Journal of Experimental Psychology:
Learning, Memory, and Cognition
2007, Vol. 33, No. 1, 197–218

0278-7393/07/$12.00 DOI: 10.1037/0278-7393.33.1.197

Within-Category Discontinuity Interacts With Verbal Rule Complexity in Perceptual Category Learning

W. Todd Maddox
University of Texas at Austin

J. Vincent Filoteo
Veterans Affairs San Diego Healthcare System and University of California, San Diego

J. Scott Lauritzen
University of Texas at Austin

A test of the predicted interaction between within-category discontinuity and verbal rule complexity on information-integration and rule-based category learning was conducted. Within-category discontinuity adversely affected information-integration category learning but not rule-based category learning. Model-based analyses suggested that some information-integration participants improved performance by recruiting more "units" in the discontinuous condition. Verbal rule complexity adversely affected rule-based category learning but not information-integration category learning. Model-based analyses suggested that the rule based effect was on both decision criterion learning and variability in decision criterion placement. These results suggest that within-category discontinuity and decision rule complexity differentially impact information-integration and rule-based category learning and provide information regarding the detailed processing characteristics of these two proposed category learning systems.

Keywords: multiple systems, classification learning, category learning, procedural-based learning, rule complexity

The ability to quickly and accurately categorize is fundamental to the survival of all organisms. Incorrect classification of an enemy as friend, or poison as healthful, could be deadly. Recently, there has been an explosion of work that suggests multiple category learning systems exist and that each has a unique neurobiological underpinning. Empirical support comes from a wide range of research areas, including animal learning (McDonald & White, 1993, 1994; Packard & McGaugh, 1992; Poldrack & Packard, 2003), neuropsychology (Filoteo, Maddox, & Davis, 2001a, 2001b; Maddox & Filoteo, 2001, 2005; Myers et al., 2003), functional neuroimaging (Filoteo et al., 2005; Nomura et al., in press; Poldrack, Prabhakaran, Seger, & Gabrieli, 1999; Reber, Stark, & Squire, 1998; Seger & Cincotta, 2002, 2005; E. E. Smith, Patalano, & Jonides, 1998), and cognitive psychology (for reviews, see Keri, 2003; Ashby & Maddox, 2005; Maddox & Ashby, 2004). Collectively this work has identified at least four unique systems. These include a hypothesis-testing system, a procedural-based system, a perceptual representation system, and an exemplar-based system (see Ashby & Maddox, 2005, for a review). Although the neurobiological bases of each system continues to be actively researched, there is good evidence to suggest that the hypothesis testing system relies upon fronto-striatal circuits (Filoteo et al., 2005; E. E. Smith et al., 1998; Ashby, Alfonso-Reese, Turken, & Waldron, 1998), the procedural-based system relies upon the posterior caudate (Ashby et al., 1998; Seger & Cincotta, 2005; Nomura et al., in press), the perceptual representation system relies upon visual cortical areas (Reber, Gitelman, Parrish, & Mesulam, 2003; Reber et al., 1998), and the exemplar-based system potentially relies on medial temporal lobe structures (Nomura et al., in press; see also Poldrack et al., 2001), although the neurobiology of this latter system is less well developed.

This article focuses on the processing characteristics associated with two of these systems: the hypothesis-testing and procedural-based learning systems. These two systems form the cornerstone of a neurobiologically motivated multiple systems model called the *competition between verbal and implicit systems model* (COVIS; Ashby et al., 1998; Ashby & Waldron, 1999). The model postulates that the two systems compete throughout learning and that rule-based category learning is dominated by the hypothesis-testing system, whereas information-integration category learning is dominated by the procedural-based learning system. In rule-based tasks, frequently the rule that maximizes accuracy (i.e., the optimal rule) is easy to describe verbally. For example, Figure 1A presents a scatter plot of stimuli from a rule-based condition with four categories. Each point in the plot denotes the length and orientation of a single line stimulus with different symbols denot-

W. Todd Maddox and J. Scott Lauritzen, Department of Psychology, University of Texas at Austin; J. Vincent Filoteo, Veterans Affairs San Diego Healthcare System and University of California, San Diego.

This research was supported in part by National Institute of Health Grant R01 MH59196 to W. Todd Maddox, National Institute of Health Grant R01 NS41372 to J. Vincent Filoteo, and a McDonnell-Pew Consortium Grant. We thank Shawn Ell, Brad Love, and Art Markman for useful insights on the verbal rule complexity measure, and Russell Poldrack for helpful comments on an earlier version of this article.

Correspondence concerning this article should be addressed to W. Todd Maddox, 1 University Station A8000, Department of Psychology, University of Texas at Austin, Austin, TX 78712. E-mail: maddox@psy.utexas.edu

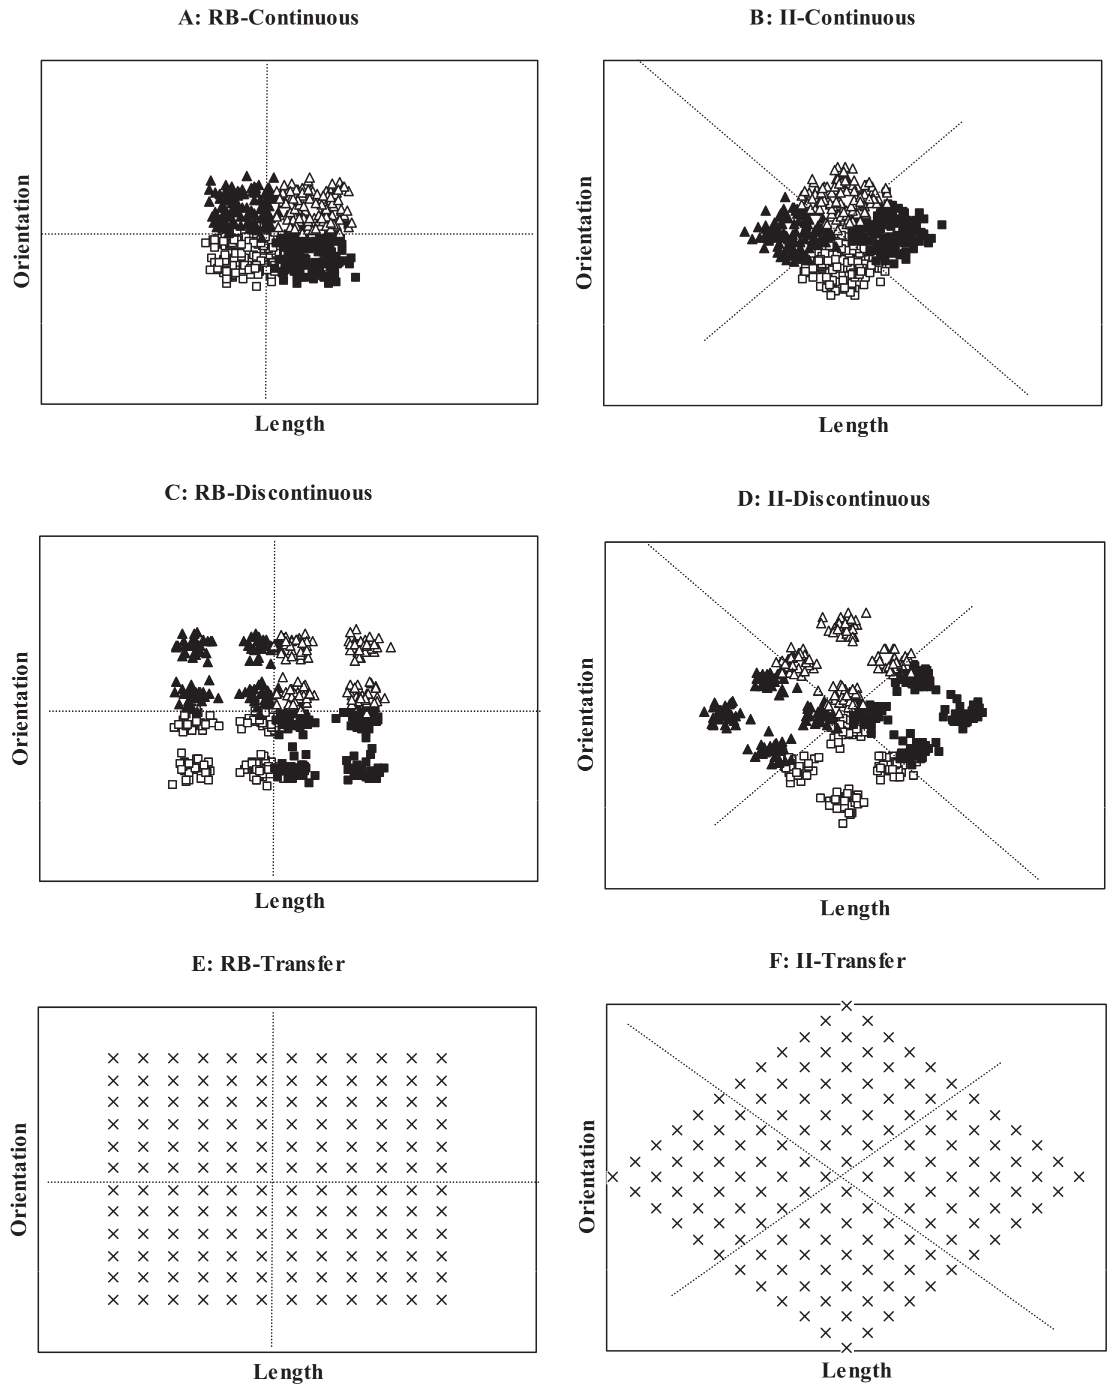

Figure 1. Scatter plots of the stimuli along with the optimal decision bounds from the four conditions of Experiment 1, along with the transfer items used in the rule-based (RB) and information-integration (II) conditions. Open squares denote stimuli from category A: Solid triangles denote stimuli from Category B. Solid squares denote stimuli from Category C. Open triangles denote stimuli from Category D. Plus signs denote transfer items.

ing different categories. In this example, the participant might set a criterion along line length to determine whether the line is "short" or "long," and one along orientation to determine whether the angle is "shallow" or "steep." The decision along each dimension might then be explicitly integrated to determine category membership (e.g., short, shallow angle lines are assigned to Category A; short, steep angle lines are assigned to Category B; long, shallow angle lines are assigned to Category C; and long, steep angle lines are assigned to Category D). This integration is explicit because a decision is first made about the value along each dimension and then that information is integrated to generate a response. In information-integration tasks, the optimal rule is not easy to describe verbally, and the integration of the two dimensions is less likely to be explicit. For example, Figure 1B presents a scatter plot of stimuli from an information-integration condition with four categories. In this case the information from both stimulus components is integrated at some implicit, predecisional stage yielding partitions of the response regions shown in Figure 1B. Because length and orientation are measured in different units, and because the response region partitions are not parallel to the coordinate axes, these response region partitions can not be described verbally.

