

HUMAN CATEGORY LEARNING

F. Gregory Ashby

*Department of Psychology, University of California, Santa Barbara,
California 93106; email: ashby@psych.ucsb.edu*

W. Todd Maddox

*Department of Psychology, University of Texas, Austin, Texas 78712;
email: maddox@psyvax.psy.utexas.edu*

Key Words exemplar, prototype, decision bound, multiple systems, striatum

■ **Abstract** Much recent evidence suggests some dramatic differences in the way people learn perceptual categories, depending on exactly how the categories were constructed. Four different kinds of category-learning tasks are currently popular—rule-based tasks, information-integration tasks, prototype distortion tasks, and the weather prediction task. The cognitive, neuropsychological, and neuroimaging results obtained using these four tasks are qualitatively different. Success in rule-based (explicit reasoning) tasks depends on frontal-striatal circuits and requires working memory and executive attention. Success in information-integration tasks requires a form of procedural learning and is sensitive to the nature and timing of feedback. Prototype distortion tasks induce perceptual (visual cortical) learning. A variety of different strategies can lead to success in the weather prediction task. Collectively, results from these four tasks provide strong evidence that human category learning is mediated by multiple, qualitatively distinct systems.

CONTENTS

INTRODUCTION	150
EARLY CATEGORY-LEARNING THEORIES	152
CATEGORY-LEARNING TASKS	152
RULE-BASED TASKS	155
Introduction	155
Neuropsychological Patient Data	156
Neuroimaging Data	156
Theories of Rule-Based Category Learning	157
INFORMATION-INTEGRATION TASKS	158
Category-Learning Limits in Information-Integration Tasks	159
Neuropsychological Patient Data	160
Neuroimaging Data	160
Theories of Information-Integration Category Learning	161

DISSOCIATIONS BETWEEN RULE-BASED AND INFORMATION-INTEGRATION CATEGORY LEARNING	163
PROTOTYPE DISTORTION TASKS	165
Neuropsychological Patient Data	166
Neuroimaging Data	166
Theories of Prototype Distortion Learning	167
WEATHER PREDICTION TASK	169
Task and Individual Differences Analysis	170
Neuropsychological Patient and Neuroimaging Studies	171
CONCLUSIONS	172

INTRODUCTION

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 lacking this ability would quickly become extinct.

Given the important role that categorization plays in our day-to-day lives, it is not surprising that there is a huge and old literature on the perceptual, cognitive, and neurobiological processes that mediate this vital skill. This article surveys that literature, with an emphasis on discoveries and developments during the past 10 years. The past decade has seen exciting and profound changes in categorization research. Two important new themes have fundamentally changed the field. First, there has been huge attention, both theoretical and empirical, devoted to the question of whether human category learning is mediated by a single system, or by multiple, qualitatively distinct learning systems. Second, the categorization field has whole-heartedly embraced the cognitive neuroscience revolution. Not only has the past 10 years seen a massive increase in neuropsychological and neuroimaging research on category learning, but this new knowledge also has permeated into, and significantly sharpened, the core theories in the field. These twin themes—multiple systems and cognitive neuroscience—organize and motivate much of the present review.

The categorization literature is enormous, and no single article could survey it all. Thus, we focus on a restricted subset of the entire literature. In particular, this article focuses on how humans learn perceptual categories. By restricting ourselves in this way, we must necessarily ignore a number of large and interesting subliterations. First, we do not discuss category learning in nonhuman animals (Smith et al. 2004). The animal literature is especially relevant to understanding the neural basis of human category learning. Unfortunately, however, this literature is fractionated, largely because, within the behavioral neuroscience literature at least, relatively few papers address animal categorization *per se*. Rather, the relevant results come from a wide variety of tasks and phenomena (e.g., discrimination learning, memory systems, long-term potentiation). For this reason, perhaps, we know of no

recent comprehensive review (although, for reviews of the behavioral literature, see Vauclair 2002 or the 2002 special issue of the *Journal of the Experimental Analysis of Behavior*).

Second, our focus on learning prevents us from considering 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, many neuropsychological groups that are impaired in category learning (e.g., frontal patients and Parkinson's disease patients) do not lose old, familiar categories (e.g., fruits and tools). Similarly, there is no evidence that people who lose a familiar category (i.e., who develop a category-specific agnosia) develop any general category-learning deficits. Readers interested in the representation of highly learned categories are referred to any of several excellent reviews of the category representation literature (e.g., Cree & McRae 2003, Humphreys & Forde 2001, Joseph 2001).

Third, we discuss how people learn new categories, but not how they use this new information in other cognitive tasks. For example, people use categorical information to make inferences about unobserved features of a stimulus, to facilitate decision making, and to problem solve (e.g., Lewandowsky et al. 2002, Markman & Ross 2003).

Finally, our focus is on perceptual categories rather than concepts. By "perceptual category," we mean a collection of similar objects belonging to the same group. Although the term "concept" often is used interchangeably with category, we refer to a concept as a group of related ideas. For example, the set of all X rays displaying a tumor forms a perceptual category, whereas the many varieties of religious experience form a concept. Although many of the results reviewed below are relevant to understanding both perceptual categorization and concept formation, the representation of categories and concepts is likely different, and a review of both literatures is beyond the scope of this article. Readers interested in concepts are referred to any of several excellent recent reviews (e.g., Barsalou 2003, Murphy 2002).

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 was paid to the manner in which these arbitrary categories were created. However, much recent evidence suggests some dramatic differences in the way people learn such categories, depending on exactly how the categories are constructed. In fact, these differences are so profound that we take the unusual step of organizing all the research that we review by the type of task that was used. Toward this end, we focus on four different kinds of category learning tasks—rule-based tasks, information-integration tasks, prototype distortion tasks, and the so-called weather

prediction task. The next section briefly reviews some important early category-learning theories. The third section describes the four basic tasks, and then sections four through eight review results from each of these tasks. Finally, we close with some general conclusions.

EARLY CATEGORY-LEARNING THEORIES

Many theories of human category learning have been proposed. The early theories virtually all assumed that humans have a single category-learning system that they use to learn all types of categories (for an exception, see Brooks 1978). A few of these theories are still important in the sense that they continue to motivate significant amounts of new research. This section briefly introduces the most important of these theories. Other sources should be consulted for a more complete discussion, and for a more thorough review of category-learning theories (e.g., Ashby & Maddox 1998, Estes 1994, Smith & Medin 1981).

Prototype theory assumes that category learning is equivalent to learning the category prototype. When an unfamiliar stimulus is then encountered, it is assigned to the category with the most similar prototype (Homa et al. 1981; Posner & Keele 1968, 1970; Reed 1972; Rosch 1973, 1975; Smith & Minda 1998).

Exemplar theory assumes that category learning is a process of learning about the exemplars that belong to the category. When an unfamiliar stimulus is encountered, its similarity is computed to the memory representation of every previously seen exemplar from each potentially relevant category. The stimulus is then assigned to the category for which the sum of these similarities is greatest (Brooks 1978; Estes 1986, 1994; Hintzman 1986; Lamberts 2000; Medin & Schaffer 1978; Nosofsky 1986).

Decision bound theory assumes subjects partition the stimulus space into response regions. When presented with an unfamiliar stimulus, the subject determines which region the percept is in, and then emits the associated response. The partition between regions associated with competing responses is called the decision bound. Category learning is the process of either learning the decision bound or else of learning the regions associated with each response (Ashby & Gott 1988, Ashby & Townsend 1986, Maddox & Ashby 1993).

CATEGORY-LEARNING TASKS

As mentioned above, this article focuses on four different kinds of category learning tasks. Rule-based tasks are those in which the categories can be learned via some explicit reasoning process. Frequently, the rule that maximizes accuracy (i.e., the optimal strategy) is easy to describe verbally (Ashby et al. 1998). 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. However, there is no requirement that the rule that maximizes accuracy (i.e., the optimal rule) in rule-based tasks is one-dimensional. For example, 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 because it is easy to describe verbally.

Information-integration category learning tasks are those in which accuracy is maximized only if information from two or more stimulus components (or dimensions) is integrated at some predecisional stage (Ashby & Gott 1988). Perceptual integration could take many forms—from computing a weighted linear combination of the dimensional values to treating the stimulus as a gestalt. In many cases, the optimal strategy in information-integration tasks is difficult or impossible to describe verbally (Ashby et al. 1998). Real-world examples of information-integration tasks are common. For example, deciding whether an X ray shows a tumor requires years of training, and expert radiologists are only partially successful at describing their categorization strategies.

Examples of rule-based and information-integration categories that might be used in experimental research are shown in Figure 1. In both cases, the two contrasting categories are composed of circular sine-wave gratings (i.e., disks in which luminance varies sinusoidally). Examples are shown in Figure 1*a*. The disks are all of equal diameter, but they differ in spatial frequency (i.e., the frequency of the sine wave) and (sine-wave) orientation. Each symbol in Figures 1*b* and 1*c* denotes the spatial frequency and orientation of a sine-wave grating. Category A exemplars are denoted by pluses and category B exemplars are denoted by circles. In each condition, there are two distinct categories that do not overlap, so perfect accuracy is possible. Also shown in Figures 1*b* and 1*c* are the decision bounds that maximize categorization accuracy. In the rule-based task (Figure 1*b*), the optimal bound, denoted by the vertical line in Figure 1*b*, requires observers to attend to spatial frequency and ignore orientation. This bound has a simple verbal description: “Respond A if the bars are thick and B if they are thin.” In the information-integration task (Figure 1*c*), which was generated by rotating the rule-based categories by 45°, equal attention must be allocated to both stimulus dimensions. In this task, there is no simple verbal description of the optimal decision bound.

Note that we use the word “rule” more narrowly than is common in the psychological literature, where it is often used to refer to any strategy from an explicit reasoning process to any algorithm that can be expressed formally. In particular, we define “rule-based strategy” narrowly to refer specifically to an explicit reasoning process. Note that according to this criterion, there is no limit on the complexity of the optimal rule in rule-based tasks. However, as the complexity of the optimal rule increases, its salience decreases and it becomes less likely that observers will learn the associated categories through an explicit reasoning process. Thus, the boundary is fuzzy between rule-based and information-integration tasks. Tasks in which the optimal rule is one-dimensional are unambiguously rule-based (at least with separable stimulus dimensions), and tasks in which the optimal rule is

significantly more complex than a conjunction rule almost never are rule-based. In between, the classification is not so clear-cut.

It is also important to emphasize that the terms “rule-based” and “information-integration” make no assumptions about how people learn these different category structures in any particular application. For example, there is evidence that pigeons can learn both types of category structures (Herbranson et al. 1999), but no one would claim that they learn rule-based categories via an explicit reasoning process. The question of how people learn rule-based and information-integration categories is strictly empirical. As such, this particular classification of categorization tasks is useful only because there are many interesting empirical dissociations between the two tasks (e.g., Ashby et al. 1999, 2002, 2003; Ashby & Waldron 1999; Maddox et al. 2003).

