological data reviewed below suggests

Box 2. Information-integration category-learning tasks

Information-integration category-learning tasks are those in which accuracy is maximized only if information from two or more stimulus components is integrated at some pre-decisional stage[a]. Generally, it is difficult or impossible to describe verbally the rule that separates the contrasting categories[b]. Neuropsychological evidence discussed in the text indicates that people might learn differently in these tasks, depending on whether the categories contain few or many exemplars. Examples are illustrated below. Figure I shows the stimuli and category structure of a recent information-integration task that used only eight exemplars per category[c]. The categorization stimuli are described in Box 1. To create information-integration category structures, one dimension is arbitrarily selected to be irrelevant. For example, in Fig. I, the irrelevant dimension is symbol shape. Next, one level on each relevant dimension is arbitrarily assigned a value of +1 and the other level is assigned a value of 0. In Fig. I, a background color of blue, a symbol color of green, and a symbol number of 2 are all assigned a value of +1. Finally, the category assignments are determined by the following rule: the stimulus belongs to category A if the sum of values on the relevant dimensions is >1.5; otherwise it belongs to category B.

This rule is readily learned by healthy young adults, but even after achieving perfect performance, they can virtually never accurately describe the rule they used.

Figure II is an abstract representation of the category structure of an information-integration task in which there are hundreds of exemplars in each category (first developed in Ref. a). In this experiment, each stimulus is a line that varies across trials in length and orientation. Each blue cross in Fig. II denotes the length and orientation of an exemplar in category A and each red dot denotes the length and orientation of an exemplar in category B. The categories overlap, so perfect accuracy is impossible in this example. Even so, the quadratic curve is the boundary that maximizes response accuracy. Note that this curve is impossible to describe verbally.

References
a Ashby, F.G. and Gott, R.E. (1988) Decision rules in the perception and categorization of multidimensional stimuli. *J. Exp. Psychol. Learn. Mem. Cognit.* 14, 33–53
b Ashby, F.G. *et al.* (1998) A neuropsychological theory of multiple systems in category learning. *Psychol. Rev.* 105, 442–481
c Waldron, E.M. and Ashby, F.G. The effects of concurrent task interference on category learning: Evidence for multiple category learning systems. *Psychol. Bull. Rev.* (in press)
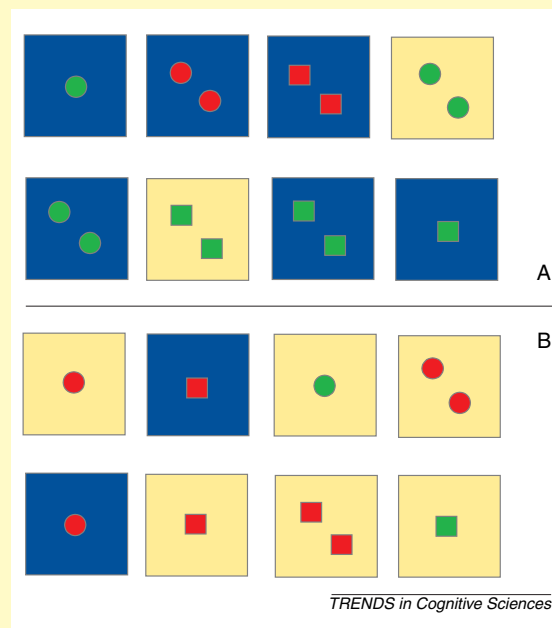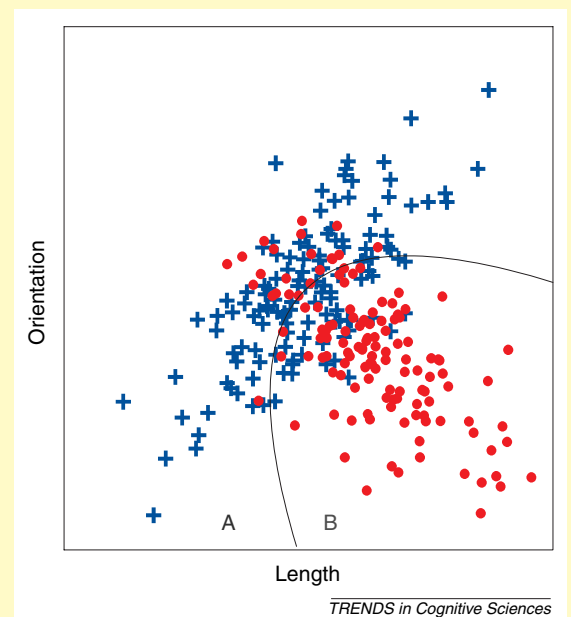
Fig. I. Category structure of an information-integration category-learning task with only a few exemplars in each category. Symbol shape is irrelevant. A background color of blue, a symbol color of green and a symbol number of two are all arbitrarily assigned a numerical value of +1, whereas all other dimensional values are assigned a value of 0. The optimal rule is: respond A if the sum of values on the relevant dimensions is >1.5, otherwise respond B.