Nearly all multiple systems models of category learning acknowledge the existence of some sort of hypothesis-testing system. These include models that make no claims about the neurobiological underpinnings, such as rule-plus-exception (RULEX; Nosofsky, Palmeri, & McKinley, 1994) and attention to rules and instances in a unified model (ATRIUM; Erickson & Kruschke, 1998), as well as models that make assumptions regarding the neurobiology such as COVIS (Ashby et al., 1998; see also E. E. Smith et al., 1998). The type of learning that we believe takes place in this system has been studied under different guises for many years, beginning with seminal work by Bruner and Shepard (Bruner, Goodnow, & Austin, 1956; Shepard, Hovland, & Jenkins, 1961) and more recently by Feldman (2001).

The most unique, and perhaps most controversial, aspect of the COVIS model is the proposed existence of a procedural-based category learning system. Of all the systems reviewed above, the procedural system is the more recently proposed and least well understood. As such, the detailed processing characteristics have not been worked out as well as other proposed category learning systems that have been examined over the last twenty years (e.g., hypothesis-testing, exemplar, or prototype systems). To address this shortcoming, we have undertaken a line of research that examines the detailed processing characteristics of this proposed system (reviewed in Maddox & Ashby, 2004). The key structure that has been implicated in procedural-based category learning is the striatum, and for visual stimuli, the key structure within the striatum is the posterior caudate. The initial focus of this structure's involvement in this form of category learning was based on studies that showed striatal involvement in procedural learning (e.g., Jahanshahi, Brown, & Marsden, 1992; Mishkin, Malamut, & Bachevalier, 1984; Willingham, Nissen, & Bullemer, 1989; Willingham, 1998) and visual discrimination learning (Fernandez-Ruiz, Wang, Aigner, & Mishkin, 2001; Malamut, Saunders, & Mishkin, 1984; Teng, Stefanacci, Squire, & Zola, 2000). More direct evidence for the role of this brain region comes from fMRI studies in normal adults that have identified striatal activation during the performance of information-integration tasks (Nomura et al., in press; Seger & Cincotta, 2005).

Although research on the detailed neuroanatomy of the striatum is ongoing, and will likely lead to modifications of any current models of striatal involvement in category learning, several properties of the striatum are generally agreed upon, and these properties make this brain region an excellent candidate for the neurobiological substrate of procedural-based category learning (reviewed in Ashby et al., 1998). Importantly, the architecture of these structures provides constraints on the potential processing characteristics of procedural-based category learning. First, the striatum is the primary input structure within the basal ganglia that receives projections from virtually all areas of the neocortex including extrastriate visual cortical areas, such as the inferotemporal lobes (Heimer, 1995; Saint-Cyr, Ungerleider, & Desimone, 1990; Van Hoesen, Yeterian, & Lavizzo-Mourey, 1981; Webster, Bachevalier, & Ungerleider, 1993). Second, within the posterior caudate, there is good evidence that these projections are convergent in the sense that many visual cortical afferents converge on relatively few striatal cells. Current estimates suggest a convergence ratio of approximately 10,000 to 1 (e.g., Wilson, 1995). Because of the many-to-one convergence of visual cortical cells onto cells in the posterior caudate, it is likely that the visual perceptual space is represented in some form in the posterior caudate. In addition to the convergent nature of the cortico-striatal projections, Wilson (1995) found that only about 5% of cells in the caudate nucleus are interneurons and that they are sparsely distributed, suggesting little interconnectivity across the caudate. This latter property would make it difficult for neurons within the caudate to communicate with one another, and one potential consequence of this would be less generalization of any learning that takes place in this region. Third, the striatum is characterized by a high degree of synaptic plasticity that is mediated by dopamine release from the substantia nigra pars compacta (Schultz, Apicella, & Ljungbert, 1993; J. Wickens, 1993), and there is considerable evidence that this dopamine signal is involved in reinforcement-based learning (Beninger, 1983; Miller, Sanghera, & German, 1981; Montague, Dayan, & Sejnowski, 1996; J. Wickens, 1993). In fact, recent studies have implicated the substantia nigra in some forms of human category learning, suggesting that dopamine likely plays a role in this process (Aron et al., 2004). Finally, support for a role of the striatum in stimulus-response learning in categorization comes from single-cell recording studies in monkeys (Merchant, Zainos, Hernandez, Salinas, & Romo, 1997; Romo, Merchant, Ruiz, Crespo, & Zainos, 1995; Romo, Merchant, Zainos, & Hernandez, 1997). In addition, the fact that the basal ganglia project to prefrontal cortex and motor output areas via connections through the thalamus (Alexander, DeLong, & Strick, 1986) further supports the notion that the striatum might be involved in linking stimuli with responses. From these basic tenets, Ashby et al. (1998) outlined the framework for COVIS' procedural-based category learning system. As currently formulated, the model assumes that, through a procedural learning process, cells in the posterior caudate begin to associate a large group of visual cortical cells (i.e., all that project to it) with a motor response (via the caudate projections through the globus pallidus to the thalamus and ultimately to premotor areas; Alexander et al., 1986).

To date, a number of behavioral experiments have been conducted using healthy young adults as participants that dissociate processing in the hypothesis-testing and procedural-based learning systems (see Ashby & Maddox, 2005; Maddox & Ashby, 2004). Behavioral manipulations designed to affect processing in the hypothesis-testing system include manipulations that test the proposed role of the frontal cortex and anterior caudate in this system. In particular, these studies manipulated attention and working memory. For example, Waldron and Ashby (2001; see also Zeithamova & Maddox, 2006) 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, Ashby, Ing, and Pickering (2004) showed that rule-based category learning was disrupted by a sequential memory scanning task whereas information-integration category learning was not. Behavioral manipulations designed to affect processing in the procedural-based learning system have primarily focused on providing evidence that dopamine and the striatum are involved this system. For example, manipulations of the nature and timing of the feedback impact information-integration category learning more than rule-based category learning (Maddox, Ashby, & Bohil, 2003; Maddox & Ing, 2005), supporting the observation that the dopamine learning signal is time dependent. In addition, manipulations of the linkage between the category of stimuli and the response location have also been examined (Ashby, Ell, & Waldron, 2003; Maddox, Bohil, & Ing, 2004; however, see Nosofsky, Stanton, & Zaki, 2005). This latter work follows directly from findings in the procedural learning literature that suggest that changing the location of the response keys interferes with learning on the serial reaction time (SRT) task even when the sequence of stimulus positions was unchanged (Willingham, Wells, Farrell, & Stemwedel, 2000). Taken together, the results of these past studies are highly consistent with the possibility that dopamine and the striatum play an important role in procedural-based category learning, and provide evidence that the mechanisms involved in this form of learning are generally distinct from those involved in hypothesis-testing category learning.

Additional evidence that the procedural-based system is distinct from the hypothesis-testing system has come from the application of computational models to participants' data in these past studies (see Maddox & Ashby, 2004, for a review). Each model instantiates the assumptions of the hypothesis-testing or procedural-based learning system. The goals of the model-based analysis are twofold. First, and foremost, the models are used to determine which system (at the computational level) best describes the pattern of responses for a particular participant. This is important because it is critical to determine how many and which participants are using the strategy that is predicted to mediate learning of a particular category structure (i.e., a procedural-based learning system for information-integration categories and a hypothesis-testing system for rule-based categories). Because the experimental manipulations are chosen to affect processing in one system, but not the other, we can examine accuracy rates for participants best fit by each model class and determine whether the manipulation affected only those participants using the relevant strategy. Although a detailed discussion of this issue is beyond the scope of this article, it is important to say there are many examples in the literature in which researchers make the implicit assumption that all participants are using a particular strategy, when follow-up research suggests a wide range of strategies are used (see Ashby & Maddox, 2005, for a discussion of this issue with regard to other categorization tasks). A lack of understanding of the specific strategy used often leads to discrepancies in the literature. Second, each model has a small set of parameters whose values are estimated from the data. An examination of these parameter values often provides insights onto the processing locus of some accuracy deficit (Maddox & Ashby, 2004).

Overview of the Current Studies

Experiment 1 examines processing in the procedural-based learning system. As outlined above, it is well established that there is a many-to-one mapping of inferotemporal cortex cells onto the posterior caudate nucleus (Wilson, 1995). It is also generally agreed upon that proximal cells within the inferotemporal cortex converge onto the same cells within the caudate and are tuned to perceptually similar stimuli (Cheng, Saleem, & Tanaka, 1997; Fujita, Tanaka, Ito, & Cheng, 1992; Tanaka, 2000). Because of the funnel-like connectivity between inferotemporal cortex and posterior caudate, we propose that a low-resolution map of the perceptual space is represented by the cells of the posterior caudate (however, see Bar-Gad, Morris, & Bergman, 2003). In short, perceptually similar stimuli should be more likely to stimulate the same caudate cell. It is also well established that the caudate nucleus has relatively few interneurons (approximately 5%) and that these are sparsely distributed (Wilson, 1995). Thus, there is very little interconnectivity and very little information transfer across cells in the caudate. If we make the reasonable assumption that category learning is easier when each category can be represented by few rather than many cells, then it follows that the procedural-based learning system should be more efficient when stimuli coming from the same category are perceptually similar and form a coherent (or continuous) group and should be less efficient when stimuli coming from the same category are perceptually dissimilar and form a less coherent (or discontinuous) group. Consequently, any experimental manipulation that increases within-category discontinuity by having categories composed of distinct subgroups of stimuli should adversely affect learning, whereas any manipulation that increases within-category coherence should improve learning.