Prototype distortion tasks are a third type of category-learning task in which each category is created by first constructing a category prototype (Posner & Keele 1968, 1970). The other exemplars of the category are then created by randomly distorting the prototype. In the most popular prototype distortion task, each stimulus is a random pattern of (often nine) dots. One pattern is selected as the category prototype and then the other category exemplars are created by randomly perturbing the location of each dot in the prototype. Examples are shown in Figure 2.

The final category-learning task we consider is the so-called weather prediction task. Stimuli in this task are tarot cards; each displays a unique geometric pattern. The subject’s task is to decide if the particular constellation of cards that is shown signals “rain” or “sun.” The actual outcome is determined by a probabilistic rule based on the individual cards.




Figure 2 Stimuli that might be used in a prototype distortion category-learning task.

RULE-BASED TASKS

Introduction

As mentioned above, rule-based tasks are those in which it is easy for subjects to describe the optimal strategy verbally. In general, several conditions must be met before a verbal description is possible. First, a semantic label must correspond to each of the stimulus properties that are relevant to the decision. In the Figure 1*b* rule-based task, the critical stimulus feature has the semantic label “width.” Second, the subject must be able to attend selectively to each relevant stimulus property. For example, it is possible to verbalize a rule such as “Respond A if the saturation of the color patch is high, respond B if saturation is low.” Even so, people are not good at attending selectively to saturation and ignoring irrelevant variations in hue and brightness, so it is unlikely that people would spontaneously experiment with such rules. In the selective attention literature, a stimulus feature that can be attended to selectively is said to be separable from the other stimulus features, whereas features for which it is difficult or impossible to attend selectively are said to be integral. A large and old literature is devoted to this topic (Ashby & Maddox 1994, Ashby & Townsend 1986, Garner 1974, Lockhead 1966, Shepard 1964).

The third critical property necessary for easy verbalization is that the rule for combining information from the relevant stimulus features is itself verbalizable. In general, this requires that separate decisions are first made about the level of each feature, and then these separate decisions are combined using logical operations, such as “and” and “or.” Perhaps the most obvious example is a conjunction rule of the sort: “Respond A if the bars are wide *and* the orientation is steep; otherwise respond B.” Note that to apply this rule the subject must first decide if the bars are narrow or wide and if the orientation is shallow or steep. Next, the outcomes of these two decisions are combined using the word “and.” Thus, information from the various relevant stimulus dimensions is combined after decisions are first made about each dimension. This is in contrast to information-integration tasks in which the raw perceptual information from the relevant stimulus dimensions is combined before any decisions are made.

Other rule-based strategies that require combining decisions from separate dimensions include disjunctive and exclusive-or rules. There is no doubt that healthy adults can learn these rules without much difficulty, at least if they are given regular feedback about their response accuracy (e.g., Salatas & Bourne 1974). However, it is also quite clear that such combination rules are much less salient than simple one-dimensional rules, in the sense that people rarely experiment with such rules unless compelled to do so by the feedback they receive (Alfonso-Reese 1996, Ashby et al. 1998).

Virtually all category-learning tasks used in neuropsychological assessment are rule-based, including the widely known Wisconsin Card Sorting Test (WCST; Heaton 1981). Stimuli in the WCST are cards containing geometric patterns that vary in color, shape, and symbol number, and in all cases, the correct categorization

rule is one-dimensional (and easy to describe verbally). Perseverative errors on the WCST are a classic symptom of frontal dysfunction (e.g., Kimberg et al. 1997).

Neuropsychological Patient Data

Although categorization has been studied in many different neuropsychological groups, the most extensive data come primarily from studies with three different groups: (a) patients with frontal lobe lesions, (b) patients suffering from a disease of the basal ganglia, typically either Parkinson's or Huntington's disease, and (c) patients with amnesia. Within this latter group, the most theoretically interesting are those whose amnesia was caused by damage to the medial temporal lobes. In almost all cases, however, studies with amnesiacs include a wide variety of patients, typically including some with Korsakoff's syndrome and some with medial temporal lobe amnesia.

One characteristic feature of the neuropsychological literature on category learning is its inconsistency. For each major patient group, some studies report deficits and some do not. However, as we will see, when the existing studies are partitioned according to the type of task that was used, the discrepancies largely disappear.

As mentioned above, perseverative responding on the WCST is among the most classic of all signs of frontal damage. Not surprisingly then, many studies have shown that frontal patients are impaired at rule-based category learning (see, e.g., Kimberg et al. 1997, Robinson et al. 1980). Another group with well-known deficits in rule-based category learning is Parkinson's disease patients (e.g., Ashby et al. 2003, Brown & Marsden 1988, Cools et al. 1984, Downes et al. 1989). Although later in the disease Parkinson's patients have frontal damage (primarily the result of cell death in the ventral tegmental area), the disease mainly targets the basal ganglia. The region most affected appears to be the head of the caudate nucleus (van Domburg & ten Donkelaar 1991), which is reciprocally connected to the prefrontal cortex. Thus, the rule-based category-learning deficits of frontal and Parkinson's disease patients are consistent with the hypothesis that rule-based category learning is mediated, in part, by frontal-striatal circuits (Ashby et al. 1998).

In contrast to frontal and basal ganglia disease patients, several studies have reported that amnesiacs with medial temporal lobe damage are normal in rule-based category learning (Janowsky et al. 1989, Leng & Parkin 1988). An obvious possibility is that many rule-based tasks are simple enough (e.g., the WCST) that working memory is sufficient for subjects to keep track of which alternative rules they have tested and rejected. If so, then a natural prediction is that medial temporal lobe amnesiacs should be impaired in complex rule-based tasks (e.g., when the optimal rule is disjunctive).

Neuroimaging Data

A number of neuroimaging studies have used the WCST or a rule-based task similar to the WCST. All of these have reported task-related activation in prefrontal cortex,

most have reported activation in the head of the caudate nucleus, and at least one has reported task-related activation in the anterior cingulate (Konishi et al. 1999, Lombardi et al. 1999, Rao et al. 1997, Rogers et al. 2000, Volz et al. 1997). 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 (Posner & Petersen 1990) and working memory (Goldman-Rakic 1987, 1995), both of which are likely to be critically 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 (Elliott et al. 1999). Third, lesion studies in rats implicate the dorsal caudate nucleus in rule switching (Winocur & Eskes 1998). Fourth, of course, are the neuropsychological data reviewed above, which show that patient groups with damage to any of these structures are impaired in rule-based tasks.

Theories of Rule-Based Category Learning

Theories of rule-based category learning can be classified according to whether they assume that learning in rule-based tasks is not fundamentally different from other category-learning tasks, or whether they assume that rule-based learning is special. Most of the attempts to account for the results of rule-based category learning with a single system model have been by exemplar theorists. According to exemplar theory, rule-based tasks, in general, are no different from any other type of category-learning task. However, one-dimensional rules, like the one depicted in Figure 1*b*, are unique in that they encourage selective attention to a single dimension, which in turn dramatically affects the stimulus-exemplar similarity computations (Kruschke 1992, Nosofsky 1991, Nosofsky et al. 1989). In particular, increasing attention to a dimension will serve to increase the perceived differences on that dimension. As a result, the perceived separation between the Figure 1*b* rule-based categories will tend to be greater than the perceived separation between the Figure 1*c* information-integration categories. It is in this way that exemplar theory is able to account for the difficulty differences between the two tasks.

The earliest theory that applies directly to rule-based tasks is the so-called classical theory of categorization (e.g., Bruner et al. 1956, Smith & Medin 1981), which assumes that every category is represented by a set of necessary and sufficient features. When a stimulus is presented for categorization, the subject is assumed to retrieve the feature list of the relevant categories and then test whether the stimulus features match one of these feature lists. This theory accounts for performance in many rule-based tasks. For example, category A in Figure 1*b* is defined by the necessary and sufficient feature “thick bars.” Similarly, the conjunction rule: “Respond A if the bars are thick and the orientation is shallow” is equivalent to the necessary and sufficient features “thick bars and shallow orientation.” On the

other hand, in some rule-based tasks the optimal rule cannot be expressed as a set of necessary and sufficient conditions. For example, given appropriate feedback, subjects can learn disjunctive-or rules such as: "Respond A if the bars are thick and the orientation is shallow or if the bars are thin and the orientation is steep" (Salatas & Bourne 1974). Thus, it was recognized long ago that the classical theory is incomplete, even for the restricted set of rule-based tasks (e.g., Ashby & Maddox 1998, Smith & Medin 1981).

As originally proposed, classical theory was meant to apply to all categorization tasks. There have been several attempts to modernize the theory. Each of these attempts, however, has assumed that rule-based category learning is only one of several category-learning systems that humans have available (Ashby et al. 1998, Brooks 1978, Erickson & Kruschke 1998, Nosofsky et al. 1994b). The rule-based components of these various models are similar. The model that perhaps is best developed, and the only one with a neuropsychological basis, was proposed by Ashby et al. (1998) as part of the COVIS (COmpetition between Verbal and Implicit Systems) model of category learning (which also includes a procedural learning component).

The COVIS explicit system assumes that rule-based category learning is mediated primarily by an explicit, hypothesis-testing system that depends heavily on working memory and executive attention. The idea is that candidate rules are stored in working memory during the time they are being tested. COVIS assumes that subjects will continue to use the active rule until feedback or other evidence disconfirms its validity. At this point, a new rule must be instantiated. COVIS assumes that activating a new rule requires two separate processes. First, a new candidate rule must be identified or selected, and second, attention must be switched from the old rule to the new rule. The probability that any given rule will be instantiated is determined by its reinforcement history (which determines its overall salience), the tendency of the subject to select novel hypotheses (which is assumed to depend on cortical dopamine levels; Ashby et al. 1999), and the tendency of the subject to perseverate (which is assumed to depend on basal ganglia dopamine levels). Ashby and his colleagues proposed, and presented evidence in support of the hypothesis, that the selection operation is mediated cortically, by the anterior cingulate and possibly also by the prefrontal cortex, and that switching is mediated by the head of the caudate nucleus. A review of this evidence is beyond the scope of this article (see Ashby et al. 1998, 1999).

INFORMATION-INTEGRATION TASKS

As mentioned above, information-integration category-learning tasks are defined as those in which accuracy is maximized only if information from two or more stimulus components (or dimensions) is integrated at some predecisional stage. Typically, the optimal rule is difficult or impossible to describe verbally (Ashby et al. 1998).