Fig. II. Category structure of an information-integration category-learning task with many exemplars per category. Each stimulus is a line that varies across trials in length and orientation. Each blue cross denotes the length and orientation of an exemplar of category A and each red dot denotes the length and orientation of an exemplar of category B. The quadratic curve is the boundary that maximizes response accuracy.

that performance in such tasks is qualitatively different depending on the size of the categories – in particular, when a category contains only a few highly distinct exemplars, memorization is feasible.

However, when the relevant categories contain many exemplars (e.g. hundreds), memorization is less efficient. Examples of these two types of information-integration tasks are described in Box 2.

In the most popular version of the prototype distortion task, the category exemplars are random dot patterns[a]. An example is shown in Fig. I. In a typical application, many stimuli are created by randomly placing a number of dots on the display. One of these dot patterns is then chosen as the prototype for category A. The others become stimuli not belonging to category A. The other exemplars in category A are then created by randomly perturbing the position of each dot in the category A prototype.

Reference

a Posner, M.I. and Keele, S.W. (1968) On the genesis of abstract ideas. *J. Exp. Psychol.* 77, 353–363
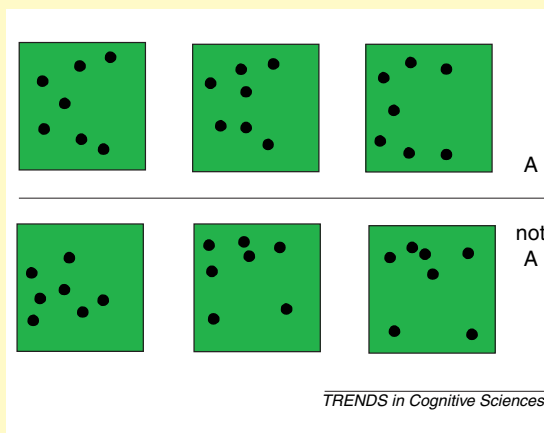

TRENDS in Cognitive Sciences

Fig. I. Some exemplars from a prototype distortion category-learning task with random dot patterns.

Information-integration tasks with few exemplars per category have been the favorites of exemplar theorists, who argue that categorization requires accessing the memory representations of every previously seen exemplar from each relevant category[10–12]. In contrast, decision bound theorists, who argue that category learning is a process of associating category labels with regions of perceptual space, have traditionally used information-integration tasks with many exemplars per category[9,13,14].

Prototype distortion tasks are a third type of category learning task in which each category is created by first defining a category prototype and then creating the category members by randomly distorting these prototypes. An example is shown in Box 3. As the name suggests, prototype distortion tasks have been commonly used by prototype theorists, who argue that categorization is the act of comparing the presented stimulus with the prototype of each contrasting category[15,16].

Neuropsychological data on category learning

A wide variety of different evidence supports the hypothesis that category learning in these three tasks

is mediated, at least in part, by different neural circuits. The most extensive of such data come from neuropsychological studies with different patient groups. Although categorization has been studied in many different neuropsychological groups, the most extensive data come primarily from studies with three different groups: (1) individuals with frontal lobe lesions; (2) individuals with amnesia, of which the most theoretically interesting are those whose amnesia was caused by damage to the medial temporal lobes; and (3) individuals suffering from a disease of the basal ganglia – typically either Parkinson's or Huntington's disease. In each case, the literature contains conflicting reports about the category-learning abilities of these populations. However, as Table 1 shows, when the existing studies are partitioned according to the type of task that was used, the discrepancies disappear.