Maddox, Filoteo, Lauritzen, Connally, and Hejl (2005) provided an initial test of this hypothesis and found that discontinuous categories did in fact adversely affect information-integration but not rule-based category learning. However, in that study the rule-based category learning task was unidimensional, whereas both dimensions were relevant for the information-integration task. It is reasonable to assume that, in general, learning a two-dimensional rule is more difficult than learning a one-dimensional rule and that any experimental manipulation that increases the difficulty of the task (such as increasing within-category discontinuity) would, by definition, affect the more complex task. Thus the effect of within-category discontinuity observed in Maddox et al. (2005) might not be due to differential processing characteristics of the two systems but rather to a simple difference in task difficulty. Some researchers have argued that task difficulty underlies many of the empirical dissociations attributed to multiple systems (e.g., Nosofsky et al., 2005). Experiment 1 of the current article examines a situation in which both the rule-based and the information-integration category learning tasks involve two-dimensional decision rules and provides a direct test of the alternative "difficulty" hypothesis.

To elaborate, consider the scatter plots of the stimuli for two information-integration conditions displayed in Figures 1B and 1D. Notice that each of the categories in Figure 1D is composed of stimuli from four separate clusters each of which is not contiguous with any of the other clusters. On the other hand, note that each of the categories in Figure 1B is also composed of stimuli from four separate clusters, but in this case the clusters are contiguous. Because the within-category discontinuity is much larger in Figure

1D than in Figure 1B, the procedural-based learning system should perform more poorly in the former case. Unlike the procedural-based learning system, the efficiency of the hypothesis-testing system should not be dependent on within-category continuity. Rather, as long as the decision rule remains the same across within-category continuity conditions, there should be little effect on rule-based category learning (Maddox et al., 2005). For example, consider the rule-based conditions displayed in Figures 1A and 1C. Despite the increased discontinuity in Figure 1C relative to Figure 1A, the decision rule is identical in both cases, and thus a hypothesis-testing system, whose aim is to identify the correct decision rule, should perform equally in the two cases. It is worth mentioning that the stimuli in Figures 1A and 1C were constructed directly from the Figure 1B and 1D stimuli, respectively, by applying a simple linear transformation. This approach is highly advantageous because it allows us to compare rule-based and information-integration category learning using structurally equivalent categories—that is, category structures for which the optimal accuracy, number of stimulus clusters, within-cluster scatter, and cluster coherence is equivalent across information-integration and rule-based conditions

It is important to point out that the working memory demand associated with learning a simple unidimensional rule is less than that for a two-dimensional conjunctive rule. Thus, it is possible that the effect of within-category discontinuity that was previously observed in Maddox et al. (2005) might only emerge when the working memory demand is higher. In the present study, the rule based category learning task in Experiment 1 is conjunctive and, thus, is more complex and has a greater working memory demand than that associated with a unidimensional rule. The possibility that increased working memory load might lead to a rule-based performance decrement under within-category discontinuity conditions is not without precedence. In two recent studies (Ell, Ing, & Maddox, 2006; Maddox, Lauritzen, & Ing, 2006), we found that experimental manipulations thought to affect only information-integration category learning did lead to rule-based category learning deficits when the situation was one with a high working memory demand.

Experiment 2 examines processing in the hypothesis-testing system, in particular that the efficiency of hypothesis-testing systems is dependent on the complexity of the verbal decision rule (Feldman, 2001; Salatas & Bourne, 1974; Shepard et al., 1961). It is a decision rule that is learned by the hypothesis-testing system. This decision rule is not derived from the specific values of the stimulus attributes or their similarity, but, rather, the rule that is learned is more abstract. Because rule-based category learning is thought to be mediated by a hypothesis-testing system, any experimental manipulation that increases the complexity of the decision rule should adversely affect rule-based category learning by increasing the load on working memory and attention. On the other hand, the procedural-based learning system is not assumed to learn "rules" in the way that the hypothesis-testing system does. Rather the procedural-based learning system is assumed to learn to assign groups of stimuli to responses. As long as the nature of the stimulus space remains relatively unchanged across decision rule complexity conditions (e.g., no discontinuity is introduced), the efficiency of the procedural-based learning system should be unaffected, and, by extension, information-integration category learning should not be impacted to a large extent.

To test this prediction we constructed a pair of rule-based tasks that differed in verbal rule complexity (Figures 2A and 2C). We then applied a simple linear transformation to each of these category structures to construct a pair of structurally equivalent information-integration categories (Figures 2B and 2D). In selecting the rule-based tasks, it was critical that the difference in verbal rule complexity be obvious and have face validity. In light of this fact, we selected as our simple task a two-category, unidimensional rule-based task in which the participant should set a criterion on length, responding *A* to short lines and *B* to long lines, while ignoring orientation (Figure 2A). We selected as our complex task, a two-category disjunctive rule-based task in which the participant should set one criterion on length and a second criterion on orientation, responding *A* if the line is short and shallow or long and steep and *B* if the line is short and steep or long and shallow. The Figure 2C rule-based task is classically referred to as an *exclusive-or (XOR) type decision rule*. For binary-valued dimensions, it is well established that XOR problems are some of the most complex (Salatas & Bourne, 1974; Shepard et al., 1961) and are more complex than a simple unidimensional problem.

In summary, we predicted an interaction between the effects of within-category discontinuity and the complexity of the verbal decision rule on processing efficiency in the procedural-based learning and hypothesis-testing systems and, by extension, on information-integration and rule-based category learning. In particular, within-category discontinuity should adversely affect information-integration category learning but not rule-based category learning as long as the verbal decision rule remains unchanged across conditions. On the other hand, an increase in the complexity of the rule-based decision rule should adversely affect rule-based category learning but not information-integration category learning. We test these predictions in the following two experiments.

Experiment 1: Within-Category Discontinuity Experiment

Experiment 1 tests the hypothesis that within-category discontinuity will adversely affect information-integration but not rule-based category learning, when both the information-integration and rule-based category structures are two-dimensional. Two aspects of this study provide important extensions of Maddox et al. (2005). First, as outlined above, both the rule-based and information-integration category structures are two-dimensional, whereas in Maddox et al. (2005) the rule-based structure was one-dimensional. By using two-dimensional structures, we alleviate possible concerns that differences in complexity, and not differences in processing characteristics, might explain the pattern of results. In addition, the working memory demand associated with learning a simple unidimensional rule is less than that for a two-dimensional conjunctive rule. Thus, it is possible that the effects of within-category discontinuity might emerge only when the working memory demand is higher. Second, extensive model-based analyses are conducted on the present data to provide insights into the types of response strategies used by participants. Maddox et al. (2005) showed only that, in general, procedural-based learning models fit information-integration category learning data best, whereas hypothesis-testing models fit rule-based category learning data best. In Experiment 1, we offer a detailed

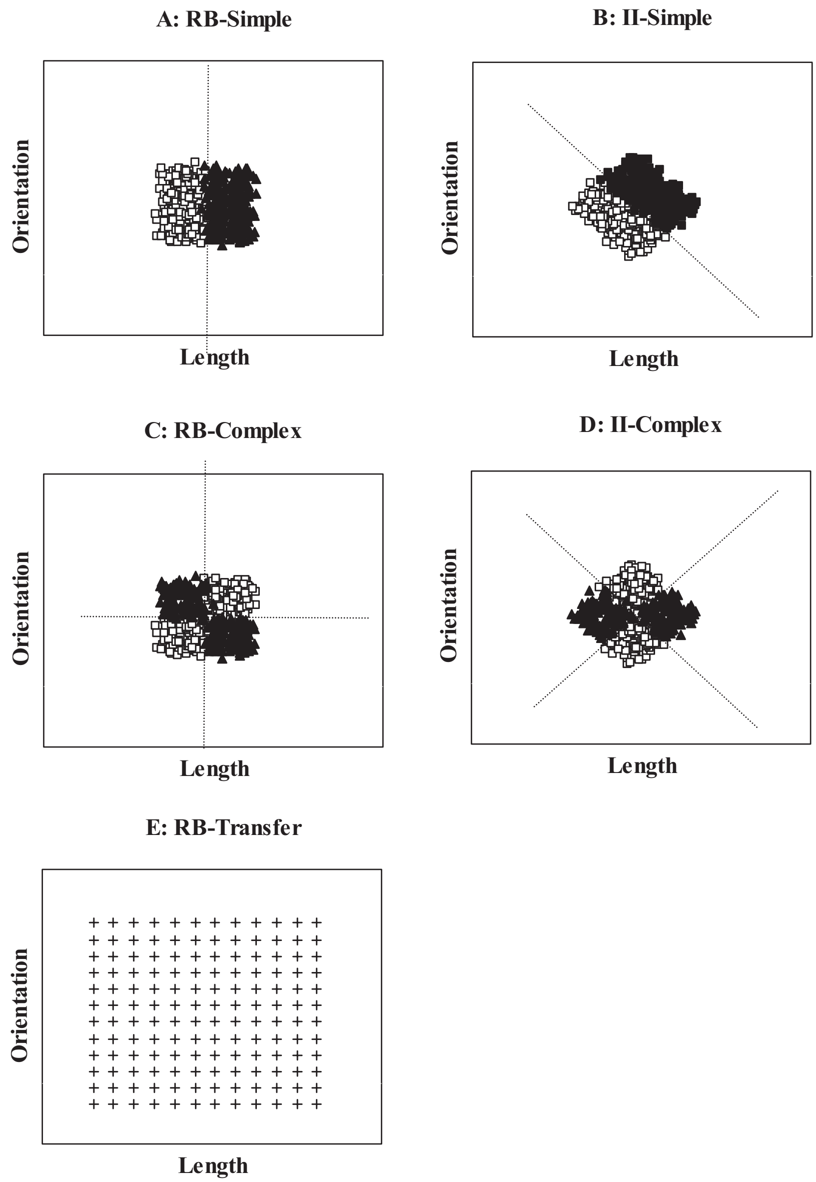

Figure 2. Scatter plots of the stimuli with the optimal decision bounds from the four conditions of Experiment 2 and with the transfer items used in the rule-based and information-integration conditions. Open squares denote stimuli from Category A. Solid triangles denote stimuli from Category B. Plus signs denote transfer items.

examination of the model parameters and, more importantly, apply a large set of procedural-based learning models to the data. We anticipated that these analyses would turn out to be quite informative and have strong implications for the detailed processing characteristics of the procedural-based learning system.