Category-Learning Limits in Information-Integration Tasks

An important theoretical and practical question is whether there are limits on the complexity of information-integration category structures that can be learned. One of the first research efforts to address this question focused on comparing learning in linearly and nonlinearly separable categories. A pair of categories is linearly separable if optimal performance can be achieved by making category decisions based on the magnitude of a linear combination of dimensional values (or equivalently, if a linear decision bound is optimal). Categories are nonlinearly separable if optimal performance depends on a nonlinear combination of dimensional values (i.e., if the optimal bound is nonlinear). Prototype theory predicts that nonlinearly separable category structures should be impossible to learn, at least if each category contains only a single prototype (Ashby & Gott 1988), whereas exemplar models predict no consistent advantage for linearly or nonlinearly separable categories. Medin & Schwanenflugel (1981) compared linearly and nonlinearly separable category learning using a small number of stimuli constructed from binary-valued dimensions and found no advantage for linearly separable structures. Ashby and colleagues (Ashby & Gott 1988; Ashby & Maddox 1990, 1992) compared linearly and nonlinearly separable category learning using a large number of stimuli constructed from continuous-valued dimensions. They found a consistent advantage for linearly separable categories, with nonlinearly separable category learning often requiring numerous experimental sessions. Despite their greater difficulty, however, the fact that people can learn nonlinearly separable categories effectively falsifies the standard prototype-theory account of information-integration category learning. Even so, there is evidence that early in learning people may abstract prototypes in some information-integration tasks (Minda & Smith 2001; Smith & Minda 1998, 2002).

McKinley & Nosofsky (1995) examined category learning when the categories were composed of a large number of unique exemplars sampled from mixtures of bivariate normal distributions. Although some subjects were able to learn the categories fairly well, a large number failed to learn even after seeing nearly 4000 exemplars over a full week of training. These data present a challenge to exemplar theory because after such extensive training, exemplar theory predicts nearly optimal performance, no matter what the category structures (Ashby & Alfonso-Reese 1995). In a related study, Ashby et al. (2001) compared two-category learning where the optimal decision bound was quadratic, with four-category learning where the optimal bound separating each pair of categories was quadratic. As in previous studies (e.g., Ashby & Maddox 1992), learning was good in the two-category case, and the best fitting model assumed subjects used quadratic bounds. In the four-category case, however, learning was worse, and the best fitting model assumed subjects used suboptimal linear bounds to separate each pair of categories. Taken together, these two data sets suggest that there is an upper bound on the complexity of information-integration category structures that can be learned (in a reasonable amount of time) and that this upper bound is greater than a single quadratic curve but less than a set of such curves.

Neuropsychological Patient Data

Over the past several years, a number of studies of information-integration category learning have been conducted in brain-damaged populations. The focus has been on patients with medial temporal lobe amnesia or striatal damage (e.g., patients with Parkinson's disease or Huntington's disease). Filoteo et al. (2001b) tested the ability of amnesiacs to learn a highly nonlinear information-integration rule when the categories were normally distributed and a large number of unique stimuli were sampled from each category. Over the full 600 trials of the experiment, the performance of amnesiacs and controls was equivalent. One patient and one control returned for a second session on the following day. During the first block of trials on the second day, the amnesiac and control again showed equivalent performance, and in fact, performance during the first block of the second session was slightly better than during the final block of trials from the first session. Some researchers have argued that amnesiacs learn categorization rules using working or short-term memory processes (Nosofsky & Zaki 1998, Palmeri & Flanery 1999). The day 2 results of Filoteo et al. (2001b) argue against this possibility. Instead, these findings indicate that the categorization rule was retained over the one-day delay period and argue strongly against the hypothesis that working memory or explicit declarative memory mediates information-integration category learning.

Filoteo et al. (2001a) and Maddox & Filoteo (2001) tested the ability of Huntington's disease and Parkinson's disease patients to learn the same category structures used by Filoteo et al. (2001b). Over the full 600 trials of the experiment, both patient groups showed a consistent performance decrement, suggesting an involvement of the striatum in nonlinear information-integration category learning. On the other hand, Ashby et al. (2003) found that Parkinson's disease patients learned as well as an age-matched control group in an information-integration task with linearly separable categories. More recently, Filoteo et al. (2004) compared the ability of Parkinson's disease patients to learn a linear and a nonlinear information-integration rule. The linear results replicated the Ashby et al. (2003) results—that is, the Parkinson's disease patients were not impaired in learning linearly separable categories. On the other hand, the same patients were impaired in the nonlinear condition, but only later in training. Thus, these studies suggest that Parkinson's disease subjects are impaired in information-integration tasks, but only if the category structures are complex (as, e.g., when the categories are nonlinearly separable).

Neuroimaging Data

To date, only one neuroimaging study of information-integration category learning has been conducted (although see the section on the weather prediction task). Seger & Cincotta (2002) reported significant striatal and lateral occipital activation by a group of subjects who had already had extensive training in the task.

Theories of Information-Integration Category Learning

There are a number of successful theories of information-integration category learning. These can be classified into two types: parametric and nonparametric (see Ashby & Alfonso-Reese 1995 for a detailed discussion). Parametric classifiers assume either that the categories have a specific type of structure (e.g., normal distributions) or that the categorization boundary has a specific functional form (e.g., linear). Nonparametric classifiers make no assumptions about category structure or categorization boundaries.

Simple prototype models (Reed 1972, Smith & Medin 1981) are parametric because they assume a linear decision bound (Ashby & Gott 1988). For this reason, as mentioned above, they can be rejected as a general theory of information-integration category learning because humans can learn nonlinear decision bounds (see, e.g., Ashby & Maddox 1992, Medin & Schwanenflugel 1981).

Exemplar models are nonparametric because they assume that every exemplar presented is stored in memory along with the appropriate category label. Because all category information is retained (although perhaps in a degraded form), exemplar models predict that, under very general conditions, subjects should eventually respond almost optimally, no matter how complex the categories (Ashby & Alfonso-Reese 1995). For this reason, despite the enormous success of this class of models, the failure of subjects to respond optimally in complex information-integration tasks (e.g., Ashby et al. 2001, McKinley & Nosofsky 1995) suggests that exemplar models may be too powerful. In addition, the only current neurobiological hypotheses about the exemplar-memory process attach a critical role to the hippocampus (Pickering 1997). As such, the finding that medial temporal lobe amnesiacs are relatively normal at information-integration category learning is problematic for the hypothesis that exemplar theory is adequate as a general theory of categorization in information-integration tasks.

Decision bound models can be parametric or nonparametric depending on whether they assume subjects learn decision bounds (parametric) or assign responses to regions (nonparametric). In this latter case, the bound is simply the partition between regions associated with contrasting responses.

Ashby & Waldron (1999) conducted a critical test of whether category learning in information-integration tasks is parametric or nonparametric. Previous research (reviewed above) showed that people can learn either linear or nonlinear (e.g., quadratic) information-integration decision bounds. If humans are parametric classifiers, then some categorical information must signal whether they should use a linear bound or a nonlinear bound. Ashby & Waldron (1999) constructed categories in which all statistical information that could be readily estimated (e.g., means, variances, and covariances) signaled that a parametric classifier should use a linear decision bound, but for which the optimal bound was quadratic. In a second condition, the statistical information signaled that a nonlinear bound should be used, but the optimal bound was linear. All known parametric classifiers predict that subjects will use a decision bound of the wrong type in these two conditions.

Yet, the data of every subject (who did not use a rule-based strategy) were best fit by a bound of the optimal type. These results provide strong evidence against all known parametric classifiers, including prototype models, and decision bound models that assume people learn decision bounds. At the same time, the Ashby & Waldron (1999) results support nonparametric classifiers, such as exemplar models and decision bound models that assume people learn to assign responses to regions of perceptual space.

As a model of the latter type, Ashby & Waldron (1999) proposed a nonparametric decision bound model called the striatal pattern classifier. Other nonparametric decision-bound-type models include Anderson's (1991) rational model, and Love et al.'s (2004) SUSTAIN (Supervised and Unsupervised Stratified Adaptive Incremental Network) model. Each of these models can be loosely described as multiple prototype (or cluster) models because perceptually similar category exemplars tend to be grouped or clustered together. For the most part, all three models can account for the observed complexity limits on the learning of information-integration category structures, but the rational and SUSTAIN models make no attempt to account for the neuropsychological and neuroimaging data, and neither proposes a neurobiological account of information-integration category learning.

The striatal pattern classifier, on the other hand, offers a computational model of the procedural-learning-based system proposed in COVIS, and it proposes a neurobiological interpretation. In brief, the model assumes that information-integration category learning is mediated primarily within the tail of the caudate nucleus (for visual stimuli). Two key neurophysiological features make this system a good candidate for information-integration category learning. First, all of visual cortex (except area V1) projects directly to the tail of the caudate, and these projections are characterized by massive convergence (i.e., of approximately 10,000 to 1; Wilson 1995). This convergence causes the decision space represented in the caudate nucleus to have much lower resolution than the perceptual space represented in visual cortex, and is assumed to account for the complexity limits on information-integration category learning reviewed above. Second, the tail of the caudate receives dopaminergic input from the substantia nigra that is widely thought to serve as a reward-mediated feedback signal (e.g., Schultz 1992, Wickens 1993). The idea is that an unexpected reward causes dopamine to be released from the substantia nigra into the tail of the caudate, and that the presence of this dopamine strengthens recently active synapses. The next section reviews a number of empirical results that are thought to be directly attributable to unique features of this feedback system.

The striatal pattern classifier is consistent with much of the neuropsychological and neuroimaging data reviewed above. For example, the model predicts that amnesiacs should show normal information-integration category learning, whereas patients with striatal damage (Parkinson's and Huntington's patients) should be impaired. In addition, the model predicts that there should be striatal activation during information-integration category learning. In addition, although speculative at this point, it seems reasonable to suppose that the learning of complex nonlinear

information-integration category structures requires a higher resolution in the tail-of-the-caudate decision space than the learning of simpler linear information-integration category structures. Because the tail of the caudate is dysfunctional in Parkinson's disease, the model makes the natural prediction that Parkinson's disease patients should be especially impaired in nonlinear information-integration category learning. Each of these predictions was supported by the data reviewed above.

DISSOCIATIONS BETWEEN RULE-BASED AND INFORMATION-INTEGRATION CATEGORY LEARNING

In a seminal study, Shepard et al. (1961; Shepard & Chang 1963) examined category learning in six tasks constructed from different stimulus-category assignments of the same 8 three-dimensional binary-valued stimuli (for replications and extensions, see Nosofsky et al. 1994a, Nosofsky & Palmeri 1996, Smith et al. 2004). These included rule-based, information-integration, and unstructured (memorization) tasks. Results showed that the one-dimensional task was easiest to learn and the unstructured, memorization task was the most difficult, with the other tasks (including the information-integration task) of intermediate difficulty. One weakness of this study is that the structural properties of the categories, such as within-category coherence and between-category discriminability, were not controlled. More recently, a number of studies have compared rule-based and information-integration category learning in a variety of settings where these and other structural properties are controlled. Collectively, these data offer a serious challenge to single-system models.