'…much recent evidence suggests there might be at least some differences in the neural circuitry that mediate category learning'

Individuals with frontal lobe or basal ganglia dysfunction are impaired in rule-based tasks[17–20], but individuals with medial temporal lobe damage are normal in this type of category learning[21,22]. Thus, an obvious first hypothesis is that the prefrontal cortex and the basal ganglia participate in rule-based category learning, but the medial temporal lobes do not. Next, note that in information-integration tasks with large categories, only individuals with basal ganglia dysfunction are known to be impaired[23,24]. In particular, individuals with medial temporal lobe dysfunction are normal[25]. So a first hypothesis should be that the basal ganglia are crucial in this task, but the medial temporal lobes are not. If the number of exemplars per category is reduced in this task to a small number (e.g. between four and eight), then medial temporal lobe amnesiacs show late training deficits – that is, they learn normally during the first 50 trials or so, but thereafter show impaired learning relative to age-matched controls[26]. An obvious possibility in this case, is that normal subjects begin memorizing responses to at least a few of the more distinctive stimuli – a strategy that is not available to the medial temporal lobe amnesiacs, and which is either not helpful or impossible when the categories contain many exemplars. As patients with basal ganglia dysfunction are also impaired with these small, information-integration categories[27,28], a first hypothesis should be that learning in such tasks depends on the basal ganglia and on medial temporal lobe structures.

Finally, none of these patient groups is impaired in the prototype distortion tasks, which suggests that

Table 1. Performance of various neuropsychological populations on three types of category learning tasks

Neuropsychological group	Task[a]	Information integration		Prototype distortion
	Rule based	Many exemplars	Few exemplars	
Frontal lobe lesions	Impaired	?	Normal	?
Basal ganglia disease	Impaired	Impaired	Impaired	Normal
Medial temporal lobe amnesia	Normal	Normal	Late training deficit	Normal
[a]? indicates no known studies				

learning on these tasks does not depend on an intact medial temporal lobe or basal ganglia[28–31]. It has been suggested instead that learning might depend on a perceptual representation memory system – through a perceptual learning process[28]. In the random dot pattern experiments this makes sense, because all category A exemplars are created by randomly perturbing the positions of the dots that form the category A prototype (see Box 3). Thus, if there are cells in the visual cortex that respond strongly to the category A prototype, they are also likely to respond to the other category A exemplars, and perceptual learning will increase their response. If this occurs, the subject could perform well in this task by responding 'yes' to any stimulus that produces a strong feeling of visual familiarity.

Neuroimaging data on category learning
Evidence for multiple neural circuits in category learning
A variety of other evidence supports the hypothesis that learning in these three tasks is mediated by different neural circuits. Some of this is from neuroimaging studies. Although neuroimaging lends itself well to the study of category representation, its use as a tool to study category learning is more problematic, largely because learning necessarily involves changes over time and neuroimaging is most straightforward when a single time point is studied. For this reason, there are currently only a few neuroimaging studies of category learning.

Nevertheless, a number of neuroimaging studies are worth noting. In one of the earliest to study categorization[32], subjects performed a task like the one shown in Box 2 (Fig. I). In one condition, subjects were instructed to memorize the stimuli while in a second condition different subjects were trained to apply a complex explicit rule that successfully partitioned the stimuli into the contrasting categories. Positron-emission tomography scans showed substantially different activation patterns in the two conditions, suggesting that different neural circuits were mediating category learning in these different conditions.

Converging evidence that different neural circuits sometimes mediate category learning comes from a recent study showing that a dual task known to activate frontal cortex (a numerical Stroop task) interfered much more with a simple rule-based task than with a much more difficult information-

integration task[33]. If the same neural processes mediated learning in both tasks, then the dual task should have interfered more strongly with the difficult task than with the easy task (i.e. it is harder to do two difficult things at once than two easy things).