In Experiment 1, we constructed information-integration and rule-based categories for which four clusters of stimuli were associated with each category. In the continuous conditions (hereafter referred to as *RB-continuous* and *II-continuous* conditions) the four clusters formed a coherent group. In the discontinuous conditions (hereafter referred to as the *RB-discontinuous* and *II-discontinuous* conditions) the four clusters were discontinuous (i.e., perceptually distinct clusters). Scatter plots of the stimuli and optimal bounds are displayed in Figure 1A–1D. It is important to note that the optimal decision bound is identical across continuous and discontinuous conditions, and the rule-based and information-integration stimuli are derived via a simple linear transformation yielding structurally equivalent categories.

As outlined above, the funnel-like nature of cortico-caudate connections and the fact that there are few interneurons in the caudate leads to the prediction that information-integration category learning will be adversely affected by discontinuous clusters because the subclusters within a category are distinct. If the hypothesis-testing system learns an abstract rule, and that rule is identical across continuous and discontinuous conditions, then rule-based category learning should be unaffected by discontinuous clusters. Category learning was examined across five, 96-trial blocks. At the end of the experimental session, we also included a set of 144 transfer trials in which no corrective feedback was provided. The transfer stimuli (see Figure 1E and 1F) included stimuli from the portion of the length-orientation space presented during training as well as items from novel (untrained) portions of the space. The transfer trials were included to examine performance generalization to portions of the stimulus space that were not encountered during training. If an abstract rule is learned by the hypothesis-testing system, then learning should generalize fairly well to novel (untrained) items in the rule-based condition (Maddox et al., 2005). In fact, untrained items that are better described by the rule (i.e., stimuli that are further from the optimal bound) should yield superior performance. On the other hand, if the procedural-based system is constrained by the neurobiology of the striatum (i.e., the many-to-one convergence and the lack of interneurons), then generalization should not be as good in the information-integration condition and should be worse for transfer items coming from regions of the stimulus space that are farther from the training items, even if they are well described by the optimal rule.

Method

Participants

Eighteen participants (9 women and 9 men) were solicited from the University of Texas community and received $30 for participating in this study. Each participant completed all four experimental conditions with the condition order being determined from a Latin square. Only one of the four conditions (approximately 60 min) was completed during a single test day, and 1 rest day was required between testing sessions. In many cases, there were 2 or more days between sessions. Visual acuity was tested in each participant, and all participants had either 20/20 vision or vision corrected to at least 20/20. One participant perseverated in one condition and responded randomly in another condition. The data from this participant were excluded from all subsequent analyses.

Stimuli and Stimulus Generation

The experiment used the randomization technique introduced by Ashby and Gott (1988). The category structures are displayed in Figure 1 along with the optimal decision bounds. The category distribution parameters are outlined in Table 1, and optimal accuracy was approximately 95%. The stimuli were computer generated lines displayed on a 21-in. monitor with

Table 1
Category Distribution Parameters From the Experiment 1 Within-Category Discontinuity Experiment

Category	μ_l	μ_o	σ_l	σ_o	cov_{lo}	Category	μ_l	μ_o	σ_l	σ_o	cov_{lo}
Rule-Based Continuous						Information-Integration Continuous					
A_1	189	75	10	10	0	A_1	321	−27	10	10	0
A_2	189	109	10	10	0	A_2	297	−3	10	10	0
A_3	223	75	10	10	0	A_3	345	−3	10	10	0
A_4	223	109	10	10	0	A_4	321	21	10	10	0
B_1	189	143	10	10	0	B_1	273	21	10	10	0
B_2	189	177	10	10	0	B_2	248	45	10	10	0
B_3	223	143	10	10	0	B_3	297	45	10	10	0
B_4	223	177	10	10	0	B_4	273	69	10	10	0
C_1	257	75	10	10	0	C_1	369	21	10	10	0
C_2	257	109	10	10	0	C_2	345	45	10	10	0
C_3	291	75	10	10	0	C_3	393	45	10	10	0
C_4	291	109	10	10	0	C_4	369	69	10	10	0
D_1	257	143	10	10	0	D_1	321	69	10	10	0
D_2	257	177	10	10	0	D_2	297	93	10	10	0
D_3	291	143	10	10	0	D_3	345	93	10	10	0
D_4	291	177	10	10	0	D_4	321	118	10	10	0
Rule-Based Discontinuous						Information-Integration Discontinuous					
A_1	155	41	10	10	0	A_1	321	−75	10	10	0
A_2	155	109	10	10	0	A_2	273	−27	10	10	0
A_3	223	41	10	10	0	A_3	369	−27	10	10	0
A_4	223	109	10	10	0	A_4	321	21	10	10	0
B_1	155	143	10	10	0	B_1	248	−3	10	10	0
B_2	155	211	10	10	0	B_2	200	45	10	10	0
B_3	223	143	10	10	0	B_3	297	45	10	10	0
B_4	223	211	10	10	0	B_4	248	93	10	10	0
C_1	257	41	10	10	0	C_1	393	−3	10	10	0
C_2	257	109	10	10	0	C_2	345	45	10	10	0
C_3	325	41	10	10	0	C_3	441	45	10	10	0
C_4	325	109	10	10	0	C_4	393	93	10	10	0
D_1	257	143	10	10	0	D_1	321	69	10	10	0
D_2	257	211	10	10	0	D_2	273	118	10	10	0
D_3	325	143	10	10	0	D_3	369	118	10	10	0
D_4	325	211	10	10	0	D_4	321	166	10	10	0

Note. The subscripts denote cluster distributions for each category. Optimal accuracy was held constant at 95% in all conditions.

1360 × 1024 resolution. Each line was presented in white on a black background. The orientation values in Table 1 were transformed to radians by multiplying each value by π/500. A total of 30 stimuli were sampled randomly from each of the four clusters for each of the four categories for a total of 480 stimuli. The resulting 480 stimuli were randomized and divided into five, 96-trial blocks. These were presented during categorization training. One hundred forty-four unique stimuli were used during the transfer phase (see Figures 1E and 1F).

Procedure

Each participant performed individually in a dimly lit testing room with an approximate viewing distance of 35 cm. The participants were informed that there were four equally likely categories. They were informed that perfect performance was impossible but that high levels of accuracy could be achieved. They were instructed to learn about the categories, to be as accurate as possible, and not to worry about speed of responding. On each trial, the stimulus appeared and remained on the screen until the participant generated a response by pressing one of four keys. The correct category label was then presented on the screen for 1 s along with the word *wrong* if their response was incorrect or the word *right* if their response was correct. Once feedback was given, the next trial was initiated. The procedure for the transfer trials was identical except that feedback was omitted.

Results

Training Block Analyses

A Category Structure (rule-based vs. information-integration) × Within-Category Discontinuity (continuous vs. discontinuous) × Block (five 96-trial blocks) within-participant analysis of variance (ANOVA) was conducted on the training block accuracy rates. The accuracy rates averaged across participants are presented in Figure 3. The main effects of within-category discontinuity, $F(1, 16) = 16.80$, $p < .001$, and block, $F(4, 64) = 15.21$, $p < .001$, were significant, and the main effect of category structure was not, $F(1, 16) = 1.61$, $p = .22$. It is important to note that the Category Structure × Within-Category Discontinuity interaction was significant, $F(4, 64) = 4.85$, $p < .05$, whereas the other two-way and the three-way interactions were nonsignificant. (To determine whether the Latin square condition orderings affected performance, we replicated the analysis with order included. It is important to note that there was no Category Structure × Within-Category Discontinuity × Order interaction, [$p = .33$].) To determine the locus of the Category Structure × Within-Category Discontinuity interaction, we examined the effects of within-category discontinuity

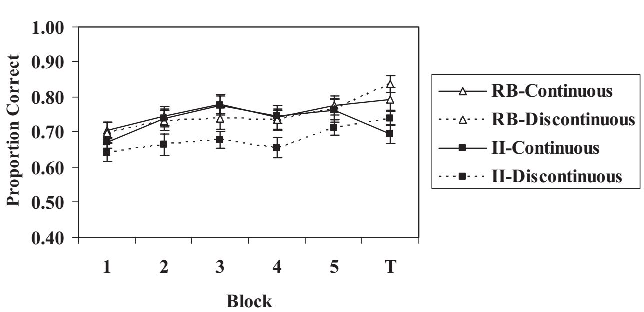

Figure 3. Proportion correct (averaged across participants) from Experiment 1 along with standard error bars. Blocks 1–5 are the training blocks and Block T is the transfer block.

separately for the rule-based and information-integration conditions. As predicted, information-integration performance was significantly worse in the discontinuous condition (.67) than in the continuous condition (.74), $t(16) = 3.83$, $p < .001$, whereas within-category discontinuity had no effect on rule-based category learning (discontinuous = .73; continuous = .75; $t[16] = 1.18$, $p = .26$). These findings suggest that discontinuous categories did affect information-integration category learning but not rule-based category learning when both category structures required the participant to process two stimulus dimensions. Notice also that performance in the two continuous conditions was approximately equal, providing evidence against any simple "difficulty" explanation of the results.

