Chapter 26

THE NEUROPSYCHOLOGY OF PERCEPTUAL CATEGORY LEARNING*

W. TODD MADDOX

University of Texas

J. VINCENT FILOTEO

University of California, San Diego &
Veterans Administration San Diego Healthcare System

Contents

* This research was supported in part by National Institute of Health Grant R01 MH59196 to WTM, National Institute of Neurological Disorders and Stroke Grant R01 41372 to JVF, and a James McDonnell Foundation Grant.

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

Abstract

There is widespread agreement that multiple qualitatively different category learning systems mediate the learning of different category structures. Two systems that have received support are (1) a frontal-based explicit system that uses logical reasoning, depends on working memory and executive attention, and is mediated primarily by the anterior cingulate, the prefrontal cortex, and the head of the caudate; and (2) a basal ganglia-mediated implicit system that uses procedural learning and requires a dopamine reward signal. This chapter reviews a large body of work conducted in our laboratories that examines the details of the two proposed systems using neurological patients as experimental participants. Collectively, the studies suggest that the medial temporal lobes are little involved in category learning with large categories. They also suggest that, in striatal-damaged patients, the need to ignore irrelevant information is predictive of a rule-based category learning deficit, whereas the complexity of the rule is predictive of an information-integration category learning deficit.

1. Introduction

Category learning involves laying down a memory trace that can be used to improve the efficiency (i.e., accuracy and speed) of responding. It is now widely accepted that mammals have multiple memory systems [Schacter (1987), Squire (1992)], and this fact alone makes it reasonable to postulate that multiple category learning systems might also exist. This chapter reviews a body of work that suggests that perceptual category learning is characterized by multiple systems, each of which involves a set of diverse neurocognitive processes. This chapter builds upon the work outlined in the previous chapter by Ashby and Valentin, who reviewed a number of studies that tested *a priori* predictions from a recently proposed neurobiologically plausible multiple systems theory called the Competition between Verbal and Implicit Systems theory [COVIS; Ashby et al. (1998), Ashby and Waldron (1999), for a review, see Maddox and Ashby (2004)]. In each study, some experimental manipulation was introduced that was predicted *a priori* to affect processing in one system, but not the other. All of these studies used healthy young adults as participants. These studies provide a nice foundation, but the next step is to examine the systems in greater detail. One way to achieve this goal is to examine category learning in neurological patients with damage to specific brain areas. The aim of this chapter is to review a body of work conducted in our laboratories that examined category learning in various patient populations.

This chapter is organized as follows. First, we briefly introduce COVIS and the proposed underlying neurobiology. A more detailed description is offered in the previous chapter and in Maddox and Ashby (2004). The second section briefly reviews the qualitative dissociations introduced in the previous chapter. The third section reviews a body of work that examines category learning in various patient populations. The third section is subdivided into sections devoted to rule-based and information-integration category learning in patients with amnesia, Parkinson's disease (PD), or Huntington's disease (HD). The final section offers a brief summary and conclusions. It is important to note that this is not a substantive review of the field. Two excellent reviews are provided by Kéri (2003) and Poldrack and Packard (2003). Rather, this chapter reviews and integrates a large body of patient work conducted in our laboratories and others that takes a systematic empirical approach, supplemented by the application of a series of quantitative models, to the study of perceptual category learning.

2. Competition between verbal and implicit systems (COVIS)

A growing body of research suggests that the learning of different types of category structures is mediated by different systems with distinct but partially overlapping neurobiological substrates [Reber and Squire (1994), Pickering (1997), Erickson Kruschke (1998), Smith, Patalano and Jonides (1998), however, see Nosofsky and Johansen (2000), Ashby and Ell (2001, 2002), Maddox and Ashby (2004)]. One of the most successful multiple systems models of category learning, and the only one that specifies the underlying neurobiology, is COVIS. COVIS postulates two systems that compete throughout learning – an explicit hypothesis-testing system that uses logical reasoning and depends on working memory and executive attention, and a procedural-learning-based system that relies more on incremental and feedback-learning processes. One intriguing aspect of the procedural-learning-based system is its association with the processes involved in motor performance [e.g., Willingham (1998), Hazeltine and Ivry (2002)], which leads to the important prediction that categories learned via a procedural-learning-based system should be closely linked to the motor response.

Much of the evidence for multiple category learning systems comes from two different types of categorization tasks. *Rule-based category learning tasks* are those in which the category structures can be learned via some explicit reasoning process that treats each stimulus dimension separately. Frequently, the rule that maximizes accuracy (i.e., the optimal rule) is easy to describe verbally [Ashby et al. (1998)]. For example, in Figure 1a, the stimuli (with one presented on each trial) are composed of a single line that varies in length and orientation across trials. Each symbol in Figure 1 denotes the length and orientation of one stimulus. Also shown in Figure 1 are the decision bounds that maximize categorization accuracy. In the rule-based task, the optimal bound requires observers to attend to length and ignore orientation. The vertical bound in Figure 1a corresponds to the rule: "Respond A if the line is short and B if it is long."

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 and Gott (1988)]. Perceptual integration could take many forms – from treating the stimulus as a Gestalt to computing a weighted linear combination of the dimensional values[1]. In many cases, the optimal rule in information-integration tasks is difficult or impossible to describe verbally [Ashby et al. (1998)]. The information-integration task in Figure 1b was generated by rotating the rule-based categories by 45°. Note that the information-integration rule is linear. Figure 1c depicts a case in which the information-integration rule is nonlinear. Category structures like these were used in the studies reviewed in the previous chapter and in some of the studies reviewed below.

COVIS assumes that learning in rule-based tasks is dominated by an explicit system that uses working memory and executive attention to generate and test hypotheses and is mediated primarily by the anterior cingulate, the prefrontal cortex, and the head of the caudate nucleus (see Figure 2 from the previous chapter). Learning in information-integration tasks is dominated by an implicit procedural-learning-based system, which is mediated largely within the tail of the caudate nucleus [see Figure 3 in Chapter 25; Ashby et al. (1998), Willingham (1998), Ashby and Ell (2001)]. (see the previous chapter for details).

[1] A conjunction rule (e.g., respond A if the stimulus is small on dimension x and small on dimension y) is a rule-based task rather than an information-integration task because separate decisions are first made about each dimension (e.g., small or large) and then the outcome of these decisions is combined (integration is post-decisional, not predecisional).

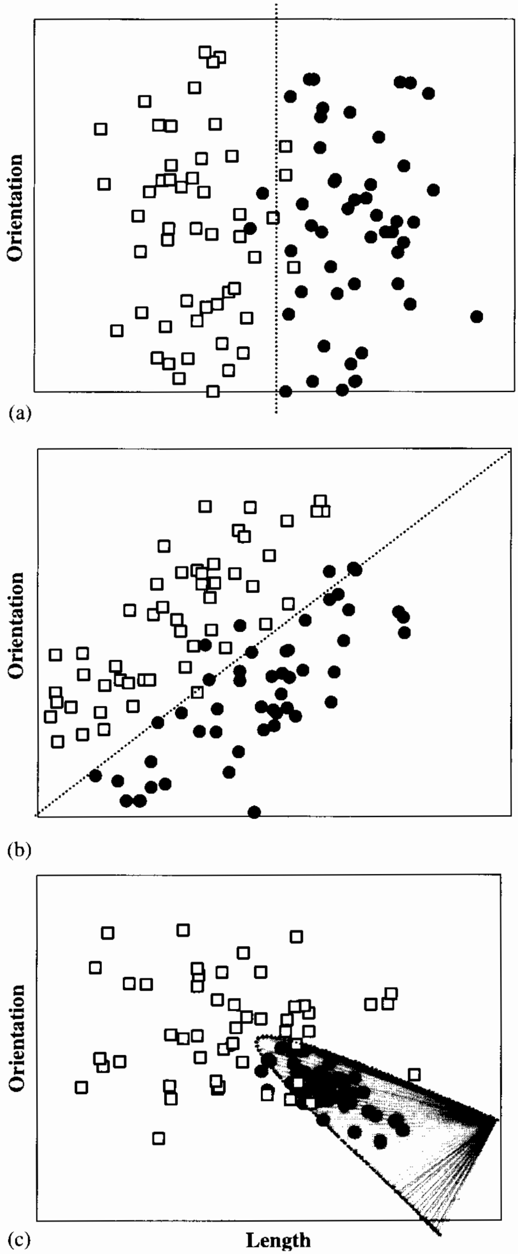

Fig. 1. Stimuli and optimal decision bound from (a) rule-based, (b) linear information-integration, and (c) nonlinear information-integration categorization condition. Open squares denote category A items, and filled circles denote category B items.