Many of the studies in question were motivated by the COVIS model of category learning (Ashby et al. 1998). As outlined earlier, COVIS assumes that learning in rule-based tasks is dominated by an explicit, hypothesis-testing system that uses working memory and executive attention and is mediated primarily by the anterior cingulate, the prefrontal cortex, and the head of the caudate nucleus. In contrast, learning in information-integration tasks is assumed to be dominated by an implicit procedural-learning-based system, which is mediated largely within the tail of the caudate nucleus (Ashby et al. 1998, Ashby & Ell 2001, Willingham 1998).

A series of studies attempted to dissociate the processes involved in rule-based and structurally equivalent information-integration category learning by introducing simple experimental manipulations that are predicted by COVIS to affect processing in the procedural-learning system but not the hypothesis-testing system, or vice versa. These tests focused on predictions that the two putative systems would be affected differently by manipulations of the nature and timing of feedback, by changes in the locations of the response keys, and by adding additional demands on working memory and executive attention. The remainder of this section briefly reviews these studies (for a more detailed review, see Maddox & Ashby 2004).

Because the COVIS hypothesis-testing system is under conscious control and has full access to working memory and executive attention, the nature and timing of the feedback signal should not be critical for rule-based category learning. In contrast, a procedural-learning system that is mediated within the tail of the caudate nucleus would not be accessible to conscious awareness and is far removed from working memory.¹ As a result, it would depend more heavily on the nature and timing of the feedback. Several studies tested these predictions. First, observational training was found to be equally effective to traditional feedback training with rule-based categories, but with information-integration categories, a distinct advantage occurred for feedback training (Ashby et al. 2002). During observational training, subjects are informed of the category membership of each stimulus just before it appears, whereas during feedback training, the stimulus is presented and then the category label is shown immediately after the subject responds. Second, some rule-based categories can be learned without feedback of any kind, whereas there is no evidence that information-integration categories can be learned without feedback (Ashby et al. 1999). Third, delaying the feedback by as little as 2.5 seconds after the response significantly interferes with information-integration category learning, but delays as long as 10 seconds have no effect on rule-based learning (Maddox et al.).

A second set of studies tested the prediction that information-integration category learning is mediated largely by a form of procedural learning. The quintessential paradigm for studying procedural learning is the serial reaction time task (Nissen & Bullemer 1987), in which subjects press keys as quickly as possible in response to stimuli that appear in various locations on the screen. A large response time improvement is observed when the stimulus sequence is repeated, even when subjects are unaware that a sequence exists. Willingham et al. (2000) showed that changing the location of the response keys interferes with serial reaction time learning, but that changing the fingers that push the keys does not. Thus, if procedural learning is used in information-integration tasks, then switching the locations of the response keys should interfere with learning, but switching the fingers that depress the keys should not. In fact, Ashby et al. (2003) reported evidence that directly supported this prediction. They also reported that neither manipulation had any effect on rule-based category learning. Maddox et al. (2004) reported a similar sensitivity of information-integration category learning to response location. On half the trials, subjects responded “Yes” or “No” depending on whether the stimulus belonged to category A, and on half the trials they responded “Yes” or “No” depending on whether the stimulus belonged to category B. Thus, there was no consistent mapping of category label to response position. Compared to a standard

¹Crick & Koch (1990, 1995, 1998) offered a cognitive neuroscience theory of conscious awareness that states one can have conscious awareness only of activity in brain areas that project directly to the prefrontal cortex. The caudate nucleus does not project to the prefrontal cortex (it first projects through the globus pallidus and then the thalamus), so the Crick-Koch hypothesis predicts that we are not aware of activity within the caudate nucleus.

control condition, learning was impaired with the information-integration categories, but not with the rule-based categories. These results provide the first direct evidence of procedural learning in perceptual categorization and suggest that the hypothesis-testing system learns abstract category labels, whereas the procedural-learning system learns response positions (for other examples of response effects in categorization, see Barsalou et al. 2003).

A third set of studies tested the prediction that rule-based category learning requires working memory and executive attention, both to select and apply the correct rule and to interpret and process the feedback signal. First, Waldron & Ashby (2001) showed that rule-based category learning was disrupted more than information-integration category learning by the simultaneous performance of a task that required working memory and executive attention (a numerical Stroop task). In a second related study, Maddox et al. (2004) required subjects to alternate categorization trials with trials of a classic memory-scanning task (Sternberg 1966). In one condition, a short delay followed categorization and a long delay followed memory scanning, whereas these delays were reversed in the other condition. Learning of the information-integration categories was unaffected by the location of the short delay, whereas rule-based category learning was significantly worse when the short delay followed categorization. This result supports the hypothesis that feedback processing requires attention and effort in rule-based categorization, but not in information-integration category learning.

PROTOTYPE DISTORTION TASKS

In prototype distortion tasks, the category exemplars are created by randomly distorting a single category prototype. As mentioned above, the most widely known example uses a constellation of dots (often 7 or 9) as the category prototype (see Figure 2 for an example), and the other category members are created by randomly perturbing the spatial location of each dot. These random dot stimuli and categories have been used in dozens of studies (e.g., Homa et al. 1979, 1981; Posner & Keele 1968, 1970; Shin & Nosofsky 1992; Smith & Minda 2002).

Two different types of prototype distortion tasks are popular—(A, B) and (A, not A). In an (A, B) task, subjects are presented a series of exemplars that are each from some category A or from a contrasting category B. The task of the subject is to respond with the correct category label on each trial (i.e., “A” or “B”). An important feature of (A, B) tasks is that the stimuli associated with both responses each have a coherent structure—that is, they each have a central prototypical member around which the other category members cluster. In an (A, not A) task, on the other hand, there is a single central category A and subjects are presented with stimuli that are either exemplars from category A or random patterns that do not belong to category A. The subject’s task is to respond “Yes” or “No” depending on whether the presented stimulus was or was not a member of category A. In an (A, not A)

task, the category A members have a coherent structure, but the stimuli associated with the “not A” (or “No”) response do not. Historically, prototype distortion tasks have been run both in (A, B) form and in (A, not A) form, although (A, not A) tasks are most common.

Neuropsychological Patient Data

Prototype distortion tasks are particularly important because the neuropsychological patient data are profoundly different from those in rule-based or information-integration tasks. In particular, a variety of patients groups that are known to have deficits in rule-based and information-integration tasks show apparently normal prototype distortion learning, at least in (A, not A) designs. This includes patients with Parkinson’s disease (Reber & Squire 1999), schizophrenia (Keri et al. 2001), or Alzheimer’s disease (Sinha 1999, although see Keri et al. 1999). Normal (A, not A) performance has also been shown in patients with amnesia (Knowlton & Squire 1993, Kolodny 1994, Squire & Knowlton 1995). These results must be interpreted with caution, however, because several studies have shown that if category A is created from low-level distortions of the category A prototype, then healthy young adults can learn in (A, not A) tasks without any feedback (i.e., training) at all (Homa & Cultice 1984, Palmeri & Flanery 1999). Thus, it is not yet clear that all these patient groups would learn normally in a difficult (A, not A) task (i.e., one that requires feedback for optimal performance).

At least two studies have compared (A, not A) and (A, B) prototype distortion learning on the same patients—and both studies report the same striking dissociation. Specifically, Sinha (1999) reported normal (A, not A) performance in Alzheimer’s disease patients, but impaired (A, B) performance, and Zaki et al. (2003) reported this same pattern of results with amnesiacs. Sinha (1999) also reported deficits in (A, B) prototype distortion learning in patients with amnesia.

Neuroimaging Data

A handful of neuroimaging studies have used prototype distortion tasks. When interpreting these results, it is vital to consider whether an (A, B) or (A, not A) task was used. As in the purely behavioral studies, the most popular choice has been the (A, not A) task. All of these studies have reported learning-related changes in occipital cortex (Aizenstein et al. 2000; Reber et al. 1998a,b)—in general, reduced occipital activation was found in response to category A exemplars, although Aizenstein et al. (2000) found this reduction only under implicit learning conditions. When subjects were given explicit instructions to learn the A category, increased occipital activation was observed.

Studies that have used (A, B) tasks have reported quite different results. Seger et al. (2000) did report categorization-related activation in occipital cortex, but they also found significant learning-related changes in prefrontal and parietal cortices. Vogels et al. (2002) reported results from a hybrid task in which subjects were to respond “A,” “B,” or “Neither.” Thus, stimuli were created from distortions of

an A prototype or a B prototype, or were just random patterns. Like Seger et al. (2000), Vogels et al. (2002) found prefrontal and parietal activation (although in different foci). However, they also reported task-related activation in orbitofrontal cortex and the neostriatum, and they failed to find any task-related activation in occipital cortex.

Theories of Prototype Distortion Learning

EXEMPLAR VERSUS PROTOTYPE THEORIES The most recent debates between prototype and exemplar theories have focused on prototype distortion tasks. Prototype theory assumes a category is represented as a prototype, and that stimuli are categorized by comparing them to the prototypes of each contrasting category (Homa et al. 1981; Posner & Keele 1968; Reed 1972; Rosch 1973, 1975; Smith & Minda 2001). This theory seems ideally suited to prototype distortion tasks where all category members are simple distortions of a central prototype. Indeed, early results seemed to support this prediction. For example, performance is generally better on the prototype and on exemplars similar to the prototype than on distortions, even when subjects are trained on the distortions but not on the prototype (Homa et al. 1979, 1981; Posner & Keele 1970; Strange et al. 1970). Even so, exemplar theorists showed that exemplar theory was also compatible with these results, and they argued that exemplar theory provides at least as good an account of data from prototype distortion tasks as does prototype theory (Hintzman 1986, Hintzman & Ludlam 1980, Shin & Nosofsky 1992).

Smith & Minda (2001) identified a critical test between exemplar and prototype accounts of prototype distortion data. Consider a space with a dimension for each stimulus component. In the case of the dot patterns there would thus be two dimensions for each dot—one to identify the horizontal position of the dot and one to identify the vertical position. In this space, each category exemplar is identified by a single point, and the category prototype is represented by a point in the center of the cloud of points denoting all exemplars of the category. Smith & Minda's (2001) analogy was to the solar system, with the sun representing the category prototype and the planets the category members that were created by distorting the prototype.

Now consider the probability of responding "A" in an (A, not A) task and how this probability changes with the position of the stimulus in the dot pattern space. According to prototype theory, the probability of responding "A" is completely determined by the similarity, or equivalently the distance, between the stimulus and the category prototype (i.e., the sun in the solar system analogy). According to exemplar theory, however, the probability of responding "A" depends on the (sum of the) similarities between the stimulus and all of the category A exemplars, or equivalently on the distances between the stimulus point and all the planets in the category A solar system.