We also examined the accuracy rates separately for each category cluster averaged over blocks. These results mirrored those from the overall training analyses. One result worth highlighting is that for the four clusters closest to the center of the space. These four clusters of items were presented in the continuous and discontinuous conditions. It is interesting to note that the accuracy for these clusters was lower in the information-integration discontinuous condition (.47) than in the information-integration continuous condition (.60). In the rule-based conditions, on the other hand, accuracy rates were much more similar (discontinuous = .60, continuous = .63). This pattern of findings is predicted from the two proposed underlying systems. In the information-integration continuous condition, each of the four central clusters is perceptually similar to many items from its own category; whereas in the discontinuous condition, each of the four central clusters is not perceptually similar to many items from its own category. The more perceptually similar items are from the same category, the more potential overlap in their projections to the caudate and the better that their category assignment is learned by the procedural-based categorization system. On the other hand, the hypothesis-testing system learns an explicit rule. Thus, perceptual similarity among category members has less effect on performance.

Transfer Block Analyses

We made two a priori predictions regarding the transfer data. First, we predicted that the increase in performance from the final training block to the transfer block would be greater for the rule-based than for the information-integration category structures. The transfer phase included a large number of stimuli far from the decision bound (see Figures 1E and 1F), and thus a large number of stimuli that are excellent examples of the rule. If it is the decision rule that is represented by the hypothesis-testing system, then transfer performance should be good because of the large number of items far from the bound. On the other hand, if procedural-based category learning is impacted by perceptual similarity and constrained by the known neurobiology, then generalization should not be as good in the information-integration conditions. As predicted, the performance increase was larger in the rule-based conditions (average increase of .045) than in the information-integration conditions (average decrease of .02), $t(16) = 3.13$, $p < .01$.

Second, we predicted a bigger increase in transfer performance in the information-integration discontinuous condition than in the information-integration continuous condition. This follows because a wider range of items was trained in the discontinuous condition than in the continuous condition, which should lead to a wider range of stimulus values being linked to each categorization response. As predicted, a .03 performance improvement was observed in the discontinuous condition transfer block relative to the final training block, but a .07 performance decrement was observed in the continuous condition. This within-category discontinuity effect was statistically significant, $t(16) = 3.28$, $p < .01$. In addition, we predicted no transfer performance differences across within-category discontinuity conditions with the rule-based category structures. The performance improvement from the final training to the transfer block was larger in the rule-based discontinuous condition (.07) than in the continuous condition (.02), and this difference was marginally significant, $t(16) = 1.78$, $p = .094$, arguing against this prediction. These analyses provide support for the hypothesis that the generalization profile in the procedural-based learning system should be affected by within-category discontinuity, but they also suggest that within-category discontinuity had some effect on generalization in the hypothesis-testing system.

Model-Based Analyses

The accuracy-based analyses suggest that within-category discontinuity adversely affects information-integration but not rule-based category learning. In this section, we attempt to gain a more detailed understanding of the locus of the information-integration deficit by applying a series of computational models to the data. These models allow us to confirm or disconfirm that participants are using the appropriate strategy to solve the task and allow us to examine the resulting parameter values in an attempt to provide insight into the possible locus of accuracy deficits. Because of concerns with modeling aggregate data (e.g., Ashby, Maddox, & Lee, 1994; Estes, 1956; Maddox, 1999; Maddox & Ashby, 1998; Smith & Minda, 1998), two classes of models were fit on a block-by-block basis separately to the data from each participant. One class of models is compatible with the assumption that participants used an explicit hypothesis-testing strategy, and one class is consistent with the assumption that participants used an implicit procedural-based learning strategy—instantiated by applying Ashby and Waldron's (1999) striatal pattern classifier (see below for details). We modeled the data from all five training blocks, but for brevity we focus our discussion on fits of the models to the final block of trials.

The model parameters were estimated by using maximum likelihood (Ashby, 1992; T. D. Wickens, 1982) and the goodness-of-fit statistic,

$$\mathrm{AIC} = 2r - 2\ln L,$$

where r is the number of free parameters and L is the likelihood of the model given the data (Akaike, 1974; Takane & Shibayama, 1992). The Akaike information criterion (AIC) statistic penalizes a model for extra free parameters in such a way that the smaller the AIC, the closer a model is to the "true model," regardless of the number of free parameters. Thus, to find the best model among a given set of competitors, one simply computes an AIC value for each model and chooses the model associated with the smallest AIC value (for a discussion of the complexities of model comparisons, see Myung, 2000; Pitt, Myung, & Zhang, 2002). Because within-category discontinuity affected information-integration but not rule-based category learning, we focus our model-based analyses on the information-integration conditions. For completeness, though, we briefly address the rule-based conditions as well.

Models Fit to the Information-Integration Condition Data

Procedural-based learning models. To examine the effects of within-category discontinuity on the strategies used by participants to solve the information-integration task, procedural-based learning models were applied. Specifically, we applied several versions of Ashby and Waldron's (1999) *striatal pattern classifier* (SPC). This model has been found to provide a good computational model of participants' responding in previous information-integration category learning studies (e.g., Ashby & Waldron, 1999; Waldron & Ashby, 2001; for applications to stimulus identification, see Ashby, Waldron, Lee, & Berkman, 2001; Maddox, 2001, 2002). In addition, the assumptions of this model are based on the neurobiology proposed to underlie the procedural-based system (Ashby et al., 2001). The model is outlined in detail in Ashby and Waldron but is briefly reviewed here. As outlined above, there is a many-to-one mapping of cortical cells onto cells in the striatum (Wilson, 1995). We assume that a low-resolution map of the perceptual space is represented among the striatal units. As learning progresses, the striatal units become associated with a category label. Thus, the striatum learns to associate a response with groups of cells in visual regions of cortex. The SPC offers a simple model of the procedural-based perceptual classification learning system that takes into account this architecture. It is important to be clear that the SPC is a computational model that is inspired by what is known about the neurobiology of the striatum. Because of this fact, the "striatal units" are hypothetical and could be interpreted within the language of some other computational model (e.g., as "prototypes" in a multiple prototype model). Even so, we take the liberty of using the term *striatal unit* but readily acknowledge that to date there is no neurobiological evidence for the actual existence of such units.

Four versions of the striatal pattern classifier (SPC) were applied: SPC-1, SPC-2, SPC-4, and the optimal model. Each of these models varied in terms of the number of "units" that were used to represent the four categories. It is worth mentioning that, at this point, the models make no claim about how the striatal units become associated with each category or how the number of units is selected. The SPC-1 assumes that there is one striatal-category "unit" in the length-orientation space for each category, yielding a total of four striatal units. On each trial, the participant determines which unit is closest to the perceptual effect and gives the associated response. Because the location (in length-orientation space) of one of the units can be fixed, and because a uniform expansion or contraction of the space will not affect the location of the resulting response region partitions, the SPC-1 contains six free parameters—five that determine the location of the units and one that represents the noise associated with the placement of the striatal units. The noise parameter estimates the variability associated with the participant's responding, with large variability estimates being associated with less deterministic responding and small variability estimates being associated with more deterministic responding. The SPC-2 and SPC-4 models are the same as the SPC-1 model, except that the SPC-2 model assumes that each category has two striatal-category units (13 parameters total), and the SPC-4 model assumes that each category has four striatal-category units (29 parameters total). The optimal SPC-1 model assumes optimal placement of the striatal units and was also applied to the data. This model contains only the single noise parameter.

These models were developed to examine the possibility that participants in the discontinuous spread condition might learn to associate the separate, and distinct, subclusters of perceptually similar stimuli with the appropriate category. Recall that in the discontinuous spread condition (see Figure 1D), each category is composed of four distinct subclusters of stimuli. The stimuli within each subcluster are perceptually similar and might be associated with a single striatal unit, but stimuli across subclusters are distinct and thus might be associated with a different striatal unit. The SPC-4 model was developed to determine whether participants required a single unit within each category to represent the different clusters. The SPC-2 model, on the other hand, was developed to determine whether some participants might use more than one striatal unit but not necessarily one for each of the four distinct subclusters. All four models were applied to both the continuous and discontinuous spread condition data. By determining whether the SPC-1, SPC-2, SPC-4, or the optimal model provides the best account of the data, we should gain some insight into the effects of within-category discontinuity on information-integration category learning. Specifically, if the SPC-4 model provides a better accounting of the data in the discontinuous spread condition, then we can speculate that the manipulation is resulting in participants having to alter the manner in which the categories (and clusters) are represented cognitively. This finding, in turn, could also have implications for the biological representation of categories within the procedural-based system (see below).

Hypothesis-testing models. Three models were compatible with the assumption that participants used a hypothesis-testing strategy. All assume a conjunctive rule. The conjunctive(1) model assumes that the participant makes one decision about the length of the line (short or long), a separate decision about the orientation of the line (shallow or steep), and then integrated this information postdecisionally (i.e., after deciding whether the line is short or long and after deciding whether the orientation is shallow or steep). The model assumes that the participant used the following decision rule: Respond *A* if the line length is short and the orientation is shallow, respond *B* if the line length is short and the orientation is steep, respond *C* if the line length is long and the orientation is shallow, and respond *D* if the line length is long and the orientation is steep. This model has three free parameters: two decision criteria parameters and a noise parameter that estimates the variability associated with the participant's trial by trial memory and application of the decision criteria. The conjunctive(2) model instantiates an "extreme values" type of decision rule. This model assumes that the participant sets two criteria along the length dimension that partitions the length dimension into three regions. The model assumes that the participant sets a criterion along the orientation dimension that is invoked only when the perceived length falls into the intermediate length region. The model assumes that the participant used the following rule: Respond *B* if the length is short, respond *C* if the length is long, if the length is intermediate then respond *A* if the orientation is shallow, and respond *D* if the orientation is steep. The conjunctive(3) model is similar. This model assumes that the participant sets two criteria along the orientation dimension that partitions the orientation dimension into three regions. The model assumes that the participant sets a criterion along the length dimension that is invoked only when the perceived orientation falls into the intermediate orientation region. The model assumes that the participant used the following rule: Respond *A* if the orientation is shallow, respond *D* if the orientation is steep, if the orientation is intermediate then respond *B* if the line is short, and respond *C* if the line is long. The conjunctive(2) and the conjunctive(3) models contain four parameters: three criteria and the noise parameter.