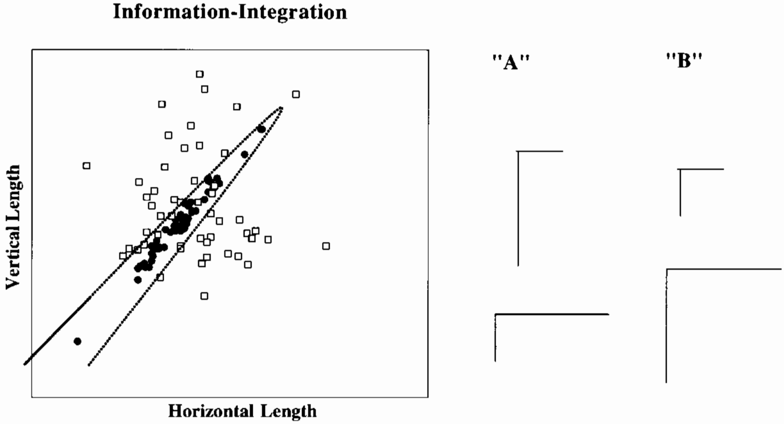

Fig. 2. Nonlinear information integration category structures and representative stimuli used in Filoteo et al. (2001a). Open squares denote category A stimuli and filled circles denote category B stimuli. The broken quadratic curve denotes the optimal decision bound.

3. Testing *a priori* Predictions of COVIS

In this section we briefly review a number of studies that provided empirical tests of *a priori* predictions derived from the proposed neurobiological underpinnings of COVIS. Because the hypothesis-testing system is under conscious control and has full access to working memory[2] and executive attention, the placement and timing of the feedback signal should not be critical for rule-based category learning because this information can be held consciously in working memory. 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. As a result, it would depend more heavily on the placement and timing of the feedback. As a test of these predictions, rule-based and information-integration category learning was compared across an observational training condition (in which observers were informed before stimulus presentation what category the ensuing stimulus would be from) and a traditional feedback training condition (in which the category label followed the response) [Ashby et al. (2002)], and across an immediate feedback condition (in which corrective feedback was provided immediately following the response) and a delayed feedback condition (in which corrective feedback was delayed by 2.5, 5, or 10 s following the response) [Maddox, Ashby and Bohil (2003)]. In line with the COVIS predictions, observational training and delayed feedback negatively impacted information-integration category learning but had little effect on rule-based category learning.

[2] Many argue for the existence of an implicit form of working memory that may not be available to conscious awareness. When we use the term "working memory," we refer to a conscious, verbalizable process.

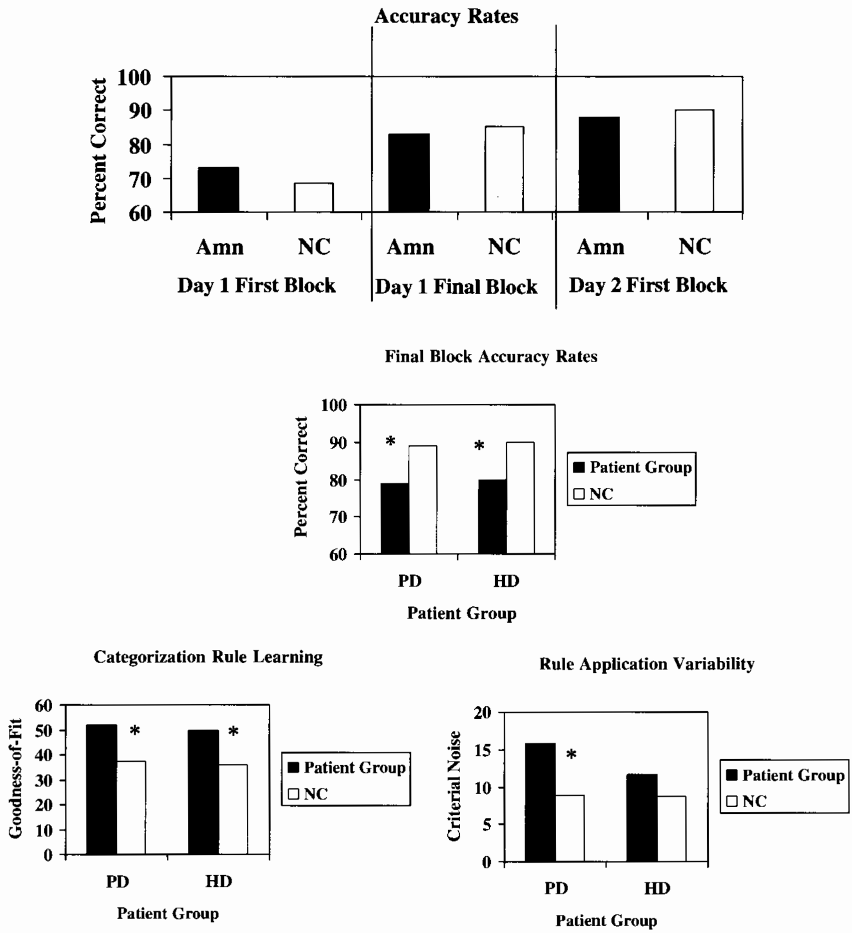

Fig. 3. (Top panel) Nonlinear information-integration percent correct for the amnesic and control participants from Filoteo et al. (2001a) during the first and final blocks of trials from Day 1 and the first block of trials from Day 2. (Bottom panels) Accuracy rates, categorization rule learning, and rule application variability estimates from the nonlinear information-integration studies conducted by Filoteo et al. (2001b) and Maddox and Filoteo (2001). *Statistically significant performance difference ($p<0.05$).

The implicit system in COVIS is assumed to be procedural-learning-based. The quintessential paradigm for studying procedural learning is the serial reaction time (SRT) task. In a typical SRT task, one of *n* stimuli is presented on each trial and each stimulus is associated with its own response key. The observer's task is to press the relevant key as quickly as possible. A large response time improvement is observed when the stimulus sequence is repeated, even though the observer is unaware that a sequence exists.

Willingham et al. (2000) showed that changing the location of the response keys interferes with SRT learning even when the sequence of stimulus positions is unchanged. In addition, they showed that SRT learning is unaffected by changing the sequence of finger movements as long as the location of the response keys remains fixed. If the implicit system in COVIS is procedural-learning-based, then changing the location of the response keys should adversely affect learning in this system, and thus information-integration category learning, whereas changing the finger press associated with each category response should not. On the other hand, hypothesis-testing systems are not typically linked to a specific motor response, and should not be especially sensitive to procedures that change the mapping between category label and response location. Two studies directly tested these hypotheses. Ashby, Ell and Waldron (2003) examined rule-based and information-integration category learning using a training-transfer procedure. There were three conditions: control, hand-switch, and button-switch. In the control condition, the response key assigned to category A was pressed with the left index finger and the response key assigned to category B was pressed with the right index finger during both training and transfer. In the hand-switch condition, the hands were crossed during training so that the response key assigned to category A was pressed with the right index finger and the response key assigned to category B was pressed with the left index finger. During transfer, the hands were uncrossed on the response keys. In the button-switch condition, training was identical to that in the control condition, but during transfer the locations of the buttons were switched. For the rule-based task, hand switching and button switching had no effect on performance. For the information-integration task, on the other hand, button switching led to a decrement in performance, but hand switching did not. These results suggest that the hypothesis-testing system learns abstract category labels, whereas the procedural-learning system learns response positions.

In a related study, Maddox, Bohil and Ing (2005) examined rule-based and information-integration category learning across two conditions. In the fixed response location condition, the response key assigned to category A was pressed with the left index finger, and the response key assigned to category B was pressed with the right index finger. In the variable response location condition, the response key assigned to category A was pressed with the left index finger on half the trials, and with the right index finger on half the trials. In line with the predictions of COVIS, information-integration category learning, but not rule-based category learning, was adversely affected in the variable response location condition.

Experimental manipulations that adversely affect rule-based, but not information-integration learning, have also been observed. Waldron and 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 addition, Maddox et al. (2004a) showed that rule-based category learning was disrupted by a sequential memory-scanning task whereas information-integration category learning was not.

Taken together, these studies provide strong support for the existence of hypothesis-testing and procedural-learning-based systems of category learning and for the neurobiological underpinnings proposed in COVIS. They provide an excellent first step and help lay the groundwork for more detailed examinations of each system. Although several methods are available for studying each system in greater detail, our work has focused on their application to individuals with various neurological conditions. We turn now to a review of this work.