When the stimulus is outside of the category A cluster or solar system, then prototype and exemplar theories both predict that the probability of responding "A" will increase sharply as the stimulus moves toward the category A prototype,

because the distance between the stimulus and the sun is decreasing as are the distances between the stimulus and each planet. The critical difference between prototype and exemplar theories emerges when the stimulus first enters the category A solar system. For example, as a meteor passes Pluto and enters our solar system, it is still moving closer to the sun and to the planets nearest to the sun, but it is now moving away from Pluto. The closer it moves toward the sun the greater this effect—that is, it is always moving steadily nearer the sun, but as it approaches the sun it begins moving away from more planets and moving nearer to fewer planets. As a result, prototype theory predicts that the probability of responding “A” continually increases as the similarity between the stimulus and the category A prototype is increased. However, exemplar theory predicts that this probability gradient will begin to flatten as the stimulus moves inside the category A cluster. Smith & Minda (2001, 2002; Smith 2002) examined these probability-of-responding-“A” profiles for a number of new and previously published studies and showed that they were steeper than predicted by exemplar theory, but were in general agreement with the predictions of prototype theory.

PERCEPTUAL LEARNING THEORIES The neuroimaging results showing learning-related changes in visual cortex in (A, not A) prototype distortion tasks motivated several proposals that the perceptual representation memory system contributes to learning under these conditions (Ashby & Casale 2002, Reber & Squire 1999). The idea is that performance may be mediated, at least in part, by perceptual learning within visual cortex.

If such visual cortical perceptual learning is important in prototype distortion tasks, then it likely will have different effects in (A, not A) and (A, B) tasks. Consider first an (A, not A) task. The category A prototype will induce a graded pattern of activation throughout visual cortex. One particular cell (or small group of cells) will fire most rapidly to the presentation of this pattern. Call this cell A. A low-level distortion of the category A prototype will be visually similar to the prototype and therefore will likely also cause cell A to fire. Thus, cell A will repeatedly fire throughout training on the category A exemplars. Perceptual learning is thought to occur any time repeated presentations of the same stimulus occur during some relatively brief time interval (Doshier & Lu 1999). As a result, perceptual learning will cause the magnitude of the cell A response to increase throughout training. In contrast, the stimuli associated with the “not A” response will be visually dissimilar to the category A prototype and therefore will be unlikely to cause cell A to fire. During the transfer or testing phase of the experiment, the subject can use the increased sensitivity of cell A to respond accurately. In particular, stimuli from category A are likely to lead to an enhanced visual response compared to stimuli that do not belong to category A. Thus, to respond with above chance accuracy, subjects need only respond “A” to any stimulus that elicits an enhanced visual response. Note that one could interpret cell A as encoding the representation of the category prototype, and thus, this type of perceptual learning could be interpreted as a neuropsychological basis of prototype theory.

Next, consider an (A, B) task. In this case, there will be some cell A maximally tuned to the category A prototype, but there will be some other cell B that is tuned to the category B prototype. During training, every presented stimulus is a distortion of either the category A or category B prototype, so it is likely that either cell A or B will fire on many trials. During the testing phase, all stimuli are again from either category A or B, and so stimuli from both categories will be equally likely to elicit an enhanced visual response. As a result, the mere existence of an enhanced visual response will not help subjects decide whether to respond "A" or "B." The conclusion, therefore, is that perceptual learning could greatly assist in (A, not A) tasks but, by itself, it would be of little help in (A, B) tasks. This is not to say that learning in (A, B) prototype distortion tasks is impossible; only that other learning systems must be used. For example, in low-distortion (A, B) tasks, the enhanced prototype responses caused by perceptual learning might facilitate an explicit memorization strategy in which subjects memorize the A and B prototype patterns and their associated responses.

Much of the cognitive neuroscience data reviewed above supports these predictions. First, neuroimaging results of (A, not A) prototype distortion tasks consistently report learning-related activation in visual cortex (Aizenstein et al. 2000; Reber et al. 1998a,b). Second, neuroimaging results of (A, B) tasks have sometimes failed to find such occipital activation (Vogels et al. 2002), and they have consistently reported task-related activation in prefrontal cortex that is not seen in (A, not A) tasks (Seger et al. 2000, Vogels et al. 2002). Third, a variety of neuropsychological studies show normal (A, not A) prototype distortion learning in patient groups that are impaired in rule-based or information-integration category learning (e.g., schizophrenics, Keri et al. 2001; Parkinson's disease patients, Reber & Squire 1999; Alzheimer's disease patients, Sinha 1999), and in amnesiacs (Knowlton & Squire 1993). Fourth, amnesiacs and patients with Alzheimer's disease are impaired in (A, B) prototype distortion learning (Sinha 1999, Zaki et al. 2003). The prefrontal activation in (A, B) tasks, and the impaired learning of amnesiacs, suggest that learning in (A, B) prototype distortion tasks may be mediated by explicit reasoning strategies and/or by explicit memorization.

WEATHER PREDICTION TASK

An important distinction in category-learning experiments is whether category membership is deterministic or probabilistic. In deterministic tasks, each stimulus is unambiguously a member of one category (i.e., optimal performance is perfect), whereas in probabilistic tasks, at least some stimuli are probabilistically associated with the contrasting categories. For example, on one trial in a probabilistic classification task a particular stimulus might belong to category A but on the next trial the same stimulus might belong to category B. Obviously, in such tasks, perfect performance is impossible. Although most category-learning studies have used deterministic tasks, probabilistic classification also has a long history (e.g.,

Ashby & Gott 1988; Ashby & Maddox 1990, 1992; Estes et al. 1989; Gluck & Bower 1988; Kubovy & Healy 1977).

One popular probabilistic classification task, which is used extensively in cognitive neuroscience, is the so-called weather prediction task (Eldridge et al. 2002; Knowlton et al. 1994, 1996a,b; Reber et al. 1996; Reber & Squire 1999). On each trial of this task, subjects are shown one, two, or three of four possible tarot cards and are asked to indicate whether the presented constellation signals rain or sun. Each card is labeled with a unique, and highly discriminable, geometric pattern. Fourteen of the 16 possible card combinations are used (the no cards and four card patterns are excluded) and each combination is probabilistically associated with the two outcomes. In the original version of the task, the highest possible accuracy was 76% (Knowlton et al. 1994). Interestingly though, a single-cue strategy in which the subject gives one response if one card is present and the other response if that same card is absent yields an accuracy of 75% correct. Knowlton et al. (1994) reported that performance of a control group increased from approximately 50% to 65% correct during the first 50 trials and continued to improve to approximately 75% correct after 350 trials.

Task and Individual Differences Analysis

In order to relate results from the weather prediction task to the other results reviewed in this article, it is important to determine its relationship to the other tasks we have discussed. Because the optimal strategy requires information-integration across cues and is nonverbalizable, the weather prediction task is technically an information-integration task. On the other hand, a single-cue strategy results in nearly optimal performance (75% for a single cue versus 76% for the optimal strategy), so nearly optimal accuracy does not rule out simple rule-based strategies. In addition, because the task uses only a few highly distinct exemplars, explicit memorization is also a plausible strategy. Because a variety of different strategies are all about equally effective, we might expect more individual differences in results obtained with the weather prediction task than with the other tasks we have considered. This possibility makes it especially important to determine what strategy each subject is using before interpreting his or her data.

Gluck et al. (2002) provided a "strategy" analysis of data collected in the weather prediction task to address this issue. In one study, they simply asked subjects what strategy they used after the experiment was over. Although many of the protocols lacked detail, those with sufficient detail generally fell into one of three types: (a) single-cue strategies in which responding was based on the presence or absence of one card (as in a rule-based task), (b) multiple-cue learning (as in an information-integration task), or (c) singleton learning in which correct responses to the single-card patterns were memorized, and guessing occurred for the remaining patterns. Based on these self-reports, Gluck et al. (2002) developed a model-based analysis to identify each subject's strategy in two follow-up studies. Based on a full 200-trial session, 90% (Experiment 1) and 80% (Experiment 2) of the subjects used

a singleton strategy (explicit memorization) to learn the categories. When broken down into 50-trial blocks, a shift from singleton toward multiple-cue strategies was observed (although the proportion of subjects using a multiple-cue strategy was still well below 0.5). Interestingly, there was little correspondence between the strategies that subjects self-reported in their protocols and the strategies identified by the modeling approach. Thus, although originally designed as an information-integration task, these results suggest that subjects adopt a variety of different strategies in the weather prediction task, and the evidence indicates that the most popular choice may be explicit memorization.

Neuropsychological Patient and Neuroimaging Studies

The weather prediction task has been used to study category learning in a number of patient groups. In one of the first such studies, Knowlton et al. (1994) found that amnesiacs performed as well as healthy controls during the first 50 trials of learning, but with extended training, amnesiacs showed a learning deficit relative to healthy controls. Declarative memory is impaired in amnesia, so amnesiacs are impaired in explicit memorization strategies. Healthy controls are not, however, so one interpretation of the late training deficit is that controls begin memorizing and the amnesiacs do not (Knowlton et al. 1994; see also Gluck et al. 1996).

Unlike amnesiacs, patients with Parkinson's or Huntington's disease show learning deficits in the weather prediction task during the first 50 trials that continue throughout training (Knowlton et al. 1996a,b). The weather prediction task has also been used to examine category learning in patients suffering from Alzheimer's disease or schizophrenia. Patients in the early stages of Alzheimer's disease are similar to amnesiacs in the sense that both show anterograde amnesia due to neurodegenerative processes in the medial temporal lobes. As one might predict given this similarity, Alzheimer's patients show intact performance during early trials of the weather prediction task, similar to that seen in amnesia (Eldridge et al. 2002). Schizophrenics exhibit marked abnormalities in executive function and explicit memory. Even so, their performance on the weather prediction task is within normal ranges (Keri et al. 2000).

Neuroimaging studies of the weather prediction task indicate that the medial temporal lobes are active early in learning, and gradually become deactivated as learning progresses (Poldrack et al. 2001). This deactivation is mirrored by a simultaneous activation of the basal ganglia. Specifically, early in learning the basal ganglia are inactive, and gradually become more active as learning progresses.

The weather prediction task has provided a useful tool for studying classification learning and has offered some important insights into the neurobiology of category learning. Even so, the fact that nearly optimal performance can be achieved by a variety of different strategies (e.g., information-integration, rule-based, explicit memorization) makes it difficult to draw strong inferences from data collected with this task. Although the strategies approach developed by Gluck et al. (2002) helps identify the type of strategy that subjects are using (however, see Shohamy

et al. 2004), a better alternative might be to use one of the other tasks discussed in this article that are not so susceptible to identifiability problems.

CONCLUSIONS

The results reviewed in this article offer important lessons. First, when interpreting a category-learning result, it is critical to consider carefully the specific task that was used. For example, Parkinson's disease patients are normal in (A, not A) prototype distortion tasks, they are mildly impaired in information-integration category learning, and they are profoundly impaired in rule-based categorization. Several studies support each of these conclusions, but without specifying the task, these studies would appear to just catalog a confusing set of contradictory results.