Model-Based Results

The number of participants in each condition whose final training block of data was best fit by each model along with the accuracy rate achieved by those participants is presented in Table 2.[1] Because a small number of participants' data was best fit by one of the three hypothesis-testing models, these participants were aggregated in Table 2. The model-based analyses revealed that only 2 of the 17 participant's data in the information-integration continuous condition and 3 of 17 participant's data in the information-integration discontinuous condition were best fit by a model that instantiates a hypothesis-testing strategy. The remaining participants' data were best fit by a procedural-based learning model (i.e., a version of the SPC). It is interesting to note that whereas the accuracy rates for the participants whose data was best fit by a hypothesis-testing model were equal across the continuous (.64) and discontinuous (.62) conditions ($t < 1$), performance for the participants whose data was best fit by some version of the SPC was marginally superior in the continuous (.79) condition relative to the discontinuous (.73) condition, $t(27) = 1.84$, $p = .077$. Thus, the overall performance decrement observed in the discontinuous condition was due only to performance for the participants whose data was best fit by a procedural-based learning model.

Next we turn to an examination of the frequency counts and accuracy rates for participants best fit by the three versions of the SPC. Several comments are in order. First, a large number of participants in the continuous ($n = 9$) and discontinuous ($n = 7$) conditions were best fit by the SPC-1 model. In addition, accuracy was higher for the SPC-1 participants in the continuous condition (.75) than in the discontinuous condition (.69), although this difference was not significant, $t(14) = 1.41$, $p = .18$. Second, the SPC-2 model provided the best account of the data from 5 participants in the continuous condition but only 3 in the discontinuous condition. Again, the accuracy was higher for SPC-2 participants in the continuous condition (.83) than for those in the discontinuous condition (.73), but, despite the small sample size, this difference was significant, $t(6) = 3.11$, $p < .05$. Third, the SPC-4 model

[1] It is important to note that AIC is not the only goodness-of-fit measure available. For nested models, likelihood ratio tests can be used (Ashby, 1992). Measures similar to AIC, such as the BIC measure, are available as well (Myung, 2000). We conducted likelihood ratio tests when the models were nested, and the results converged nicely with those in Table 2. The BIC measure tends to be more conservative, with a bias toward simpler models. As expected, the BIC measure shifted the distribution of frequency counts in Table 2 toward the simpler models. Even so, the overall pattern of results was similar. The same approach was taken in Experiment 2 and yielded similar results.

Table 2
Modeling Results From the Information–Integration Conditions of Experiment 1

Best Fitting Model Type	Continuous	Discontinuous
Hypothesis-testing		
Frequency	2	3
Accuracy	0.64	0.62
SPC-1		
Frequency	9	7
Accuracy	0.75	0.69
SPC-2		
Frequency	5	3
Accuracy	0.83	0.73
SPC-4		
Frequency	0	3
Accuracy		0.78
Optimal		
Frequency	1	1
Accuracy	0.96	0.88

Note. SPC = striatal pattern classifier.

provided the best account of the data from none of the participants in the continuous condition but did provide the best account of the data from 3 participants in the discontinuous condition. Finally, in both the continuous and discontinuous conditions, there was a monotonic increase in accuracy as one moves from the hypothesis-testing models, to the SPC-1, to the SPC-2, to the SPC-4 and finally to the optimal model. For the continuous condition, the increase in accuracy from the participants best fit by a hypothesis-testing model to those best fit by the SPC-1 approached significance, $t(9) = 1.84$, $p = .10$, but was significant from the SPC-1 to the SPC-2, $t(12) = 2.48$, $p < .05$. For the discontinuous condition, the increase in accuracy approached significance from the participants best fit by a hypothesis-testing model to those best fit by the SPC-2, $t(4) = 2.58$, $p = .06$, from participants best fit by the SPC-1 to those best fit by the SPC-4, $t(8) = 6.49$, $p < .05$, and from participants best fit by a hypothesis-testing model to those best fit by the SPC-4, $t(4) = 5.27$, $p < .01$. Taken together, these results suggest that (a) the data of participants who were more accurate in each group were best fit by models with a greater number of striatal units and (b) reasonable accuracy can be obtained with only two striatal units per category in the continuous condition, but it requires four units per category in the discontinuous condition.

Recall that accuracy was higher (albeit not significantly) for participants whose data were best fit by the SPC-1 model in the continuous condition relative to those in the discontinuous condition. In addition, accuracy was significantly higher for participants whose data was best fit by the SPC-2 model in the continuous condition relative to those in the discontinuous condition. One advantage of the model-based approach is that we can examine the model fits and parameters in an attempt to localize the processes responsible for this accuracy difference. We begin by examining the fit value for the SPC-1 and SPC-2 across spread conditions. In both cases, the AIC values did not differ ($t < 1.0$ for the SPC-1 and SPC-2). This finding is important because it suggests that each specific model is providing an equivalent account of the participants' data across spread conditions. It is important to note that the models are also providing a good account of the data (average proportion of responses accounted for: SPC-1, continuous = .79; SPC-1, discontinuous = .79; SPC-2, continuous = .88; and SPC-2, discontinuous = .86). Next, we examined the noise parameter from each model that provides a measure of response determinism. For the SPC-1 model, the noise standard deviation was larger (albeit marginally) in the discontinuous condition (32.6) than in the continuous condition (20.1), $t(14) = 1.66$, $p < .10$; but for the SPC-2 model, the noise values did not differ (discontinuous = 10.2; continuous = 8.9; $t[6] < 1.0$). Thus, for the SPC-1 model, at least part of the accuracy deficit in the discontinuous condition appears to be due to less deterministic responding, but for the SPC-2 model this is not the case. Because of the large accuracy discrepancy across spread conditions for participants whose data was best fit by the SPC-2 model, it follows that the location of the striatal units are more suboptimal in the discontinuous than in the continuous spread condition. Finally, we examine the noise values across versions of the SPC model within each spread condition. This comparison provides some insight into the locus of the performance advantage for participants whose data were best fit by a more general SPC model. For the continuous condition, the average noise standard deviation was 20.1, and it was 8.9 for the SPC-1 and SPC-2 models, respectively. This difference was significant, $t(12) = 3.55$, $p < .01$, and suggests that participants whose data were best fit by the SPC-2 model evidenced more deterministic responding. For the discontinuous condition, the average noise standard deviation was 32.6, 10.2, and 6.3 for the SPC-1, SPC-2, and SPC-4 models, respectively. The noise difference between the SPC-1 and SPC-2 models was significant, $t(8) = 3.05$, $p < .05$, and the noise difference between the SPC-2 and SPC-4 models approached significance, $t(4) = 2.71$, $p = .06$. These data suggest that participants who were more deterministic in their responding were the same participants whose data were best fit by SPC models with a greater number of units.

Models Fit to the Rule-Based Condition Data

For completeness, we fit a number of models that instantiate the assumptions of the hypothesis-testing and procedural-based learning systems to the rule-based data as well. The SPC-1 model was applied to the data along with the conjunctive(1) model. Three special cases of the conjunctive(1) model were also applied. The optimal-length model assumes that the participant uses the same decision strategy as defined by the conjunctive(1) model, but it also assumes that the decision criterion along the length dimension is the optimal decision criterion. The optimal-orientation model is similar, but it assumes that the decision criterion along the orientation dimension is the optimal decision criterion. Both of these models have two free parameters (one decision criterion and the noise). The optimal model assumes that the length and orientation criteria are optimal. This model has one free parameter (the noise). As expected, no model differences emerged. The use of hypothesis-testing strategies was high in both conditions (11 of 17 in each condition were best fit by a hypothesis-testing model), and the model parameters did not differ across conditions, providing further evidence that within-category discontinuity does not affect rule-based category learning when the decision rule remains fixed.[2]

Discussion

The results from Experiment 1 suggest that rule-based category learning is unaffected by within-category discontinuity when the optimal decision bound remains constant. On the other hand, information-integration category learning is adversely affected by within-category discontinuity. The transfer results provide some evidence for a difference in the nature of the learning in the two systems. Rule-based category learning is abstract in that a "rule" is learned directly and learning is less tied to the specific stimulus regions trained, resulting in a strong distance-to-the-bound effect. In contrast, information-integration learning is more directly tied to the specific regions of the stimulus space that were trained, a rule is not learned directly, and distance from the trained regions strongly affects performance. Because a wider range of the stimulus space was trained in the information-integration discontinuous condition, generalization was superior.

The model based analyses of the information-integration condition data yielded a number of interesting results. First, the use of procedural-based learning strategies in the discontinuous spread condition was high and was equivalent to that observed in the continuous spread condition. Second, individual differences emerged in the number of striatal units needed to best fit participants' data. Although many participants in the continuous and discontinuous conditions were best fit by an SPC model that assumed one striatal-category unit per category (SPC-1), several participants in the continuous condition required two units per category (SPC-2), and some participants in the discontinuous condition required two (SPC-2) or four units per category (SPC-4). Third, for all three versions of the SPC, accuracy rates for participants best fit by each model were lower in the discontinuous condition than in the continuous condition. This effect was due to suboptimal striatal unit placement, less deterministic responding or both. Finally, a comparison of performance across the three versions of the SPC suggested that participants whose data were best fit by an SPC model with more units evidenced responding that was more deterministic and more nearly optimal. The implications of these findings will be addressed in greater detail in the General Discussion.