4. Perceptual category learning in neurological patients

In the 1980s and 1990s, one of the most successful models of category learning and recognition memory was exemplar theory [e.g., Nosofsky (1992)]. Exemplar models assume that people access memory traces (perhaps subconsciously) of exemplars when asked to recognize or categorize. This theory is parsimonious because it assumes that the same memory representation underlies both recognition and category learning. If exemplar theory is correct, people with impaired memory storage or consolidation processes should show deficits in recognition and category learning [Pickering (1997)]. Amnesic patients have storage and consolidation problems, and generally have damage to the hippocampus and connected structures (e.g., surrounding medial temporal lobe regions and the diencephalon), and so these patients should show both types of deficits. In a classic study, Knowlton, Squire and Gluck (1994) examined probabilistic classification learning in a group of amnesic patients. The task (referred to as the weather prediction task) required participants to classify stimuli into one of two categories based on the relationship (or association) between multiple stimulus attributes. Specifically, participants were presented with one to three visual cues and asked to predict whether there would be "rain" or "sun." Corrective feedback was provided following each response. There were 14 different combinations of four cues, and each combination was differentially associated with the probability of "rain" or "sun." Knowlton, (1994) found that amnesic patients with damage to the hippocampus or diencephalon performed normally on this "weather prediction" task (at least for the first 50 trials), whereas these same patients were impaired when asked specific questions about the learning context (i.e., an explicit memory task). Thus, amnesic patients were able to learn categories but were unable to recall consciously the circumstances surrounding their learning, suggesting that the hippocampus does not mediate early category learning.

It is important to note that amnesiacs showed normal category learning during the first 50 trials, but they did not perform as well as controls later in learning (i.e., during the last 200 trials). Because only 14 unique cue-stimulus combinations were utilized, Knowlton et al. (1994) suggested that this "late-training deficit" resulted because normal controls used explicit memory for the stimuli that arose from multiple stimulus presentations, whereas the amnesic patients were unable to use such information [for a related explanation see Gluck, Oliver and Nyers (1996)]. Recent work also suggests that a number of qualitatively different strategies can be used to solve this task, ranging from strategies involving attention to a single stimulus attribute to the optimal strategy, which involves attention to all attributes [Gluck, Shohamy and Myers, 2002; see also Ashby and Maddox (2005)].

The Knowlton et al. (1994) study is important because it was one of the first to suggest that category learning and recognition memory might be mediated by different neural substrates. Even so, there are at least two problems with this study. First, the use of only 14 unique cue–stimulus combinations is problematic. This small number of stimuli allows a participant with intact explicit memory to use explicit memory processes to improve categorization performance. To control for the possibility that explicit memory processes will be invoked, and to provide a better test of category learning in amnesia, categories that contain a large number of stimuli should be used. Second, the fact that a large number of qualitatively different strategies can be used to accurately solve the task is problematic. A better approach would be to use a task in which a single, uniquely identifiable, optimal rule (i.e., the rule that maximizes long-run accuracy) can be identified, and for which other strategies yield worse accuracy. Similarly, it would be advantageous to utilize a model-based approach to identify the types of strategies that are being used and to help localize the cognitive processes that lead to any performance decrement.

4.1. Nonlinear information-integration category learning in amnesia

Filoteo, Maddox and Davis (2001a) took such an approach to study category learning in amnesic patients. They utilized the perceptual categorization task [also called the general recognition randomization technique; Ashby and Gott (1988)], which has been used extensively to study category learning in healthy young adults and attentional processes in healthy older adults and patients with PD [see Maddox and Filoteo (in press) for a review]. In a typical perceptual categorization task, participants are presented with simple stimuli, such as a horizontal and a vertical line connected at the upper left corner (see Figure 2), and are asked to assign the stimuli to one of two categories. Prior to the experiment, two bivariate normally distributed categories are specified, and a large number of stimuli are sampled randomly from each bivariate normal distribution. In the Filoteo et al. (2001a) study, 50 unique stimuli were sampled from category A and 50 from category B. With such a large number of unique stimuli, explicit memory processes cannot easily be invoked to improve categorization performance. Because the stimuli are two-dimensional, a unique point in a two-dimensional space can represent each one. Figure 2 depicts the distribution of stimuli used in this study in this two-dimensional space, where the x-axis represents the length of the horizontal line and the y-axis the length of the vertical line. Open squares denote category A stimuli and filled circles denote category B stimuli.

Because the categories are normally distributed, they overlap and a single experimenter-defined categorization rule (i.e., the rule that maximizes long-run accuracy) can be derived [e.g., Maddox and Ashby (1993)]. The form of the rule is determined by the relationship between the two category distributions and thus depends on the relationship between the two stimulus attributes. Filoteo et al. (2001a) examined a complex information-integration categorization rule that was based on a highly nonlinear relationship between the two stimulus attributes. The broken quadratic curve in Figure 2 denotes the optimal categorization rule (or boundary), and yields 95% correct responses.

Because the category structures are defined *a priori*, the experimenter has a great deal of control over potentially important aspects of the categories, such as the optimal accuracy rate and the shape of the optimal categorization rule (e.g., linear or nonlinear), to name a few. An additional advantage of the perceptual categorization task is that a number of quantitative models of category learning have been developed specifically for application to data collected in this task [Maddox and Filoteo (in press)]. Categorization accuracy (i.e., percent correct) is the typical metric used in neuropsychological studies of category learning, and although its use has several strengths, accuracy analyses have at least two weaknesses. First, because accuracy analyses generally focus on averaged performance (e.g., ANOVA), important individual differences may be obscured. The model-based approach utilized by Filoteo et al. (2001a), on the other hand, allows one to identify and quantify performance at the individual participant level. Second, accuracy-based analyses do not allow the researcher to tease apart the separate effects of various cognitive processes on performance. For example, categorization accuracy is affected not only by participants' ability to learn the experimenter-defined categorization rule, but also by their ability to accurately apply the learned rule on each trial[3]. The first process we refer to as *categorization rule learning*. Difficulty learning the experimenter-defined categorization rule (denoted by the broken curve in Figure 2) will lead to a reduced accuracy level. The second process we refer to as *rule application variability*. This has to do with participants' ability to apply consistently from trial to trial whatever categorization rule they might have learned. Greater variability in rule application can also lead to reduced accuracy. Both processes, categorization rule learning and rule application variability, will affect accuracy measures, and thus at the level of accuracy these two processes are nonidentifiable. The model-based approach utilized by Filoteo et al. (2001a) alleviates this problem because it allows one to *separate* categorization rule learning from rule application variability. The modeling approach will be summarized briefly after we review the experimental findings.

Filoteo et al. (2001a) had two amnesic patients and five matched controls complete six 100-trial blocks of trials in the perceptual categorization task using the Figure 2 category structures. On each trial a stimulus was selected at random and was presented on the computer screen, the participant generated a category A or category B response, and corrective feedback was provided.

The top panel of Figure 3 displays the proportion correct for the amnesic and control participants during the first 100 trials and the final 100 trials (i.e., 501–600) from the first experimental session. One of the amnesic patients and a matched control also completed a second session and the data from the first 100 trials are also presented in Figure 3. Several comments are in order. First, during the first and final blocks from Day 1, the amnesic patients and controls showed equivalent performance. In fact, performance did not differ in any of the six blocks of trials. This finding is important

[3] We are using the term "rule" more generally here than in COVIS. In the current application, the "rule" might be verbalizable or nonverbalizable. It might involve learning a decision bound or assigning responses to regions of perceptual space

because it suggests that amnesiacs can learn to categorize, and that the late training deficit observed in the weather prediction task was likely due to the use of explicit memory processes by the control participants. Second, during the first block of trials from Day 2, the amnesic patient and the control again showed equivalent performance, and in fact, performance during the first block of the second session was slightly better than that during the final block of trials from the first session. Some have suggested that amnesic patients learn categorization rules using working or short-term memory processes [Nosofsky and Zaki (1999)]. For example, it has been suggested that amnesic patients are able to take advantage of the repeating stimuli during some categorization tasks and this information is then used to categorize further stimuli [Nosofsky and Zaki (1999)]. The Day 2 results argue against this possibility because it is highly unlikely that subjects were able to make use of working or short-term memory processes between the 2 days. Instead, these findings indicate that the categorization rule was retained over the 1-day delay period, and given the severity of the memory deficit in our amnesic patient (e.g., on Day 2, the patient did not recall having been given the test on the previous day), brain systems not involved in explicit memory likely mediated the retention of this categorization rule.