Results from rule-based, information-integration and prototype distortion tasks
A number of neuroimaging studies have shed light on the neural structures that mediate learning of the three tasks shown in Table 1. To a remarkable degree, the results of these studies agree with the neuropsychological data summarized in Table 1.

A functional magnetic resonance imaging (fMRI) study of a rule-based task similar to the Wisconsin Card Sorting Test showed activation in the right dorsolateral prefrontal cortex, the anterior cingulate and the right caudate nucleus (head)[34]. Converging evidence for the hypothesis that these are important structures in rule-based category learning comes from several sources. First are the many studies that have implicated these structures as key components of executive attention[35] and working memory[36,37], both of which are likely to be crucially important to the explicit processes of rule formation and testing that are assumed to mediate rule-based category learning.

Second, a recent neuroimaging study identified the (dorsal) anterior cingulate as the site of hypothesis generation in a rule-based category-learning task[38]. Third, lesion studies in rats implicate the dorsal caudate nucleus in rule switching[39]. Fourth, of course, are the neuropsychological data reviewed in Table 1, which show that groups of individuals with damage to any of these structures are impaired in rule-based tasks. Note also, however, that these conclusions suggest that the rule-based deficits seen in Parkinson's disease are due primarily to dysfunction in the head of the caudate nucleus. This conclusion is consistent with postmortem autopsy, which reveals that damage to the head of the caudate is especially severe in Parkinson's disease[40]. In fact, because of its reciprocal connections to the prefrontal cortex, many of the well-documented 'frontal-like' symptoms of Parkinson's disease might actually be due to damage in the head of the caudate nucleus.

Poldrack *et al.* used fMRI to measure neural activation at four different time points of learning in a probabilistic version of the information-integration

Table 2. Brain regions that current data implicate in the learning of various categorization tasks

Neuropsychological group	Task Rule based	Information integration Many exemplars	Few exemplars	Prototype distortion
Prefrontal cortex	Yes	–	–	–
Visual cortex	–	–	–	Yes
Basal ganglia	Yes	Yes	Yes	–
Medial temporal lobe	–	–	Yes	–

task with few exemplars per category[41]. They reported learning-related changes within prefrontal cortex and in the tail of the right caudate nucleus. Interestingly, they also reported a simultaneous suppression of activity within the medial temporal lobes. Thus, the available neuroimaging data predict that the deficits of individuals with basal ganglia disease in information-integration tasks might arise from dysfunction in the tail of the caudate nucleus.

Finally, recent fMRI studies of subjects in prototype distortion tasks show learning-related changes in the visual cortex[42], and are thus consistent with the hypothesis that learning in this task depends on the perceptual representation memory system.

Table 2 summarizes the neural implications of all these data. It indicates neural structures that available data suggest are especially important in the learning of the different categorization tasks. Of course, many other structures (e.g. the visual cortex) will be crucial, even if they are not the locus of learning.

Implications for computational modeling
These results have a number of implications for future attempts to build computational models of category learning. First, computational models

must become considerably more general than they have been in the past, in order to account for the great diversity in category-learning data discussed in this article. These include neuropsychological and neuroimaging data, as well as the more traditional cognitive–behavioral data.

> **'many of the well-documented 'frontal-like' symptoms of Parkinson's disease might actually be due to damage in the caudate nucleus'**

Second, largely because of this diversity, it is unrealistic to expect any single study to determine the correct type of computational model. Instead, it is imperative that all available evidence be evaluated simultaneously. For example, a vigorous debate is currently being waged in the category-learning literature between single-system and multiple-systems models. When comparing such different model types, it is vital to compare their ability to account for multiple data sets, rather than to focus on one set of data at a time. Given three data sets, it is not valuable to show that there exist three different single-system models that are each consistent with one set of data. The important question is – does the single model that best accounts for all three data sets simultaneously, postulate one or multiple systems of category learning? Third, anyone interested in the computational modeling of category learning should look seriously towards cognitive neuroscience as a way to add more constraints to the existing models, and as a mechanism for building bridges to other related areas of cognitive science.