As expected, the use of hypothesis-testing strategies was high in the rule-based conditions, and no model parameter differences emerged.

Experiment 2: Verbal Rule Complexity Experiment

The goal of this article is to test the prediction that within-category discontinuity and verbal rule complexity interact in their effects on information-integration and rule-based category learning. Experiment 1 revealed that within-category discontinuity adversely affects information-integration but not rule-based category learning when the decision rule remains constant across within-category discontinuity conditions. Experiment 2 tests the hypothesis that verbal rule complexity will affect rule-based category learning but not information-integration category learning. Experiment 2 examined four category learning conditions. In the Figure 2A simple rule-based condition (hereafter referred to as RB-Simple), the participant had to learn one decision criterion along the length dimension, while ignoring the orientation dimension, and used the following decision rule: "Respond A if the length is short; respond B if the length is long." In the Figure 2C complex decision rule condition (hereafter referred to as RB-Complex) the stimuli were identical to those from the RB-Simple condition. However, in the complex condition the participant had to learn one decision criterion on length and a second on orientation and apply a disjunctive decision rule. The decision rule was as follows: "Respond A if the length is short and the orientation is shallow *or* if the length is long and the orientation is steep. Respond B if the length is short and the orientation is steep *or* if the length is long and the orientation is shallow." The latter is more complex because the number of operations necessary to instantiate the rule is larger, and it is disjunctive.

The Figure 2B, information-integration simple decision rule condition (hereafter referred to as II-Simple) was constructed from the rule-based simple condition stimuli by applying a linear transformation. This changed the nature of the optimal decision rule qualitatively while keeping the conditions structurally equivalent. The information-integration complex condition was constructed from the rule-based complex condition via the same linear transformation and is shown in Figure 2D (the II-Complex condition).

It is important to be note that the terms "simple verbal rule" and "complex verbal rule" are misnomers with respect to the information-integration conditions since we hypothesize that it is not a decision rule that is learned by the procedural-based learning system, but rather learning involves assigning responses to regions of the stimulus space. Even so, we will use the terms simple and complex decision rule to denote these two information-integration conditions. It is possible though that the two information-integration conditions might differ in complexity but for very different reasons. For example, the categories in the II-Simple condition are somewhat more coherent than those in the II-Complex condition. To be clear, the categories in the II-Complex condition are not distinct (i.e., discontinuous or separated in perceptual space) as they were in Experiment 1, but they are less coherent than in the II-Simple condition. In light of this fact, some performance deficit may emerge in the II-Complex condition relative to the II-Simple condition, albeit for a very different reason from that predicted for the RB-Complex and RB-Simple conditions. Even so, the effect (if any) is predicted to be small relative to the effect of verbal rule complexity in the rule-based conditions. Category learning was examined across five 96-trial blocks and a 144-trial transfer block in which no corrective feedback was provided. The transfer stimuli (see Figure 2E and 2F) were identical to those from Experiment 1.

[2] The fact that only 65% of the participants in each condition were best fit by a hypothesis-testing model was somewhat surprising. Usually 80%–85% or more are best fit by a hypothesis-testing strategy. A closer examination of the model fits suggests that in three of the six cases from the continuous and discontinuous conditions, the percentage of responses accounted for was less than 1% better for the procedural-based learning model than for the best fitting hypothesis-testing model, suggesting that a hypothesis-testing strategy was providing as good an account of the data.

Method

Participants

Sixty participants were solicited from the University of Texas community and received $6 or course credit for participating in this study. Fifteen participants completed each experimental condition. No participant completed more than one condition. Visual acuity was tested in each participant, and all participants had 20/20 vision, or vision corrected to at least 20/20.

Stimuli and Stimulus Generation

The stimuli and stimulus generation procedure were identical to those from Experiment 1 except that the category distribution parameters from Table 3 were used.

Procedure

The procedure was identical to that from Experiment 1.

Results

Training Block Analyses

A Category Structure (rule-based vs. information-integration) × Decision Rule Complexity (simple vs. complex) × Block (five 96-trial blocks) mixed design ANOVA was conducted on the training block accuracy rates. (It is worth mentioning that Experiment 1 used a within-particpant design whereas Experiment 2 used a between-participant design. Because the analyses of condition order in Experiment 1 failed to indicate any effect of order, the use of different designs does not appear to be a concern.) The accuracy rates averaged across participants are presented in Figure 4. The main effects of rule complexity [$F(1, 56) = 59.47$, $p < .001$] and block [$F(4, 224) = 13.63$, $p < .001$] were significant, but the main effect of category structure was not [$F < 1.0$]. Block did not interact with any other factor [Fs < 1.0], and the three-way interaction was nonsignificant [$F < 1.0$]. Category structure interacted with decision rule complexity [$F(1, 56) = 30.18$, $p < .001$]. To determine the locus of the interaction we compared perfor-

Table 3
Category Distribution Parameters From the Experiment 2 Decision Rule Complexity Experiment

Category	μ_1	μ_0	σ_1	σ_0	cov_{1o}	Category	μ_1	μ_0	σ_1	σ_0	cov_{1o}
Rule-Based Simple						Information-Integration Simple					
A_1	189	75	10	10	0	A_1	321	−27	10	10	0
A_2	189	109	10	10	0	A_2	297	−3	10	10	0
A_3	189	143	10	10	0	A_3	345	−3	10	10	0
A_4	189	177	10	10	0	A_4	321	21	10	10	0
A_5	223	75	10	10	0	A_5	273	21	10	10	0
A_6	223	109	10	10	0	A_6	248	45	10	10	0
A_7	223	143	10	10	0	A_7	297	45	10	10	0
A_8	223	177	10	10	0	A_8	273	69	10	10	0
B_1	257	75	10	10	0	B_1	369	21	10	10	0
B_2	257	109	10	10	0	B_2	345	45	10	10	0
B_3	257	143	10	10	0	B_3	393	45	10	10	0
B_4	257	177	10	10	0	B_4	369	69	10	10	0
B_5	291	75	10	10	0	B_5	321	69	10	10	0
B_6	291	109	10	10	0	B_6	297	93	10	10	0
B_7	291	143	10	10	0	B_7	345	93	10	10	0
B_8	291	177	10	10	0	B_8	321	118	10	10	0
Rule-Based Complex						Information–Integration Complex					
A_1	189	75	10	10	0	A_1	321	−27	10	10	0
A_2	189	109	10	10	0	A_2	297	−3	10	10	0
A_3	223	75	10	10	0	A_3	345	−3	10	10	0
A_4	223	109	10	10	0	A_4	321	21	10	10	0
A_5	257	143	10	10	0	A_5	321	69	10	10	0
A_6	257	177	10	10	0	A_6	297	93	10	10	0
A_7	291	143	10	10	0	A_7	345	93	10	10	0
A_8	291	177	10	10	0	A_8	321	118	10	10	0
B_1	189	143	10	10	0	B_1	273	21	10	10	0
B_2	189	177	10	10	0	B_2	248	45	10	10	0
B_3	223	143	10	10	0	B_3	297	45	10	10	0
B_4	223	177	10	10	0	B_4	273	69	10	10	0
B_5	257	75	10	10	0	B_5	369	21	10	10	0
B_6	257	109	10	10	0	B_6	345	45	10	10	0
B_7	291	75	10	10	0	B_7	393	45	10	10	0
B_8	291	109	10	10	0	B_8	369	69	10	10	0

Note. The subscripts denote cluster distributions for each category. Optimal accuracy was held constant at 95% in all conditions.

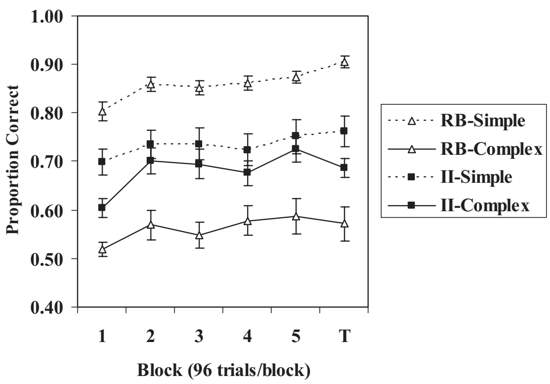

Figure 4. Proportion correct (averaged across participants) from Experiment 2 along with standard error bars. Blocks 1–5 are the training blocks and Block T is the transfer block.

mance across the simple and complex decision rules separately for rule-based and information-integration categories. For the rule-based categories, simple condition performance was .29 better than in the complex condition [$F(1, 28) = 128.61$, $p < .001$]. For the information-integration categories, simple condition performance was .05 better than in the complex condition, but this difference was nonsignificant [$F(1, 28) = 1.86$, $p = .18$]. Taken together these data suggest that increasing verbal rule complexity adversely affects rule-based category learning because the hypothesis-testing system learns a decision rule and more complex decision rules are more poorly learned. On the other hand, no performance difference emerged across the two information-integration conditions even though the categories were more coherent in one condition than in the other. Although future work should test this more fully, apparently the fact that there was no discontinuity allowed the procedural-based learning system to learn less coherent (albeit not discontinuous) categories nearly as well as it learned more coherent categories.

Transfer Block Analyses

Our only strong prediction with these data was that the increase in performance from the final training block to the transfer block would be greater for the rule-based than for the information-integration category structures. Performance increased in the rule-based conditions (average increase of .01) and decreased in the information-integration conditions (average decrease of .015), but this difference was nonsignificant [$F(1, 56) = 1.66$, $p = .20$]. Interestingly, there was a general performance increase in the simple relative to the complex condition for both category structures. This finding was unexpected.