4.1.1. Model-based analyses

The specifics of the modeling procedure are outlined in numerous articles [e.g., Ashby (1992), Maddox and Filoteo (in press)]. In this section we provide only an overview of the approach, highlighting aspects of the modeling that are relevant to the Filoteo et al. (2001a) study. Filoteo et al. tested amnesic patients and controls in their ability to learn a rule in which correct classification was based on a unique nonlinear (quadratic) relationship between the horizontal and vertical line lengths. This rule is depicted as the broken curve in Figure 2. The aim of the modeling approach with these data was twofold. First, to determine how well a participant learned the optimal decision rule, we fit the optimal decision bound model to each block of data separately for each participant[4]. As a measure of *categorization rule learning* we examined the goodness-of-fit value (i.e., the

[4] In two dimensions, a quadratic function takes the form $ax^2 + by^2 + cxy + dx + ey + f$, where a to f denote the coefficients of the quadratic function, and x and y denote the horizontal and vertical line lengths, respectively. In the experimenter-defined optimal quadratic categorization rule, the coefficients, a to f, are fixed, and are determined from the category structures. Maximum likelihood criteria were used to estimate model parameters [see Ashby (1992)]. In essence, the maximum likelihood procedure attempts to maximize the "fit" of the model to the data by attempting to generate predictions from the model that most closely match the observed data. In Filoteo et al. (2001a), the data were the participant's categorization responses for each presented exemplar. Thus, for each exemplar the observed probability of responding "Category A" was either 1 or 0. Assuming the optimal categorization rule is applied, and for a fixed value of the rule application variance (an estimate of the variability associated with a participant's inability to accurately apply the same rule on every trial), the model generated a predicted probability of responding "Category A" for each exemplar. Because the coefficients are fixed in the optimal model, the rule application variance was the only parameter adjusted iteratively until the difference between the observed and predicted "Category A" response probabilities was minimized.

maximum likelihood value, -lnL, negative log likelihood) from the optimal model. The smaller the fit, the better the optimal rule describes the data. Second, we were interested in quantifying the magnitude of any variability in the application of the participant's rule. To achieve this goal we fit a suboptimal model that assumed a quadratic decision bound, but allowed the coefficients of the quadratic decision bound to be estimated from the data. As a measure of *rule application variability*, we examined the noise variability estimate from this suboptimal model. The smaller the magnitude of the rule application variability, the less variable was the participant's trial-by-trial application of the rule. To reiterate, at the level of accuracy rates, these very different processes are nonidentifiable. Only with the model-based approach can these two subprocesses be teased apart and be made identifiable.

Because no accuracy differences were observed across amnesic and control participants, Filoteo et al. (2001a) predicted no differences in categorization rule learning or rule application variability. None were observed. Since publication of the Knowlton et al. (1994) and Filoteo et al. (2001a) studies showing intact category learning in amnesia, several challenges to these findings have been offered. For example, Nosofsky and Zaki (1998) suggested that exemplar theory can be used to account for the category learning/recognition memory dissociation observed in amnesia. Smith and Minda (2000) argued against these claims. Although an interesting topic of further work, our laboratories have shifted attention away from amnesia toward a study of category learning in striatal-damaged patients. COVIS does not implicate the medial temporal lobes directly in category learning and so the Filoteo et al. findings were expected. COVIS does predict that the striatum is directly involved in category learning and so an examination of striatal contributions to category learning is in order.

4.2. Nonlinear information-integration category learning in striatal-damaged patients

Knowlton, Mangels and Squire [1996a; see also Knowlton et al. (1996b)] suggested that the striatum may play an important role in category learning. They found that patients with PD, whose neuropathology results in a decrement in striatal functioning, demonstrate impaired probabilistic classification learning in the weather prediction task but intact recognition memory. These observations are also consistent with animal studies that implicate the striatum in certain aspects of category learning [McDonald and White (1993), Packard and McGaugh (1993)]. Notice that these data, along with those from amnesic patients, represent a double dissociation. PD patients show a deficit in category learning but intact recognition memory, whereas amnesic patients show intact category learning but a deficit in recognition memory. These results provide strong evidence for the involvement of the striatum, but not the medial temporal lobes, in information-integration category learning.

Although the Knowlton et al.(1996a) study suggests that the striatum is involved in category learning, other reports argue against this position. One important study was conducted by Reber and Squire (1999), who found that PD patients performed normally in learning to classify dot patterns and artificial grammars. These results contradict the findings of Knowlton et al. (1996a), who demonstrated that PD patients were impaired in probabilistic classification. In an attempt to reconcile these findings, Reber and Squire (1999) suggested that probabilistic classification is different from artificial grammar or dot pattern classification because the participants must learn the cue-outcome relations through trial-by-trial feedback [for an alternative explanation, see Ashby and Maddox (2005)]. In artificial grammar and dot pattern classification, on the other hand, individuals are simply exposed to members of a category, and are required to study these items. They are then tested on items that either "fit" or "do not fit" the trained category structure. Reber and Squire argue that the need to learn cue-outcome relations in probabilistic classification requires the striatal learning system, which is impaired in PD. Since dopamine is dramatically reduced in the striatum of patients with PD [Cornford, Chang and Miller (1995)], this interpretation would also be consistent with the proposed role of dopamine in reward-based learning mechanisms [Ashby et al. (1998)].

It is very likely that the need to learn cue-outcome relations in probabilistic classification at least partially accounts for the poor performance of PD patients. Unfortunately, the artificial grammar and dot pattern classification tasks differ from the probabilistic classification task in a number of ways, any of which might fully or partially explain the performance dissociation observed in PD. First, the artificial grammar and dot pattern classification tasks usually require the learning of only a single category, whereas two categories must be learned in the probabilistic classification task [Ashby and Maddox (2005)]. Second, perfect performance is possible in the artificial grammar and dot pattern classification tasks, whereas it is not possible in the probabilistic classification task. Third, all three tasks differ with regard to the nature of the stimuli (dot patterns, artificial words, or cards with geometric forms), and the response requirements (point at the center dot, reproduce the artificial word on a piece of paper, choose a category "rain" or "sun"). Finally, the nature of the categorization rule is very different across the three tasks. Dot pattern classification can be solved by prototype extraction, and the artificial grammar task might involve perceptual priming of letter string chunks [Knowlton et al. (1996b)]. The probabilistic classification task, on the other hand, appears to involve the learning of a complex categorization rule [although simpler strategies will suffice; Gluck et al. (2002)]. Thus, the performance dissociation could be due to any one of the following: reliance on striatal learning, effects of maximum attainable accuracy, differential stimulus characteristics and task requirements, categorization rule complexity, or any combination of these factors.

Because of these problems, we decided to examine striatal involvement in category learning using the perceptual categorization task. Maddox and Filoteo (2001) had PD patients and matched controls complete six 100-trial blocks of trials using the nonlinear information-integration category stimuli displayed in Figure 2, and Filoteo, Maddox and Davis (2001b) had HD patients and matched controls perform the same task. Like PD, HD also impacts striatal functioning [Vonsattel et al. (1985)].

The asymptotic accuracy rates obtained during the final block of trials (i.e., trials 501–600) for the PD, HD, and relevant control participants are depicted in the middle panel of Figure 3. Note that both the PD and HD participants showed clear category learning deficits. To determine the locus of the nonlinear information-integration category learning deficit in PD and HD participants, we examined the categorization rule learning and rule application variability estimates from the final block of trials. These values are displayed in the bottom two panels of Figure 3. The PD patients evidenced categorization rule learning and rule application variability deficits suggesting that their accuracy deficit was due to an inability to learn the optimal rule, and to greater variability in the application of the rule that they had learned. The HD patients showed categorization rule learning deficits but not rule application variability deficits (although the trend is in that direction) suggesting that their performance deficit was due to an inability to learn the optimal rule.

Taken together with the Filoteo et al. (2001a) study, the results support the prediction that the striatum, but not the medial temporal lobes, is involved in nonlinear information-integration category learning when a large number of unique stimuli are utilized to minimize the influence of explicit memory processes. This follows since information-integration category learning is assumed to be mediated within the tail of the caudate, which is impacted in HD, and involves a dopamine-mediated reward signal, which is likely deficient in patients with PD. The results also suggest that the locus of the PD and HD subjects' deficits was in their ability to learn the optimal decision bound, with the additional difficulty, for PD patients only, of accurately applying the rule that they have learned.

4.3. Rule-based category learning in PD

Maddox and Filoteo (2001) also examined rule-based category learning in PD using the perceptual categorization task and the same two-line stimuli. Their task required the participant to attend to both stimulus dimensions, and to use the following rule: respond A if the vertical line is longer than the horizontal line and respond B if the vertical line is shorter than the horizontal line. They found that PD patients were as good as controls at performing this task[5].