Conclusions
Much work remains to be undertaken in order to understand human category learning. For example, owing to a paucity of data, several of the conclusions in Table 1 must be considered tentative. Nevertheless, the cognitive neuroscience revolution has led to some explosive progress within the past few years. Most important of all is the increasingly impressive evidence that humans might use different neural circuits to learn different types of category structures. In particular, different brain

> **Questions for future research**
>
> - Are there multiple category-learning systems in the human brain, or is there a single system that might use different types of memory in different category-learning tasks?
> - Regardless of how many systems there are, what is the role of medial temporal lobe structures in human category learning?
> - If there are multiple systems, how do they interact? Do they operate independently? Do they compete or do they cooperate?
> - What is the neural basis of experienced categorization? Is it purely cortical? If experienced categorization does have a different neural basis from novice category learning, then what processes mediate the transition?
> - Are there hemispheric asymmetries in the neural basis of category learning?

regions are implicated according to whether the category-learning task involves rules, information integration, or prototype distortion. At the very least, this finding cautions researchers to be extremely careful about the type of categorization experiments they run, because their conclusions could change dramatically with seemingly minor changes in category structure.

References

1 Ashby, F.G. and Lee, W.W. (1993) Perceptual variability as a fundamental axiom of perceptual science. In *Foundations of Perceptual Theory* (Masin, S.C., ed.), pp. 369–399, Elsevier

2 Ashby, F.G. and Maddox, W.T. (1998) Stimulus categorization. In *Handbook of Perception and Cognition: Measurement, Judgment, and Decision Making* (Birnbaum, M.H., ed.), pp. 251–301, Academic Press

3 Humphreys, G.W. and Forde, E.M.E. Hierarchies, similarity and interactivity in object recognition: on the multiplicity of 'category-specific' deficits in neuropsychological populations. *Behav. Brain Sci.* (in press)

4 Gaffan, D. and Heywood, C.A. (1993) A spurious category-specific visual agnosia for living things in normal human and nonhuman primates. *J. Cogn. Neurosci.* 5, 118–123

5 Rolls, E.T. *et al.* (1977) Activity of neurones in the inferotemporal cortex of the alert monkey. *Brain Res.* 130, 229–238

6 Ashby, F.G. *et al.* (1998) A neuropsychological theory of multiple systems in category learning. *Psychol. Rev.* 105, 442–481

7 Grant, D.A. and Berg, E.A. (1948) Behavioral analysis of degree of reinforcement and ease of shifting to new responses in a Weigl-type card-sorting problem. *J. Exp. Psychol.* 38, 404–411

8 Smith, E.E. and Medin, D.L. (1981) *Categories and Concepts*, Harvard University Press

9 Ashby, F.G. and Gott, R.E. (1988) Decision rules in the perception and categorization of multidimensional stimuli. *J. Exp. Psychol. Learn. Mem. Cognit.* 14, 33–53

10 Estes, W.K. (1994) *Classification and Cognition*, Oxford University Press

11 Medin, D.L. and Schaffer, M.M. (1978) Context theory of classification learning. *Psychol. Rev.* 85, 207–238

12 Nosofsky, R.M. (1986) Attention, similarity, and the identification-categorization relationship. *J. Exp. Psychol. Gen.* 115, 39–57

13 Ashby, F.G. and Maddox, W.T. (1992) Complex decision rules in categorization: contrasting novice and experienced performance. *J. Exp. Psychol. Hum. Percept. Perform.* 18, 50–71

14 Maddox, W.T. and Ashby, F.G. (1993) Comparing decision bound and exemplar models of categorization. *Percep. Psychophys.* 53, 49–70

15 Posner, M.I. and Keele, S.W. (1968) On the genesis of abstract ideas *J. Exp. Psychol.* 77, 353–363

16 Homa, D. *et al.* (1981) Limitations of exemplar-based generalization and the abstraction of categorical information. *J. Exp. Psychol. Hum. Learn. Mem.* 7, 418–439

17 Brown, R.G. and Marsden, C.D. (1988) Internal versus external cues and the control of attention in Parkinson's disease. *Brain* 111, 323–345