Second, although the issue is far from resolved, the results presented here make a strong case that human category learning is mediated by multiple qualitatively distinct systems. To a large extent, it is also becoming clear that this issue—of whether there are one or more category-learning systems—is tied to the historically older issue of whether there are one or more memory systems. Learning is, by definition, the process of laying down some sort of memory trace, and there is certainly no reason to suspect that any of the separate memory systems that have been hypothesized are incapable of storing memories about categories. Although research efforts to resolve this debate will continue, other work is already attacking the next set of questions. For example, the next decade will likely see a flurry of research activity directed at determining the conditions under which the various systems contribute to category learning, at determining how the different systems interact, and at fleshing out their underlying neurobiology.

ACKNOWLEDGMENTS

Preparation of this article was supported by Public Health Service Grants MH3760 and MH59196. We thank Michael Casale and David Smith for their helpful comments.

The *Annual Review of Psychology* is online at <http://psych.annualreviews.org>

LITERATURE CITED

- Aizenstein HJ, MacDonald AW, Stenger VA, Nebes RD, Larson JK, et al. 2000. Complementary category learning systems identified using event-related functional MRI. *J. Cogn. Neurosci.* 12:977–87
- Alfonso-Reese L. 1996. *Dynamics of category learning*. Unpubl. PhD dissert. Univ. Calif., Santa Barbara
- Anderson JR. 1991. The adaptive nature of human categorization. *Psychol. Rev.* 98:409–29
- Ashby FG, Alfonso-Reese LA. 1995. Categorization as probability density estimation. *J. Math. Psychol.* 39:216–33
- Ashby FG, Alfonso-Reese LA, Turken AU, Waldron EM. 1998. A neuropsychological

- theory of multiple systems in category learning. *Psychol. Rev.* 105:442–81
- Ashby FG, Casale MB. 2002. The cognitive neuroscience of implicit category learning. In *Attention and Implicit Learning*, ed. L. Jiménez, pp. 109–41. Amsterdam/Erdenheim, PA: Benjamins
- Ashby FG, Ell SW. 2001. The neurobiology of category learning. *Trends Cogn. Sci.* 5:204–10
- Ashby FG, Ell SW, Waldron EM. 2003. Procedural learning in perceptual categorization. *Mem. Cogn.* 31:1114–25
- Ashby FG, Gott RE. 1988. Decision rules in the perception and categorization of multidimensional stimuli. *J. Exp. Psychol.: Learn. Mem. Cogn.* 14:33–53
- Ashby FG, Isen AM, Turken AU. 1999. A neuropsychological theory of positive affect and its influence on cognition. *Psychol. Rev.* 106:529–50
- Ashby FG, Maddox WT. 1990. Integrating information from separable psychological dimensions. *J. Exp. Psychol.: Hum. Percept. Perform.* 16:598–612
- Ashby FG, Maddox WT. 1992. Complex decision rules in categorization: contrasting novice and experienced performance. *J. Exp. Psychol.: Hum. Percept. Perform.* 18:50–71
- Ashby FG, Maddox WT. 1994. A response time theory of separability and integrality in speeded classification. *J. Math. Psychol.* 38:423–66
- Ashby FG, Maddox WT. 1998. Stimulus categorization. In *Handbook of Perception and Cognition: Measurement, Judgment, and Decision Making*, ed. MH Birnbaum, pp. 251–301. New York: Academic
- Ashby FG, Maddox WT, Bohil CJ. 2002. Observational versus feedback training in rule-based and information-integration category learning. *Mem. Cogn.* 30:666–77
- Ashby FG, Noble S, Filoteo J, Waldron EM, Ell SW. 2003. Category learning deficits in Parkinson's disease. *Neuropsychology* 17:115–24
- Ashby FG, Queller S, Berretty PM. 1999. On the dominance of unidimensional rules in unsupervised categorization. *Percept. Psychophys.* 61:1178–99
- Ashby FG, Townsend JT. 1986. Varieties of perceptual independence. *Psychol. Rev.* 93:154–79
- Ashby FG, Waldron EM. 1999. On the nature of implicit categorization. *Psychon. Bull. Rev.* 6:363–78
- Ashby FG, Waldron EM, Lee WW, Berkman A. 2001. Suboptimality in human categorization and identification. *J. Exp. Psychol.: Gen.* 130:77–96
- Barsalou LW. 2003. Situated simulation in the human conceptual system. *Lang. Cogn. Process. Spec. Iss.: Concept. Rep.* 18:513–62
- Barsalou LW, Simmons WK, Barbey AK, Wilson CD. 2003. Grounding conceptual knowledge in modality-specific systems. *Trends Cogn. Sci.* 7:84–91
- Brooks L. 1978. Nonanalytic concept formation and memory for instances. In *Cognition and Categorization*, ed. E Rosch, BB Lloyd, pp. 169–211. Hillsdale, NJ: Erlbaum
- Brown RG, Marsden CD. 1988. Internal versus external cues and the control of attention in Parkinson's disease. *Brain* 111:323–45
- Burner JS, Goodnow J, Austin G. 1956. *A Study of Thinking*. New York: Wiley
- Cools AR, van den Bercken JHL, Horstink MWI, van Spaendonck KPM, Berger HJC. 1984. Cognitive and motor shifting aptitude disorder in Parkinson's disease. *J. Neurol. Neurosurg. Psychiatry* 47:443–53
- Cree GS, McRae K. 2003. Analyzing the factors underlying the structure and computation of the meaning of chipmunk, cherry, chisel, cheese, and cello (and many other such concrete nouns). *J. Exp. Psychol.: Gen.* 132:163–201
- Crick F, Koch C. 1990. Towards a neurobiological theory of consciousness. *Semin. Neurosci.* 2:263–75
- Crick F, Koch C. 1995. Are we aware of neural activity in primary visual cortex? *Nature* 375:121–23
- Crick F, Koch C. 1998. Consciousness and neuroscience. *Cereb. Cortex* 8:97–107

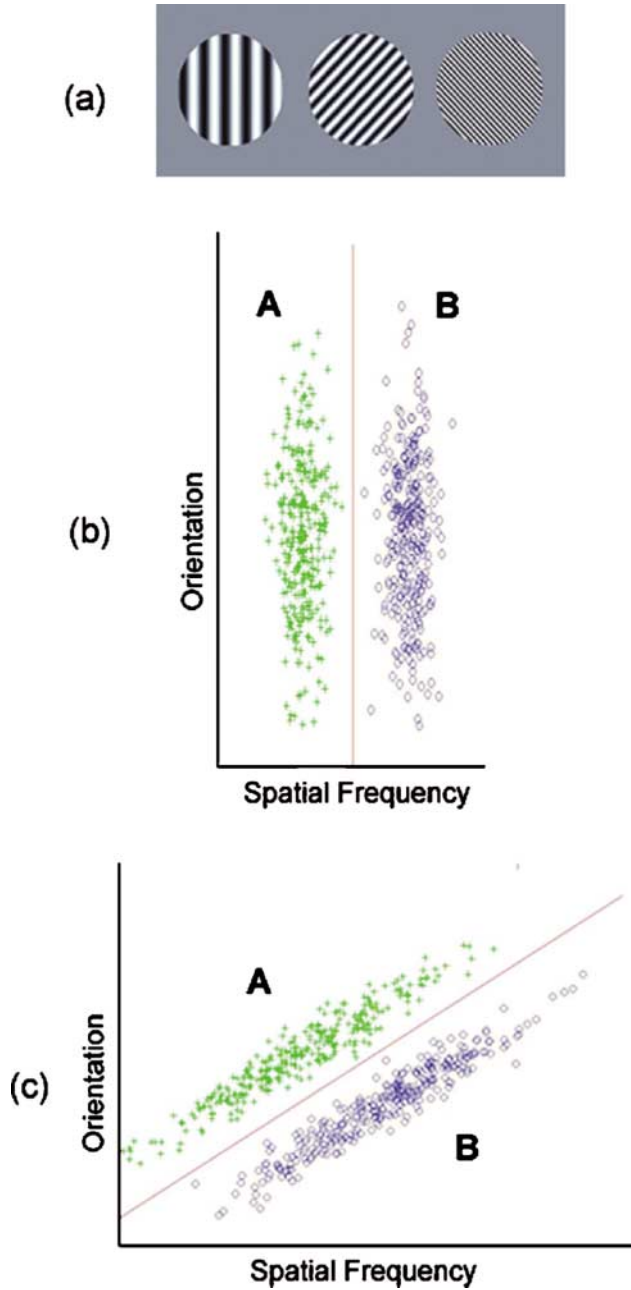
- Dosher B, Lu ZL. 1999. Mechanisms of perceptual learning. *Vision Res.* 39:3197–221
- Downes JJ, Roberts AC, Sahakian BJ, Evenden JL, Morris RG, Robbins TW. 1989. Impaired extra-dimensional shift performance in medicated and unmedicated Parkinson's disease: evidence for a specific attentional dysfunction. *Neuropsychologia* 27:1329–43
- Eldridge LL, Masterman D, Knowlton BJ. 2002. Intact implicit habit learning in Alzheimer's disease. *Behav. Neurosci.* 116:722–26
- Elliott R, Rees G, Dolan RJ. 1999. Ventromedial prefrontal cortex mediates guessing. *Neuropsychologia* 37:403–11
- Erickson MA, Kruschke JK. 1998. Rules and exemplars in category learning. *J. Exp. Psychol.: Gen.* 127:107–40
- Estes WK. 1986. Array models for category learning. *Cogn. Psychol.* 18:500–49
- Estes WK. 1994. *Classification and Cognition*. New York: Oxford Univ. Press
- Estes WK, Campbell JA, Hatsopoulos N, Hurwitz JB. 1989. Base-rate effects in category learning: a comparison of parallel network and memory storage-retrieval models. *J. Exp. Psychol.: Learn. Mem. Cogn.* 15:556–71
- Filoteo JV, Maddox WT, Davis JD. 2001a. A possible role of the striatum in linear and non-linear categorization rule learning: evidence from patients with Huntington's disease. *Behav. Neurosci.* 115:786–98
- Filoteo JV, Maddox WT, Davis JD. 2001b. Quantitative modeling of category learning in amnesic patients. *J. Int. Neuropsychol. Soc.* 7:1–19
- Filoteo JV, Maddox WT, Salmon DP, Song DD. 2004. Information-integration category learning in patients with striatal dysfunction. *Neuropsychology*. In press
- Garner WR. 1974. *The Processing of Information and Structure*. New York: Wiley
- Gluck MA, Bower GH. 1988. From conditioning to category learning: an adaptive network model. *J. Exp. Psychol.: Gen.* 117:227–47
- Gluck MA, Oliver LM, Myers CE. 1996. Late-training amnesic deficits in probabilistic category learning: a neurocomputational analysis. *Learn. Mem.* 3:326–40
- Gluck MA, Shohamy D, Myers C. 2002. How do people solve the “weather prediction” task? Individual variability in strategies for probabilistic category learning. *Learn. Mem.* 9:408–18
- Goldman-Rakic PS. 1987. Circuitry of the prefrontal cortex and the regulation of behavior by representational knowledge. In *Handbook of Physiology*, ed. F Plum, V Mountcastle, pp. 373–417. Bethesda, MD: Am. Physiol. Soc.
- Goldman-Rakic PS. 1995. Cellular basis of working memory. *Neuron* 14:477–85
- Heaton RK. 1981. *A Manual for the Wisconsin Card Sorting Test*. Odessa, FL: Psychol. Assess. Resourc.
- Herbranson WT, Fremouw T, Shimp CP. 1999. The randomization procedure in the study of categorization of multidimensional stimuli by pigeons. *J. Exp. Psychol. Anim. Behav. Process.* 25:113–34
- Hintzman DL. 1986. Schema abstraction in a multiple-trace memory model. *Psychol. Rev.* 93:411–28
- Hintzman DL, Ludlam G. 1980. Differential forgetting of prototypes and old instances: simulation by an exemplar-based classification model. *Mem. Cogn.* 8:378–82
- Homa D, Cultice JC. 1984. Role of feedback, category size, and stimulus distortion on the acquisition and utilization of ill-defined categories. *J. Exp. Psychol.: Learn. Mem. Cogn.* 10:83–94
- Homa D, Rhoads D, Chambliss D. 1979. Evolution of conceptual structure. *J. Exp. Psychol.: Learn. Mem. Cogn.* 5:11–23
- Homa D, Sterling S, Trepel L. 1981. Limitations of exemplar-based generalization and the abstraction of categorical information. *J. Exp. Psychol.: Learn. Mem. Cogn.* 7:418–39
- Humphreys GW, Forde EME. 2001. Hierarchies, similarity and interactivity in object recognition: “category-specific” neuropsychological deficits. *Behav. Brain Sci.* 24:453–509
- Janowsky JS, Shimamura AP, Kritchewsky M,