This finding is at odds with the predictions of COVIS and of a recent study conducted by Ashby et al. (2003b). COVIS predicts that PD patients will show impaired rule-based category learning because of depleted dopamine projections from the substantia nigra into the head of the caudate nucleus. As a test of this hypothesis, Ashby et al. had PD patients and controls learn a rule-based category structure. The stimuli were constructed from four highly discriminable binary dimensions (e.g., red vs. blue, circle vs. square, etc.), factorially combined to create a set of 16 stimuli. One dimension was chosen to be relevant while the three remaining dimensional values varied randomly across trials. Using a learning criterion of 10 correct responses in a row, Ashby et al. found that ~50% of the PD patients failed to learn this rule-based task within 200 trials, whereas only 10% of the controls failed to learn, suggesting a PD deficit in rule-based category learning.

[5] Filoteo et al. (2001b) examined HD patient learning in the same condition and found a small deficit in these patients.

The findings of Ashby et al. (2003b) and Maddox and Filoteo (2001) challenge the simplistic notion that PD patients will always show a rule-based category learning deficit.

Unfortunately, the Maddox and Filoteo (2001) and Ashby et al. (2003b) studies differ in a number of important ways, each of which might explain the contradictory findings. For example, the Maddox and Filoteo study used a large number of overlapping, continuous-valued dimension stimuli that required that a unique decision criterion be learned but did not require that any dimensional information be filtered. On the other hand, the Ashby et al. study used a small number of nonoverlapping, binary-valued dimension stimuli that did not require that a unique decision criterion be learned, but required that variation along three irrelevant dimensions be filtered.

To begin to shed some light on the locus of potential PD deficits in rule-based category learning, Maddox et al. (2005) examined PD patients' learning in two rule-based category learning conditions. Like Maddox and Filoteo (2001), they used the perceptual categorization task. Thus, they used a large number of overlapping, continuous-valued dimension stimuli that required that a unique decision criterion be learned (similar to that in Figure 1a, but with different stimulus dimensions). However, unlike Maddox and Filoteo (2001), who required attention to both stimulus dimensions, Maddox et al. (2005) required the participant to learn a decision criterion along one stimulus dimension while filtering out (or ignoring) information from a second stimulus dimension. Each participant completed five 50-trial blocks of trials in each condition. For ease of exposition we focus on asymptotic performance collapsed across the two rule-based conditions.

The asymptotic accuracy rates obtained during the final block of trials (i.e., trials 201–250) for the PD and control participants are depicted in the top panel of Figure 4. Notice that the PD patients showed a clear rule-based category learning deficit. To determine the locus of the rule-based category learning deficit in PD, we first applied a series of models to determine whether the PD patients' deficit was due to an inability to ignore variation along the irrelevant dimension. In other words, were PD patients more likely than controls to use a strategy that was qualitatively different from the optimal strategy – namely an information-integration strategy, when the optimal strategy was to attend selectively to only one dimension? To test this hypothesis, we fit three models. The optimal rule-based model assumes that the participant attended selectively and used the optimal decision criterion. The suboptimal rule-based model assumes that the participant attended selectively, but used a suboptimal decision criterion. The information-integration model assumes that the participant was unable to attend selectively and instead used a linear decision bound constructed from a weighted linear combination of the two dimensional values[6]. The results were clear. For 84% of the PD patients and

[6] The models are nested in the sense that the suboptimal rule-based model can be derived from the information-integration model by setting the slope of the information-integration model to the optimal slope (zero or infinity), and the optimal rule-based model can be derived from the suboptimal rule-based model by setting the decision criterion to the optimal value. Using nested modeling techniques [see Ashby (1992)], we identified the best-fitting model, defined as the simplest model for which a more general model did not provide a statistically significant improvement.

81% of the control participants, a rule-based model provided the best fit of the data suggesting that the PD deficit was not due to a bias towards integrating information (i.e., using a qualitative different strategy from that of the optimal classifier) when the optimal strategy was to attend selectively.

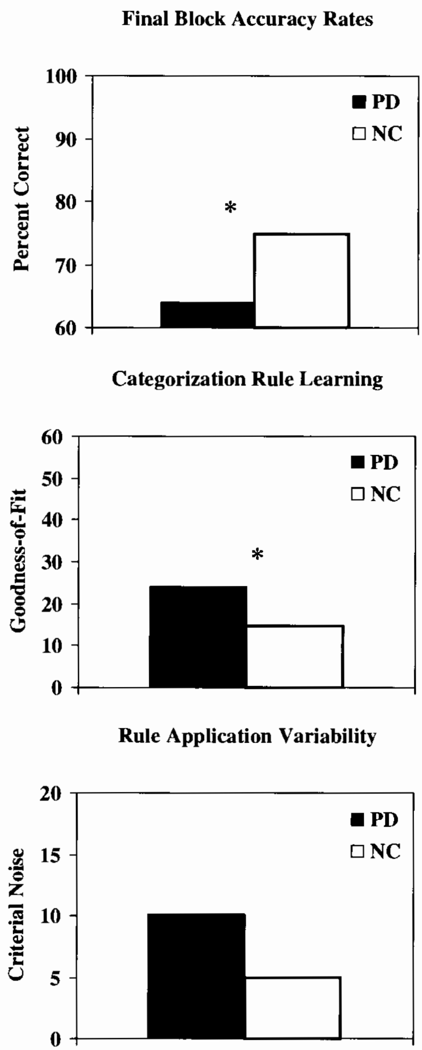

Fig. 4. Accuracy rates, categorization rule learning, and rule application variability estimates from a rule-based category learning study conducted by Maddox et al. (2005) (see text for details). *Statistically significant performance difference ($p<0.05$).

To further examine the locus of the rule-based category learning deficit in PD, we examined the categorization rule learning and rule application variability estimates from the final block of trials. The categorization rule learning index was determined from the fit of the optimal rule-based model, and the rule application variability index was determined from the noise parameter estimate from the suboptimal rule-based models. These values are displayed in the bottom two panels of Figure 4. The PD patients evidenced categorization rule learning deficits, but not rule application variability deficits (although this difference was marginally significant).

These analyses suggest that PD patients were as likely to select a task-relevant rule as control participants, but their use of this rule was less optimal. One way in which the rule might be used suboptimally is in the placement of the decision criterion. To assess this question, we examined the decision criterion estimates from the suboptimal rule-based model. Specifically, we examined the absolute deviation between the best-fitting decision criterion and the optimal decision criterion. By using the absolute deviation, the analyses assess the extent of a response bias independent of the direction of that bias. For the PD patients, the absolute deviation from optimal was 34 pixels, whereas for controls the absolute deviation was 20 pixels, a reduction of nearly 50%. Thus, PD patients tended to exhibit larger suboptimalities in decision criterion placement than the control participants.

It is worth mentioning that a group of patients with cerebellar lesions (CB) were also tested. CB patients showed no deficits in accuracy, categorization rule learning, or rule application variability, suggesting that the cerebellum is not involved in rule-based category learning.

The Maddox et al. (2005) study offered a straightforward extension of Maddox and Filoteo (2001) to a situation in which irrelevant dimensional information had to be filtered. In a related study, Filoteo et al. (in press) offered a straightforward extension of Ashby et al. (2003b). In Ashby et al. (2003b) 16 four-dimensional, highly discriminable, binary-valued dimension stimuli were used. One dimension was selected as the relevant one, and the other three were irrelevant but varied randomly across trials. Ashby et al. (2003b) found a large performance deficit for PD patients relative to matched controls. Filoteo et al. (in press) were interested in determining whether the number of randomly varying irrelevant dimensions might impact the magnitude of the rule-based category learning deficit in PD. In all conditions, the stimuli were four-dimensional, one dimension was relevant to solving the task, and the other three were irrelevant. However, the number of irrelevant dimensions that could vary randomly was manipulated across conditions. In one condition, all three irrelevant dimension values could vary randomly across trials. This is analogous to the Ashby et al. (2003b) condition. In a second condition, two of the three irrelevant dimension values varied randomly and the third was held fixed across trials. In the third and fourth conditions, one and none of the irrelevant dimension values varied randomly with the remaining ones (two or three, respectively) held fixed across trials. Using the Ashby et al. (2003b) learning criterion of 10 correct responses, Filoteo et al. (in press) found that when no irrelevant dimensional variation occurred, PD patients and controls did not differ in the number of trials it took to learn the criterion, but as the number of irrelevant dimensions increased, PD patients' rule-based learning was impacted to a greater extent than that of controls.

Clearly more work is needed to fully understand the properties that affect rule-based category learning in PD. Even so, this collection of studies provides a nice starting point and suggests that the need to filter irrelevant dimension information is predictive of rule-based category learning deficits in PD for both continuous and binary values stimuli.