18 Cools, A.R. *et al.* (1984) Cognitive and motor shifting aptitude disorder in Parkinson's disease. *J. Neurol. Neurosurg. Psychiatry* 47, 443–453

19 Kolb, B. and Whishaw, I.Q. (1990) *Fundamentals of Human Neuropsychology*, 3rd edn, Freeman

20 Robinson, A.L. *et al.* (1980) The utility of the Wisconsin Card Sorting Test in detecting and localizing frontal lobe lesions. *J. Consult. Clin. Psychol.* 48, 605–614

21 Janowsky, J.S. *et al.* (1989) Cognitive impairment following frontal lobe damage and its relevance to human amnesia. *Behav. Neurosci.* 103, 548–560

22 Leng, N.R. and Parkin, A.J. (1988) Double dissociation of frontal dysfunction in organic amnesia. *Br. J. Clin. Psychol.* 27, 359–362

23 Maddox, W.T. and Filoteo, J.V. Striatal contribution to category learning: quantitative modeling of simple linear and complex non-linear rule learning in patients with Parkinson's disease. *J. Int. Neuropsychol. Soc.* (in press)

24 Filoteo, J.V. *et al.* A possible role of the striatum in linear and nonlinear categorization rule learning: evidence from patients with Huntington's disease. *Behav. Neurosci.* (in press)

25 Filoteo, J.V. *et al.* Quantitative modeling of category learning in amnesic patients. *J. Int. Neuropsychol. Soc.* (in press)

26 Knowlton, B.J. *et al.* (1994) Probabilistic classification learning in amnesia. *Learn. Mem.* 1, 106–120

27 Knowlton, B.J. *et al.* (1996) A neostriatal habit learning system in humans. *Science* 273, 1399–1402

28 Knowlton, B.J. *et al.* (1996) Dissociations within nondeclarative memory in Huntington's disease. *Neuropsychology* 10, 538–548

29 Knowlton, B.J. *et al.* (1992) Intact artificial grammar learning in amnesia: Dissociation of classification learning and explicit memory for specific instances. *Psychol. Sci.* 3, 172–179

30 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

31 Meulemans, T. *et al.* (1998) Preserved artificial grammar learning in Parkinson's disease. *Brain Cognit.* 37, 109–112

32 Smith, E.E. *et al.* (1998) Alternative strategies of categorization. *Cognition* 65, 167–196

33 Waldron, E.M. and Ashby, F.G. The effects of concurrent task interference on category learning: evidence for multiple category learning systems. *Psychonomic Bull. Rev.* (in press)

34 Rao, S. *et al.* (1997) Functional MRI evidence for subcortical participation in conceptual reasoning skills. *NeuroReport* 8, 1987–1993

35 Posner, M.I. and Petersen, S.E. (1990) Attention systems in the human brain. *Annu. Rev. Neurosci.* 13, 25–42

36 Goldman-Rakic, P.S. (1987) Circuitry of the prefrontal cortex and the regulation of behavior by representational knowledge. In *Handbook of Physiology* (Plum, F. and Mountcastle, V., eds), pp. 373–417, American Physiological Society

37 Goldman-Rakic, P.S. (1995) Cellular basis of working memory. *Neuron* 14, 477–485

38 Elliott, R. *et al.* (1999) Ventromedial prefrontal cortex mediates guessing. *Neuropsychologia* 37, 403–411

39 Winocur, G. and Eskes, G. (1998) Prefrontal cortex and caudate nucleus in conditional associative learning: dissociated effects of selective brain lesions in rats. *Behav. Neurosci.* 112, 89–101

40 van Domburg, P.H.M.F. and ten Donkelaar, H.J. (1991) *The Human Substantia Nigra and Ventral Tegmental Area*, Springer Verlag

41 Poldrack, R.A. *et al.* (1999) Striatal activation during acquisition of a cognitive skill. *Neuropsychology* 13, 564–574

42 Reber, P.J. *et al.* (1998) Contrasting cortical activity associated with category memory and recognition memory. *Learn. Mem.* 5, 420–428