- Squire LR. 1989. Cognitive impairment following frontal lobe damage and its relevance to human amnesia. *Behav. Neurosci.* 103:548–60
- Joseph JE. 2001. Functional neuroimaging studies of category specificity in object recognition: a critical review and meta-analysis. *Cogn. Affect. Behav. Neurosci.* 1: 119–36
- J. Exp. Anal. Behav.* 2002. Vol. 78, No. 3 (entire issue)
- Kéri S, Kalman J, Rapcsak SZ, Antal A, Benedek G, Janka Z. 1999. Classification learning in Alzheimer's disease. *Brain* 122: 1063–68
- Kéri S, Kelemen O, Benedek G, Janka Z. 2001. Intact prototype learning in schizophrenia. *Schizophr. Res.* 52:261–64
- Keri S, Kelemen O, Szekeres G, Bagoczky N, Erdelyi R, et al. 2000. Schizophrenics know more than they can tell: probabilistic classification learning in schizophrenia. *Psychol. Med.* 30:149–55
- Kimberg DY, D'Esposito M, Farah MJ. 1997. Frontal lobes: neuropsychological aspects. In *Behavioral Neurology and Neuropsychology*, ed. TE Feinberg, MJ Farah, pp. 409–18. New York: McGraw-Hill
- Knowlton BJ, Mangels JA, Squire LR. 1996a. A neostriatal habit learning system in humans. *Science* 273:1399–402
- Knowlton BJ, Squire LR. 1993. The learning of natural categories: parallel memory systems for item memory and category-level knowledge. *Science* 262:1747–49
- Knowlton BJ, Squire LR, Gluck MA. 1994. Probabilistic classification learning in amnesia. *Learn. Mem.* 1:106–20
- Knowlton BJ, Squire LR, Paulsen JS, Swerdlow NR, Swenson M, Butters N. 1996. Dissociations within nondeclarative memory in Huntington's disease. *Neuropsychology* 10:538–48
- Kolodny JA. 1994. Memory processes in classification learning: an investigation of amnesic performance in categorization of dot patterns and artistic styles. *Psychol. Sci.* 5:164–69
- Konishi S, Karwazu M, Uchida I, Kikyo H, Asakura I, Miyashita Y. 1999. Contribution of working memory to transient activation in human inferior prefrontal cortex during performance of the Wisconsin Card Sorting Test. *Cereb. Cortex* 9:745–73
- Kruschke JK. 1992. ALCOVE: an exemplar-based connectionist model of category learning. *Psychol. Rev.* 99:22–44
- Kubovy M, Healy AF. 1977. The decision rule in probabilistic categorization: what it is and how it is learned. *J. Exp. Psychol.: Gen.* 106:427–46
- Lamberts K. 2000. Information-accumulation theory of speeded categorization. *Psychol. Rev.* 107:227–60
- Leng NR, Parkin AJ. 1988. Double dissociation of frontal dysfunction in organic amnesia. *Br. J. Clin. Psychol.* 27:359–62
- Lewandowsky S, Kalish M, Ngang SK. 2002. Simplified learning in complex situations: knowledge partitioning in function learning. *J. Exp. Psychol.: Gen.* 131:163–93
- Lockhead GR. 1966. Effects of dimensional redundancy on visual discrimination. *J. Exp. Psychol.* 72:95–104
- Lombardi WJ, Andreason PJ, Sirocco KY, Rio DE, Gross RE, et al. 1999. Wisconsin Card Sorting Test performance following head injury: Dorsolateral fronto-striatal circuit activity predicts perseveration. *J. Clin. Exp. Neuropsychol.* 21:2–16
- Love BC, Medin DL, Gureckis TM. 2004. SUSTAIN: a network model of category learning. *Psychol. Rev.* 111:309–32
- Maddox WT, Ashby FG. 1993. Comparing decision bound and exemplar models of categorization. *Percept. Psychophys.* 53:49–70
- Maddox WT, Ashby FG. 2004. Dissociating explicit and procedural-learning based systems of perceptual category learning. *Behav. Process.* 66:309–32
- Maddox WT, Ashby FG, Bohil CJ. 2003. Delayed feedback effects on rule-based and information-integration category learning. *J. Exp. Psychol.: Learn. Mem. Cogn.* 29:650–62

- Maddox WT, Ashby FG, Ing AD, Pickering AD. 2004. Disrupting feedback processing interferes with rule-based but not information-integration category learning. *Mem. Cogn.* 32:582–91
- Maddox WT, Bohil CJ, Ing AD. 2004. Evidence for a procedural learning-based system in category learning. *Psychon. Bull. Rev.* In press
- Maddox WT, Filoteo JV. 2001. 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.* 7: 710–27
- Markman AB, Ross BH. 2003. Category use and category learning. *Psychol. Bull.* 129:592–613
- McKinley SC, Nosofsky RM. 1995. Investigations of exemplar and decision bound models in large, ill-defined category structures. *J. Exp. Psychol.: Hum. Percept. Perform.* 21:128–48
- Medin DL, Schaffer MM. 1978. Context theory of classification learning. *Psychol. Rev.* 85:207–38
- Medin DL, Schwanenflugel PJ. 1981. Linear separability in classification learning. *J. Exp. Psychol.: Learn. Mem. Cogn.* 7:355–68
- Minda JP, Smith JD. 2001. Prototypes in category learning: the effects of category size, category structure, and stimulus complexity. *J. Exp. Psychol.: Learn. Mem. Cogn.* 27:775–99
- Murphy GL. 2002. *The Big Book of Concepts*. Cambridge, MA: MIT Press
- Nissen MJ, Bullemer P. 1987. Attentional requirements of learning: evidence from performance measures. *Cogn. Psychol.* 19:1–32
- Nosofsky RM. 1986. Attention, similarity, and the identification categorization relationship. *J. Exp. Psychol.: Gen.* 115:39–57
- Nosofsky RM. 1991. Typicality in logically defined categories: exemplar-similarity versus rule instantiation. *Mem. Cogn.* 19:131–50
- Nosofsky RM, Clark SE, Shin HJ. 1989. Rules and exemplars in categorization, identification, and recognition. *J. Exp. Psychol.: Learn. Mem. Cogn.* 15:282–304
- Nosofsky RM, Gluck MA, Palmeri TJ, McKinley SC, Glauthier P. 1994a. Comparing models of rule-based classification learning: a replication and extension of Shepard, Hovland, and Jenkins 1961. *Mem. Cogn.* 22:352–69
- Nosofsky RM, Palmeri TJ. 1996. Learning to classify integral-dimension stimuli. *Psychon. Bull. Rev.* 3:222–26
- Nosofsky RM, Palmeri TJ, McKinley SC. 1994b. Rule-plus-exception model of classification learning. *Psychol. Rev.* 101:53–79
- Nosofsky RM, Zaki SR. 1998. Dissociations between categorization and recognition in amnesic and normal individuals: an exemplar-based interpretation. *Psychol. Sci.* 9:247–55
- Palmeri TJ, Flanery MA. 1999. Learning about categories in the absence of training: profound amnesia and the relationship between perceptual categorization and recognition memory. *Psychol. Sci.* 10:526–30
- Pickering AD. 1997. New approaches to the study of amnesic patients: What can a neurofunctional philosophy and neural network methods offer? *Memory* 5:255–300
- Poldrack RA, Clark J, Pare-Blagoev EJ, Shohamy D, Moyano JC, et al. 2001. Interactive memory systems in the human brain. *Nature* 414:546–50
- Posner MI, Keele SW. 1968. On the genesis of abstract ideas. *J. Exp. Psychol.* 77:353–63
- Posner MI, Keele SW. 1970. Retention of abstract ideas. *J. Exp. Psychol.* 83:304–8
- Posner MI, Petersen SE. 1990. Attention systems in the human brain. *Annu. Rev. Neurosci.* 13:25–42
- Rao SM, Bobholz JA, Hammeke TA, Tosen AC, Woodley SJ, et al. 1997. Functional MRI evidence for subcortical participation in conceptual reasoning skills. *Neuroreport* 8:1987–93
- Reber PJ, Knowlton BJ, Squire LR. 1996. Dissociable properties of memory systems: differences in the flexibility of declarative and nondeclarative knowledge. *Behavioral Neuroscience* 110:861–71