4.4. Further study of information-integration category learning in PD

Recall that Maddox and Filoteo [2001; see also Filoteo et al. (2001b)] examined PD patients' ability to learn a nonlinear information-integration categorization rule using the perceptual categorization task. The stimulus was a horizontal and vertical line connected at the upper left corner that varied in length across trials (see Figure 2). A large number of unique stimuli were presented and the categories overlapped. PD patients showed a consistent performance deficit across six 100-trial blocks. Ashby et al. (2003b) also examined information-integration category learning using the same 16 four-dimensional, highly discriminable, binary-valued dimension stimuli used to study rule-based category learning. The only difference was that the stimulus-to-category assignments were modified to construct an information-integration condition. In their information-integration condition, one dimension was irrelevant, and category assignment was based on the combination of information from the three remaining stimulus components. Using a learning criterion of 10 correct responses in a row, Ashby, Noble, et al. found that similar percentages of PD patients and controls (50%) failed to learn the task, suggesting that PD does not result in a deficit in information-integration category learning.

As with the rule-based tasks described above, the results from Maddox and Filoteo (2001) and Ashby et al. (2003b) lead to different conclusions. However, the tasks differ along a number of important dimensions, making it difficult to determine the locus of these contradictory findings. One interesting aspect of the results is the difference in complexity or difficulty of the two tasks. In Ashby et al. (2003b), PD patients and controls who learned the task learned it in 80 trials on average. In Maddox and Filoteo (2001), on the other hand, PD patients showed a deficit relative to the controls across all 600 trials, and never reached 80% accuracy even after 600 trials. These data suggest that the information-integration rule used by Maddox and Filoteo (2001) is more complex and that this complexity might impact the likelihood of observing an information-integration category learning deficit in PD. Even so the tasks differ along too many other dimensions to make any definitive claims.

To further explore the effects of information-integration rule complexity on PD patients' category learning, Filoteo et al. (2005) tested a group of PD patients and healthy elderly controls using the linear and nonlinear information-integration conditions displayed in Figures 1b and 1c, respectively. The stimulus was a line that varied in length and orientation across trials. Each participant completed six 100-trial blocks of trials in each condition. It is important to reiterate that, because this work is couched within the framework of the perceptual categorization task, a number of important factors are equated across conditions (e.g., optimal accuracy), while only the form of the optimal decision bound is manipulated.

Before summarizing the results of this study, let us first describe the methods we used to model the data. Since we published our early work on category learning in amnesia, PD, and HD [Filoteo et al. (2001a,b), Maddox and Filoteo (2001)], the focus of our modeling approach has changed slightly. Instead of focusing on estimates of categorization rule learning and rule application variability, we have begun attempting to characterize the strategy that participants are actually using. One conclusion that we have drawn from our parallel work using healthy young adults is that participants often try to solve information-integration tasks using hypothesis-testing strategies when the experimental conditions are not conducive to learning with the procedural-learning system [see Maddox and Ashby (2004) for a review]. This might also occur with PD patients. Since the neurobiological machinery necessary to solve information-integration tasks is deficient in PD, it might be the case that PD patients attempt to use hypothesis-testing strategies. To investigate this possibility, we developed a large number of models that were applied to the data from each block of trials separately for each participant. Some of these models were hypothesis-testing models and some were information-integration models.

Figure 5 displays hypothetical decision bounds and the resulting response regions from specific response strategies that might be applied in the linear information-integration condition. The four models in the leftmost column are hypothesis-testing models and the three on the right are information-integration models. The top two hypothesis-testing models instantiate *one-dimensional* strategies. One assumes that the participant sets a criterion on length and ignores orientation, whereas the other assumes that the participant sets a criterion on orientation and ignores length. The bottom two hypothesis-testing models instantiate *two-dimensional, conjunctive* strategies. Each assumes that the participant sets a criterion on the length dimension and a separate criterion on the orientation dimension. In the first case, the participant responds A if the length is "short" and the orientation is "steep," and B otherwise. In the second case, the participant responds B if the length is "long" and the orientation is "shallow," and A otherwise. The topmost information-integration strategy assumes that the participant uses the optimal decision bound. The middle model assumes that the participant uses a suboptimal linear decision bound, but allows the slope and intercept to be suboptimal. The bottom model assumes that the participant uses two linear decision bounds[7].

The hypothesis-testing models applied to the linear information-integration condition (Figure 5) were also applied in the nonlinear information-integration condition.

[7] This model is called the *Striatal Pattern Classifier* [(SPC; Ashby and Waldron (1999)] and was developed as a computational model of the tail of the caudate. The model assumes that there are four "units" in the length-orientation space with two being assigned to category A and two to category B. On each trial the observer determines which unit is closest to the perceptual effect and gives the associated response. The model results in two "minimum-distance-based" decision bounds. This model has been found to provide a good computational model of observers' response regions in previous information-integration category learning studies [(e.g., Ashby and Waldron (1999), Maddox et al. (2004)].

The information-integration models were identical as well except that the optimal and suboptimal models assumed quadratic bounds.

4.4.1. Brief summary of the results

The asymptotic accuracy rates obtained during the final block of trials (i.e., trials 501–600) for the PD patients and controls are depicted in the top panel of Figure 6. Notice that PD patients showed normal linear information-integration learning, but a deficit in nonlinear information-integration learning. Importantly, PD patients' deficit in the

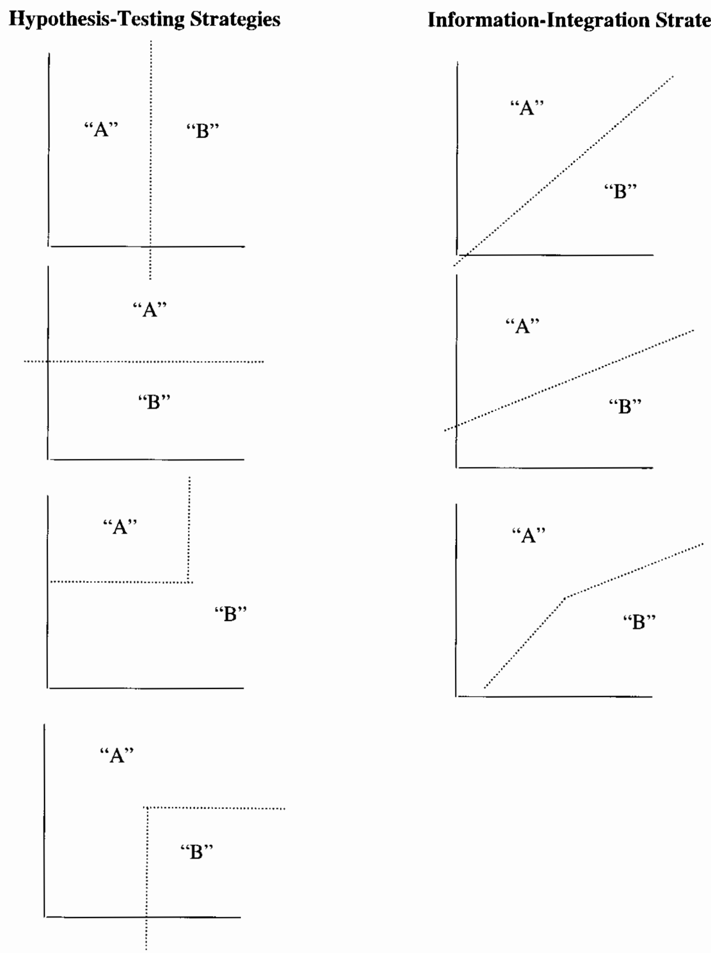

Fig. 5. Hypothetical response regions from participants using hypothesis-testing and information-integration strategies to solve the linear information-integration task.

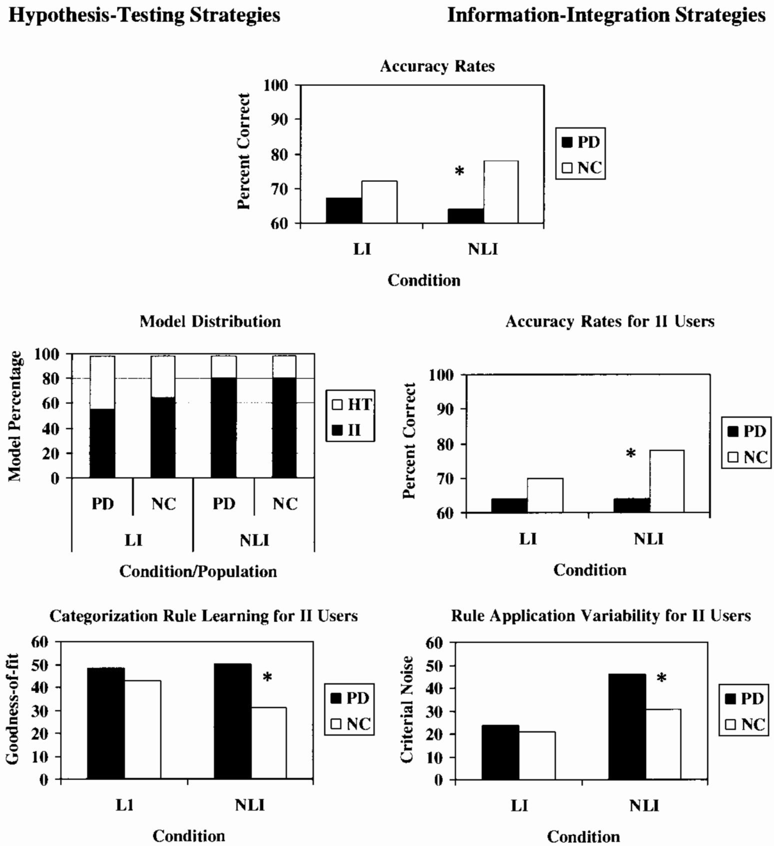

Fig. 6. Final block overall accuracy rates, model distributions, accuracy rates, categorization rule learning, and rule application variability estimates for participants, whose data were best fit by an information-integration model from Filoteo et al. (2004) (see text for details). *Statistically significant performance difference ($p<0.05$).

nonlinear condition, but not the linear condition, did not appear to be due to task difficulty per se in that the controls performed somewhat better in the nonlinear condition relative to the linear condition.

The results of the model analyses can also be seen in Figure 6. The middle left panel shows the percent of participants whose final block data were best fit by an information-integration model or a rule-based model. Notice that the model percentages are quite similar across PD and NC participants for the linear and nonlinear information-integration conditions, but many fewer PD and NC participants attempted to use hypothesis-testing strategies in the nonlinear condition. To gain additional insight into the locus of the PD patients' nonlinear information-integration learning deficit, we focused only on participants whose data were best fit by an information-integration model. We computed the accuracy rate for these participants (displayed in the middle right panel of Figure 6), as well as the categorization rule learning and rule application variability indices (displayed in the bottom two panels of Figure 6). These analyses suggest that PD patients, who used information-integration strategies to solve the linear information-integration task were as accurate as controls, and showed equivalent rule learning, and equivalent variability in the application of their rule. For PD patients, who used information-integration strategies to solve the nonlinear information-integration task, on the other hand, accuracy was lower, categorization rule learning was poorer, and variability in the application of their rule was higher than that for controls.

Taken together these results suggest that PD patients are relatively normal at solving linear information-integration tasks, but show clear and large deficits in their ability to solve more complex nonlinear information-integration tasks. Although speculative at this point, it seems reasonable to suppose that the "resolution" of the perceptual-decision space in the tail of the caudate must be higher to solve more complex nonlinear information-integration tasks than to solve simpler linear information-integration tasks. It also seems reasonable to suppose that the resolution in the caudate of PD patients is not as high as in controls (possibly due to a weaker dopamine reward signal), leading to difficulty with complex decision rules, but not with simpler ones. Clearly, more work is needed to determine exactly why PD patients are more susceptible to impairment when learning complex information-integration rules.

5. General discussion

The work presented in this chapter builds upon that presented in the previous chapter by Ashby and Valentin. They reviewed the neurobiological underpinnings of a multiple systems model called COVIS. COVIS assumes that learning in rule-based tasks is dominated by an explicit 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, whereas 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, and requires a dopamine-mediated reward signal. Ashby and Valentin reviewed a number of studies conducted using healthy young adults that tested and provided support for several *a priori* predictions derived from an examination of the neurobiological underpinnings of the two systems.

The current chapter reviewed a body of literature concerning studies conducted in our laboratories that provides a more detailed examination of the systems used by neurological patients with damage to the medial temporal lobes or the striatum. In addition, a quantitative model-based approach was taken in most of these studies that allowed us to tease apart the separate effects of various cognitive processes on performance that are nonidentifiable at the level of accuracy.

The ability of medial temporal lobe amnesic and striatal-damaged patients (PD and HD) to solve a nonlinear information-integration task was examined in three studies [Filoteo et al. (2001a,b); Maddox and Filoteo (2001)]. In each case, a large number of stimuli were presented to alleviate the possibility that participants might recruit explicit memory processes, and participants completed many experimental trials. As predicted by COVIS, medial temporal lobe amnesiacs showed no performance deficit throughout the learning session, whereas both PD and HD patients showed a consistent deficit. Because the optimal nonlinear information-integration rule was unique, we used a model-based approach to determine whether the locus of the accuracy deficit in PD and HD was an inability to learn the rule (categorization rule learning), an inability to apply consistently from trial to trial whatever categorization rule they might have learned (rule application variability), or both. PD and HD patients evidenced deficits in categorization rule learning and rule application variability suggesting that damage to the striatum affects both subprocesses.

The ability of PD patients to solve rule-based category learning tasks was examined across two pairs of related studies [Maddox and Filoteo (2001), Ashby et al. (2003b), Filoteo et al. (2004), Maddox et al. (2005)]. One pair of studies used the perceptual categorization task that utilizes a large number of unique continuous-valued stimuli sampled from overlapping bivariate normally distributed categories. In Maddox and Filoteo (2001), the optimal rule-based strategy required the participant to attend to both stimulus dimensions, whereas in Maddox, et al. (2005) the participant was required to attend to only one stimulus dimension while filtering out (or ignoring) information about the second. The two experiments were identical in all other important aspects. When no filtering was required, PD patients performed normally on rule-based category learning, but when filtering was required they showed an accuracy deficit. This accuracy deficit was not due to PD patients' inability to attend selectively (PD patients used rule-based strategies to the same degree as controls), but rather to the use of a suboptimal decision criterion. These data suggest that PD patients show deficits when the rule-based task requires dimensional filtering.

Unlike the first pair of studies, the second pair of studies utilizes a small number of binary-valued stimuli that were highly discriminable (i.e., no category overlap). Ashby et al. (2003b) used 16 stimuli composed of four binary-valued dimensions. One dimension was relevant to solving the task and the other three were irrelevant. PD patients showed a large rule-based category-learning deficit. Filoteo et al. (in press) utilized similar stimuli and a one-dimensional rule-based category structure, but manipulated the number of irrelevant dimensions, which varied across trials from three [as in Ashby et al. (2003b)] down to zero, with the remaining dimensional values held fixed across trials. Filoteo et al. (2004) found that the magnitude of the PD deficit increased as the number of randomly varying irrelevant dimensions increased. Specifically, when zero or one irrelevant dimension varied, PD patients showed normal rule-based category learning, but when two dimensions varied, PD patients evidenced a rule-based category learning deficit.

The chapter ended with a review of some recent work that examined the effect of the complexity of the information-integration categorization task on PD performance.

Complexity was examined by comparing a linear information-integration rule with a nonlinear information-integration rule. A major focus was also on identifying the types of response strategies that PD patients and controls used to solve these problems. Work with healthy young adults suggests that hypothesis-testing strategies are often used to solve information-integration category-learning problems when the learning situation is non-optimal. Because PD patients have depleted dopamine, it is reasonable to suppose that they might show similar effects. PD patients showed normal category learning when the optimal decision bound was linear, but a deficit in category learning when the bound was nonlinear. The nonlinear information-integration deficit in PD was not due to an increase in the use of hypothesis-testing strategies (nearly the same proportion of PD and control participants used this type of strategy), but rather was due to worse performance by those PD patients using information-integration strategies relative to controls. In summary, the studies reviewed in this chapter represent a first step toward a more detailed understanding of the neurobiological underpinnings of the two category learning systems proposed in COVIS.

References

Ashby, F.G. (1992), "Multivariate probability distributions", in: F.G. Ashby, ed., Multidimensional Models of Perception and Cognition (Lawrence Erlbaum Associates, Inc., Hillsdale, NJ) 1–34.

Ashby, F.G., L.A. Alfonso-Reese, A.U. Turken and E.M. Waldron (1998), "A neuropsychological theory of multiple systems in category learning", Psychological Review 105:442–481.

Ashby, F.G., and S.W. Ell (2001), "The neurobiological basis of category learning", Trends in Cognitive Sciences 5:204–210.