- Reber PJ, Squire LR. 1999. Intact learning of artificial grammars and intact category learning by patients with Parkinson's disease. *Behav. Neurosci.* 113:235–42
- Reber PJ, Stark CEL, Squire LR. 1998a. Contrasting cortical activity associated with category memory and recognition memory. *Learn. Mem.* 5:420–28
- Reber PJ, Stark CEL, Squire LR. 1998b. Cortical areas supporting category learning identified using functional MRI. *Proc. Natl. Acad. Sci. USA* 95:747–50
- Reed SK. 1972. Pattern recognition and categorization. *Cogn. Psychol.* 3:382–407
- Robinson AL, Heaton RK, Lehman RA, Stilson DW. 1980. The utility of the Wisconsin Card Sorting Test in detecting and localizing frontal lobe lesions. *J. Consult. Clin. Psychol.* 48:605–64
- Rogers RD, Andrews TC, Grasby PM, Brooks DJ, Robbins TW. 2000. Contrasting cortical and subcortical activations produced by attentional-set shifting and reversal learning in humans. *J. Cogn. Neurosci.* 12:142–62
- Rosch E. 1973. Natural categories. *Cogn. Psychol.* 4:328–50
- Rosch E. 1975. Cognitive reference points. *Cogn. Psychol.* 7:532–47
- Salatas H, Bourne LE. 1974. Learning conceptual rules: III. Processes contributing to rule difficulty. *Mem. Cogn.* 2:549–53
- Schultz W. 1992. Activity of dopamine neurons in the behaving primate. *Semin. Neurosci.* 4:129–38
- Seger CA, Cincotta CM. 2002. Striatal activity in concept learning. *Cogn. Affect. Behav. Neurosci.* 2:149–61
- Seger CA, Poldrack RA, Prabhakaran V, Zhao M, Glover GH, Gabrieli JDE. 2000. Hemispheric asymmetries and individual differences in visual concept learning as measured by functional MRI. *Neuropsychologia* 38:1316–24
- Shepard RN. 1964. Attention and the metric structure of the stimulus space. *J. Math. Psychol.* 1:54–87
- Shepard RN, Chang JJ. 1963. Stimulus generalization in learning of classifications. *J. Exp. Psychol.* 65:94–102
- Shepard RN, Hovland CI, Jenkins HM. 1961. Learning and memorization of classifications. *Psychol. Monogr.* 75(13, Whole No. 517)
- Shin HJ, Nosofsky RM. 1992. Similarity-scaling studies of dot-pattern classification and recognition. *J. Exp. Psychol.: Gen.* 121:278–304
- Shohamy D, Myers CE, Grossman S, Sage J, Gluck MA, Poldrack RA. 2004. Corticostriatal contributions to feedback-based learning: converging data from neuroimaging and neuropsychology. *Brain* 127:1–9
- Sinha RR. 1999. Neuropsychological substrates of category learning. *Dissert. Abstr. Int.: Sect. B: Sci. Engineer.* 60(5-B), 2381 (UMI No. AEH9932480)
- Smith EE, Medin DL. 1981. *Categories and Concepts*. Cambridge, MA: Harvard Univ. Press
- Smith JD. 2002. Exemplar theory's predicted typicality gradient can be tested and disconfirmed. *Psychol. Sci.* 13:437–42
- Smith JD, Minda JP. 1998. Prototypes in the mist: the early epochs of category learning. *J. Exp. Psychol.: Learn. Mem. Cogn.* 24:1411–30
- Smith JD, Minda JP. 2001. Journey to the center of the category: the dissociation in amnesia between categorization and recognition. *J. Exp. Psychol.: Learn. Mem. Cogn.* 27:984–1002
- Smith JD, Minda JP. 2002. Distinguishing prototype-based and exemplar-based processes in category learning. *J. Exp. Psychol.: Learn. Mem. Cogn.* 28:800–11
- Smith JD, Minda JP, Washburn DA. 2004. Category learning in Rhesus monkeys: a study of the Shepard, Hovland, and Jenkins tasks. *J. Exp. Psychol. Gen.* 133:398–414
- Squire LR, Knowlton BJ. 1995. Learning about categories in the absence of memory. *Proc. Natl. Acad. Sci. USA* 92:12470–74
- Sternberg S. 1966. High-speed scanning in human memory. *Science* 153:652–54

- Strange W, Keeney T, Kessel FS, Jenkins JJ. 1970. Abstraction over time of prototypes from distortions of random dot patterns: a replication. *J. Exp. Psychol.* 83:508–10
- van Domburg PHMF, ten Donkelaar HJ. 1991. *The Human Substantia Nigra and Ventral Tegmental Area*. Berlin: Springer-Verlag
- Vauclair J. 2002. Categorization and conceptual behavior in nonhuman primates. In *The Cognitive Animal: Empirical and Theoretical Perspectives on Animal Cognition*, ed. M Bekoff, C Allen, GM Burghardt, pp. 239–45. Cambridge, MA: MIT Press
- Vogels R, Sary G, Dupont P, Orban GA. 2002. Human brain regions involved in visual categorization. *Neuroimage* 16:401–14
- Volz H-P, Gaser C, Haeger F, Rzanny R, Mentzel H-J, et al. 1997. Brain activation during cognitive stimulation with the Wisconsin Card Sorting Test—a functional MRI study on healthy volunteers and schizophrenics. *Psychiatry Res.: Neuroimaging* 75:45–157
- Waldron EM, Ashby FG. 2001. The effects of concurrent task interference on category learning: evidence for multiple category learning systems. *Psychon. Bull. Rev.* 8:168–76
- Wickens J. 1993. *A Theory of the Striatum*. New York: Pergamon
- Willingham DB. 1998. A neuropsychological theory of motor skill learning. *Psychol. Rev.* 105:558–84
- Willingham DB, Wells LA, Farrell JM, Stemwedel ME. 2000. Implicit motor sequence learning is represented in response locations. *Mem. Cogn.* 28:366–75
- Wilson CJ. 1995. The contribution of cortical neurons to the firing pattern of striatal spiny neurons. In *Models of Information Processing in the Basal Ganglia*, ed. JC Houk, JL Davis, DG Beiser, pp. 29–50. Cambridge, MA: Bradford
- Winocur G, 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
- Zaki SR, Nosofsky RM, Jessup NM, Unversagt FW. 2003. Categorization and recognition performance of a memory-impaired group: evidence for single-system models. *J. Int. Neuropsychol. Soc.* 9:394–406



See legend on next page

Figure 1 (a) Three stimuli that might be used in a perceptual categorization experiment. Each stimulus is a circular sine-wave grating that varies across trials in spatial frequency and orientation. (b) A plot of stimuli that might be used in a rule-based category-learning task. Each plus denotes the spatial frequency and orientation of an exemplar from category A and each circle denotes a category B exemplar. The vertical line is the optimal category boundary. (c) A plot of stimuli that might be used in an information-integration category-learning task.

CONTENTS

Frontispiece— <i>Richard F. Thompson</i>	xviii
PREFATORY	
In Search of Memory Traces, <i>Richard F. Thompson</i>	1
DECISION MAKING	
Indeterminacy in Brain and Behavior, <i>Paul W. Glimcher</i>	25
BRAIN IMAGING/COGNITIVE NEUROSCIENCE	
Models of Brain Function in Neuroimaging, <i>Karl J. Friston</i>	57
MUSIC PERCEPTION	
Brain Organization for Music Processing, <i>Isabelle Peretz</i> and <i>Robert J. Zatorre</i>	89
SOMESTHETIC AND VESTIBULAR SENSES	
Vestibular, Proprioceptive, and Haptic Contributions to Spatial Orientation, <i>James R. Lackner and Paul DiZio</i>	115
CONCEPTS AND CATEGORIES	
Human Category Learning, <i>F. Gregory Ashby and W. Todd Maddox</i>	149
ANIMAL LEARNING AND BEHAVIOR: CLASSICAL	
Pavlovian Conditioning: A Functional Perspective, <i>Michael Domjan</i>	179
NEUROSCIENCE OF LEARNING	
The Neuroscience of Mammalian Associative Learning, <i>Michael S. Fanselow and Andrew M. Poulos</i>	207
HUMAN DEVELOPMENT: EMOTIONAL, SOCIAL, AND PERSONALITY	
Behavioral Inhibition: Linking Biology and Behavior Within a Developmental Framework, <i>Nathan A. Fox, Heather A. Henderson,</i> <i>Peter J. Marshall, Kate E. Nichols, and Melissa A. Ghera</i>	235
BIOLOGICAL AND GENETIC PROCESSES IN DEVELOPMENT	
Human Development: Biological and Genetic Processes, <i>Irving I. Gottesman and Daniel R. Hanson</i>	263

SPECIAL TOPICS IN PSYCHOPATHOLOGY

- The Psychology and Neurobiology of Suicidal Behavior,
Thomas E. Joiner Jr., Jessica S. Brown, and LaRicka R. Wingate 287

DISORDERS OF CHILDHOOD

- Autism in Infancy and Early Childhood, *Fred Volkmar,
Kasia Chawarska, and Ami Klin* 315

CHILD/FAMILY THERAPY

- Youth Psychotherapy Outcome Research: A Review and Critique
of the Evidence Base, *John R. Weisz, Amanda Jensen Doss,
and Kristin M. Hawley* 337

ALTRUISM AND AGGRESSION

- Prosocial Behavior: Multilevel Perspectives, *Louis A. Penner,
John F. Dovidio, Jane A. Piliavin, and David A. Schroeder* 365

**INTERGROUP RELATIONS, STIGMA, STEREOTYPING,
PREJUDICE, DISCRIMINATION**

- The Social Psychology of Stigma, *Brenda Major
and Laurie T. O'Brien* 393

PERSONALITY PROCESSES

- Personality Architecture: Within-Person Structures and Processes,
Daniel Cervone 423

PERSONALITY DEVELOPMENT: STABILITY AND CHANGE

- Personality Development: Stability and Change, *Avshalom Caspi,
Brent W. Roberts, and Rebecca L. Shiner* 453

WORK MOTIVATION

- Work Motivation Theory and Research at the Dawn of the Twenty-First
Century, *Gary P. Latham and Craig C. Pinder* 485

GROUPS AND TEAMS

- Teams in Organizations: From Input-Process-Output Models to IMOI
Models, *Daniel R. Ilgen, John R. Hollenbeck, Michael Johnson,
and Dustin Jundt* 517

LEADERSHIP

- Presidential Leadership, *George R. Goethals* 545

PERSONNEL EVALUATION AND COMPENSATION

- Personnel Psychology: Performance Evaluation and Pay for Performance,
Sara L. Rynes, Barry Gerhart, and Laura Parks 571

**PSYCHOPHYSIOLOGICAL DISORDERS AND PSYCHOLOGICAL EFFECTS
ON MEDICAL DISORDERS**

- Psychological Approaches to Understanding and Treating Disease-Related
Pain, *Francis J. Keefe, Amy P. Abernethy, and Lisa C. Campbell* 601

TIMELY TOPIC

- Psychological Evidence at the Dawn of the Law's Scientific Age,
David L. Faigman and John Monahan 631

INDEXES

- Subject Index 661
Cumulative Index of Contributing Authors, Volumes 46–56 695
Cumulative Index of Chapter Titles, Volumes 46–56 700

ERRATA

An online log of corrections to *Annual Review of Psychology* chapters
may be found at <http://psych.annualreviews.org/errata.shtml>