Ashby, F.G., and S.W. Ell (2002), "Single versus multiple systems of learning and memory", in: J. Wixted and H. Pashler, eds, Stevens' Handbook of Experimental Psychology, Third Edition: volume 4: Methodology in Experimental Psychology (Wiley, New York) 655–692.

Ashby, F.G., S.W. Ell and E.M. Waldron (2003a), "Procedural learning in perceptual categorization", Memory & Cognition 31:1114–1125.

Ashby, F.G., and R.E. Gott (1988), "Decision rules in the perception and categorization of multidimensional stimuli", Journal of Experimental Psychology: Learning, Memory, and Cognition 14:33–53.

Ashby, F.G., and W.T. Maddox (2005), "Human category learning", Annual Review of Psychology. 56:149–78.

Ashby, F.G., W.T. Maddox and C.J. Bohil (2002), "Observational versus feedback training in rule-based and information-integration category learning" Memory & Cognition 30:666–677.

Ashby, F.G., S. Noble, J.V. Filoteo, S.W. Ell and E.M. Waldron (2003b), "Category learning deficits in Parkinson's Disease", Neuropsychology 17:115–124.

Ashby, F.G., and E.M. Waldron (1999), "The nature of implicit categorization" Psychonomic Bulletin & Review 6:363–378.

Cornford, M.E., L. Chang and B.L. Miller (1995), "The neuropathology of Parkinsonism: an overview", Brain and Cognition 28:321–341.

Erickson, M.A., and J.K. Kruschke (1998), "Rules and exemplars in category learning" Journal of Experimental Psychology: General 127:107–140.

Filoteo, J.V., W.T. Maddox and J.D. Davis (2001a), "Quantitative modeling of category learning in amnesic patients" Journal of the International Neuropsychological Society 7:1–19.

Filoteo, J.V., W.T. Maddox and J.D. Davis (2001b), "A possible role of the striatum in linear and nonlinear category learning: Evidence from patients with Huntington's disease", Behavioral Neuroscience 115:786–798.

Filoteo, J.V., W.T. Maddox, A.D. Ing, V. Zizak and D.D. Song (in press), "The impact of irrelevant dimensional variation on rule-based category learning in patients with Parkinson's disease", Journal of the International Neuropsychological Society.

Filoteo, J.V., W.T. Maddox, D.P. Salmon and D.D. Song (in press), "Information-integration category learning in patients with striatal dysfunction," Neuropsychology.

Gluck, M.A., L.M. Oliver and C.E. Myers, (1996), "Late-training amnesic deficits in probabilistic category learning: a neurocomputational analysis," Learning and Memory 3:326–340.

Gluck, M.A., D. Shohamy and C. Myers (2002), "How do people solve the "weather prediction" task? Individual variability in strategies for probabilistic category learning," Learning and Memory 9: 408–418.

Hazeltine, E., and R.B. Ivry (2002), "Motor skill in V.S. Ramachandran, ed., Encyclopedia of the Brain (Academic Press, San Diego) 183–200.

Kéri, S. (2003), "The cognitive neuroscience of category learning," Brain Research Reviews 43:85–109.

Knowlton, B.J., J.A. Mangels and L.R. Squire (1996a), "A neostriatal habit learning system in humans," Science 273:245–254.

Knowlton, B.J., L.R. Squire and M.A. Gluck (1994), "Probabilistic classification learning in amnesia," Learning and Memory 1:106–120.

Knowlton, B.J., L.R. Squire, J.S. Paulsen, N.R. Swerdlow, M. Swenson and N. Butters, (1996), "Dissociations within nondeclarative memory in Huntington's disease," Neuropsychology 10:169–181.

Maddox, W.T., P. Aparicio, N. Marchant and R.B. Ivry (2005), "Rule-based category learning is impaired in patients with Parkinson's disease but not patients with cerebellar disorders," Journal of Cognitive Neuroscience. 17(5), 707–723.

Maddox, W.T., and F.G. Ashby (1993), "Comparing decision bound and exemplar models of categorization," Perception & Psychophysics 53:49–70.

Maddox, W.T., and F.G. Ashby (2004), "Dissociating explicit and procedural-learning based systems of perceptual category learning," Behavioral Processes 66:309–332.

Maddox, W.T., F.G. Ashby and C.J. Bohil (2003), "Delayed feedback effects on rule-based and information-integration category learning," Journal of Experimental Psychology: Learning, Memory, and Cognition 29:650–662.

Maddox, W.T., F.G. Ashby, A.D. Ing and A.D. Pickering (2004a), "Disrupting feedback processing interferes with rule-based, but not information-integration category learning," Memory & Cognition 32:582–591.

Maddox, W.T., C.J. Bohil and A.D. Ing (2004), "Evidence for a procedural-learning based system in perceptual category learning," Psychonomic Bulletin and Review. 11(5), 945–952.

Maddox, W.T., and J.V. Filoteo (2001), "Striatal contributions to category learning: quantitative modeling of simple linear and complex nonlinear rule learning in patients with Parkinson's Disease," Journal of the International Neuropsychological Society 7:710–727.

Maddox, W.T., and J.V. Filoteo (in press), "Modeling visual attention and category learning in normal aging, amnesiac and striatal-damaged patients". R.W.J. Neufeld (Ed.) Advances in Clinical–Cognitive Science: Formal Modeling and Assessment of Process and Symptons.

Maddox, W.T., J.V. Filoteo, K.D. Hejl and A.D. Ing (2004b), "Category number impacts rule-based but not information-integration category learning: further evidence for dissociable category learning systems," Journal of Experimental Psychology: Learning, Memory, and Cognition 30:227–235.

McDonald, R.J., and N.M. White (1993), "A triple dissociation of memory systems: hippocampus, amygdala, and dorsal striatum," Behavioral Neuroscience 107:3–22.

Nosofsky, R.M. (1992), "Exemplar-based approach to relating categorization, identification, and recognition," in: F.G. Ashby, ed., Multidimensional models of perception and cognition (Erlbaum, Hillsdale, NJ) 363–393.

Nosofsky, R.M., and M.K. Johansen (2000), "Exemplar-based accounts of "multiple-system" phenomena in perceptual categorization," Psychonomic Bulletin & Review 7:375–402.

Nosofsky, R.M., and S.R. Zaki (1998), "Dissociations between categorization and recognition in amnesic and normal individuals: an exemplar-based interpretation," Psychological Science 9:247–255.

Nosofsky, R.M., and S.R. Zaki (1999), "Math modeling, neuropsychology, and category learning: response to B. Knowlton (1999)," Trends in Cognitive Sciences 3:125–126.

Packard, M.G., and J.L. McGaugh (1993), "Double dissociation of fornix and caudate nucleus lesions on acquisition of two water maze tasks: further evidence for multiple memory systems," Behavioral Neuroscience 106:439–446.

Pickering, A.D. (1997) "New approaches to study of amnesic patients: what can a neurofunctional philosophy and neural network methods offer?," Memory 5:255–300.

Poldrack, R.A., and M.G. Packard (2003), "Competition among multiple memory systems: converging evidence from animal and human brain studies," Neuropsychologia 41:245–251.

Reber, P.J., and L.R. Squire (1994), "Parallel brain systems for learning with and without awareness," Learning & Memory 1:217–229.

Reber, P.J., and L.R. Squire (1999), "Intact learning of artificial grammars and intact category learning by patients with Parkinson's disease," Behavioral Neuroscience 113:235–242.

Schacter, D.L. (1987) "Implicit memory: history and current status," Journal of Experimental Psychology: Learning, Memory, and Cognition 13:501–518.

Smith, E.E., A. Patalano, and J. Jonides (1998), "Alternative strategies of categorization," Cognition 65:167–196.

Smith, J.D., and J.P. Minda (2000), "Journey to the center of the category: the dissociation in amnesia between categorization and recognition," Journal of Experimental Psychology: Learning, Memory, and Cognition 27:984–1002.

Squire, L.R. (1992), "Memory and the hippocampus: a synthesis from findings with rats, monkeys and humans," Psychological Review 99:195–231.

Vonsattel, J.D., R.H. Myers, T.J. Stevens, R.J. Ferrante and E.D. Bird (1985), "Neuropathological classification of Huntington's disease," Journal of Neuropathology and Experimental Neurology 44:559–577.

Waldron, E.M., and F.G. Ashby (2001), "The effects of concurrent task interference on category learning," Psychonomic Bulletin & Review 8:168–176.

Willingham, D.B. (1998), "A neuropsychological theory of motor skill learning," Psychological Review 105:558–584.

Willingham, D.B., L.A. Wells, J.M. Farrell and M.E. Stemwedel (2000), "Implicit motor sequence learning is represented in response locations," Memory & Cognition 28:366–375.