Model-Based Analyses

Following the approach taken in Experiment 1, we fit models that instantiated the assumption that participants used an explicit hypothesis-testing strategy and models that instantiated the assumption that participants used an implicit procedural-based learning strategy. We applied two hypothesis-testing models to the simple rule-based condition. One was the optimal model that assumed the optimal decision criterion in the presence of noise (1 parameter), and the second was the suboptimal unidimensional model that estimated the decision criterion value and the noise from the data (2 parameters). The SPC-1 procedural-based learning model was also applied. We applied four hypothesis-testing models to the simple information-integration condition. One assumed a unidimensional rule on length, and a second assumed a unidimensional rule on orientation (2 parameters). The third and fourth models were variants of the conjunctive(1) model from Experiment 1. Both assumed that the participant learned a decision criterion on length and a decision criterion on orientation in the presence of noise (3 parameters total). One version assumed the decision rule "respond A if the line is short and shallow, otherwise respond B," and a second assumed the decision rule "respond B if the line is long and steep, otherwise respond A." The models fit to the data from the complex rule-based and information-integration conditions were identical to those fit to the data from Experiment 1 except that the response region to category label assignments were adjusted to reflect the fact that only two categories were relevant. Again, we focus on fits to the final block of training trials.

In the rule-based conditions, we found that 13 of the 15 simple condition participant's data and 12 of the 15 complex condition participant's data were best fit by a model that instantiated a hypothesis-testing strategy. As expected from an examination of Figure 4, participants whose data was best fit by a hypothesis-testing model in the simple condition were more accurate (88%) than those best fit by a hypothesis-testing model in the complex condition (58%) [$t(23) = 4.63$, $p < .01$]. To provide some insight into the locus of this performance decrement, we examined the parameters of the most general hypothesis-testing model [for the simple condition this is the suboptimal unidimensional model and for the complex condition this is the conjunctive(1) model]. The model accounted on average for .88 and .64 proportion of the responses in the simple and complex conditions, respectively. The ability of the model to account for the simple data was quite good and thus, we can draw meaningful inference about performance from fits of this model to these data. On the other hand, the proportion of responses accounted for in the complex condition is somewhat low. This makes parameter interpretation more precarious, and these data should be interpreted with caution.

We computed the average absolute deviation between the best fitting and optimal length criterion in the simple rule-based condition. The best fitting length criterion was 240.2 pixels and the optimal value was 240 pixels, suggesting excellent learning of the decision rule. In the complex condition the average best fitting length was 271 pixels compared with the optimal value of 240 pixels, and the average best fitting orientation was 46 degrees compared with the optimal value of 45 degrees. Thus, in the complex rule-based condition, participants did a good job of learning the criterion along orientation but did a poorer job of learning the decision criterion along length. The average noise variance was 20.5 units and 80.8 units in the simple and complex conditions respectively. These data suggest that the memory for, and trial by trial application of, the decision criteria was poorer in the complex condition. Taken together, these data suggest that participants in the complex rule-based condition found learning of the decision criteria more difficult, and the noise associated with its recollection and application more noisy than in the simple condition.

For completeness we fit a number of models that instantiated a hypothesis-testing strategy and a number of models that instantiated a procedural-based learning strategy to the information-integration data as well. The use of information-integration strategies was high in both conditions (12 of 15 in the II-Simple and 12 of 15 in the II-Complex conditions), and the model parameters did not differ across conditions, providing further evidence that decision rule complexity does not affect information-integration category learning.

Discussion

The results from Experiment 2 suggest that increasing the verbal complexity of the decision rule affects rule-based category learning, but structurally equivalent information-integration category learning conditions show no performance difference. This pattern is predicted if a hypothesis-testing system mediates the learning of rule-based categories and a procedural-based learning system mediates the learning of information-integration categories. The hypothesis-testing system learns the decision rule directly, and thus, verbally more complex decision rules should be more difficult to learn. On the other hand, the procedural-based learning system learns to assign categorization responses to regions of perceptual space and does not learn a decision rule directly. Thus, decision rule complexity should be irrelevant to information-integration category learning. We predicted that there might be a small performance difference across the two information-integration conditions because the categories in the II-Simple condition were more coherent. Although a small difference of .05 emerged, it was nonsignificant and paled in comparison to the large .30 difference observed across the two rule-based conditions. The model-based analyses suggest that the use of hypothesis-testing strategies was high in the simple and complex rule-based category learning conditions and that the accuracy deficit was due to poor learning of the optimal decision criteria and an increase in the variability associated with the participants' memory for and trial-by-trial application of the decision criteria. Because the rule in the complex verbal rule condition requires more working memory and attention to learn and apply, it seems reasonable that the performance deficit would be due to poor memory for and application of the decision criteria.

General Discussion

The aim of the current studies was to examine the detailed processing characteristics of a proposed procedural-based category learning system and a hypothesis-testing category learning system. Two processing aspects of the proposed systems were critical to the current report. The first is the proposal that procedural-based learning systems are more efficient when the stimuli are perceptually similar and form a coherent group, and are less efficient when stimuli are perceptually dissimilar and form distinct clusters (e.g., Maddox et al., 2005), whereas hypothesis-testing systems are unaffected by the distinctness of category clusters as long as the decision rule remains constant. The second is that the efficiency of hypothesis-testing systems is highly dependent on the complexity of the decision rule as a decision rule is learned directly, and thus more verbally complex decision rules should lead to worse learning (Feldman, 2001; Salatas & Bourne, 1974; Shepard et al., 1961). Procedural-based learning systems, on the other hand, are assumed not to learn a decision rule directly, but instead learn to assign responses to regions of the space. The response region boundaries are not learned directly but rather are an emergent property. Thus, as long as the nature of the stimulus space remains relatively unchanged (e.g., no discontinuity is introduced), learning in the procedural-based system should be unaffected by verbal decision rule complexity.

Taken together, these hypotheses led us to predict an interaction between the effects of within-category discontinuity and verbal decision rule complexity on the processing efficiency of the procedural-based learning and hypothesis-testing systems, and by extension on the learning of information-integration and rule-based categories. Specifically, we tested the prediction that an increase in within-category discontinuity will lead to a deficit in information-integration category learning but will have no effect on rule-based learning, whereas an increase in verbal rule complexity will lead to a deficit in rule-based category learning but will have no effect on information-integration category learning. Both predictions were supported by the data. Model-based analyses were conducted to provide some insights into the locus of these performance deficits. In the information-integration continuous and discontinuous conditions (Experiment 1), computational level descriptions of the procedural-based learning model (i.e., SPC) best fit the vast majority of the time. Interestingly, there were large individual differences in the number of SPC units needed to account for participants' responding, and in both the continuous and discontinuous spread conditions, accuracy increased as the number of units recruited increased. Even so, across all versions of the SPC, accuracy in the continuous spread condition was higher than accuracy in the discontinuous spread condition. This accuracy deficit in the discontinuous condition was due to suboptimal striatal unit placement, less deterministic responding or both. In the simple and complex rule-based conditions (Experiment 2), hypothesis-testing models fit best the vast majority of the time, but hypothesis-testing participants in the complex decision rule condition performed worse. The model parameters indicated that this performance deficit was due to poor decision criterion learning and an increase in the variability associated with the participant's memory for, and trial-by-trial application of, the decision criteria. This finding was not unexpected as the working memory and attention load necessary to learn and apply the complex rule is greater than for the simple rule. We turn now to a discussion of a number of relevant issues.

Model-Based Analyses of Within-Category Discontinuity Participant's Strategies (Experiment 1)

One of the most interesting findings from the model-based analyses of Experiment 1 was that the vast majority of the information-integration continuous (88%) and discontinuous condition participant's final block of training data (82%) were best fit by a computational model that instantiated the assumptions of the procedural-based learning system. This result is at odds with previous work that examined the effects of various manipulations on information-integration category learning (Ashby, Maddox, and Bohil, 2002; Ashby et al., 2003; Maddox et al., 2003). In these previous studies, participants were asked to learn information-integration categories either under observational conditions, delayed feedback conditions, or when the response locations assigned to each category switched. Participants in these conditions were more likely to yield data best fit by a hypothesis-testing model than by a procedural-based learning model as compared to participants attempting to learn information-integration categories under traditional feedback, immediate feedback, or no button switch conditions. COVIS assumes that there is an initial bias toward the hypothesis-testing system, but when information-integration categories are being learned, it is generally the case that, with enough experience, the procedural-based learning system eventually dominates because it generates a more accurate response. Under observational, delayed feedback, or button switch conditions, it appears that the procedural-based learning system is simply unable to learn and the hypothesis-testing system continues to dominate throughout the experiment. When the manipulation is one of within-category discontinuity, as in Experiment 1, this does *not* appear to be the case, at least not when the feedback follows the response immediately, and the button locations are held constant. Under these conditions, it seems reasonable that the procedural-based learning system should learn the task well enough to dominate the hypothesis-testing system. The model fits support this claim.

Even so, it was still the case that information-integration learning, albeit driven by the procedural-based learning system, was poorer in the discontinuous than in the continuous spread condition. This finding has important implications for our understanding of the neurocognitive processes involved in the procedural-based learning system. Specifically, one proposed characteristic of this system is that it prefers to associate a category label with perceptually similar stimuli. This proposal is based on the funnel-like connectivity between inferotemporal cortex and the posterior caudate that has a many-to-one cell convergence, and the relative lack of interneurons that would decrease the likelihood of communication among cells in this region (Wilson, 1995). Figure 5 provides an illustration of this system and how discontinuity might impact information-integration category learning. In Figure 5A, which depicts normal information-integration category learning with continuous categories, cells within inferotemporal cortex (represented as rectangles) project in a converging fashion to cells in the posterior caudate (represented as ovals). In this example, black rectangles represent cells processing category A stimuli and white rectangles represent cells processing category B stimuli. [Note that we are not arguing that the stimuli are categorized at the level of inferotemporal cortex (see Freedman, Riesenhuber, Poggio, & Miller, 2003); we simply distinguish between black and white rectangles at this level to help with the example.] At the level of the posterior caudate, black ovals represent a pattern of activation that should become associated with a category A response, whereas white ovals represent a pattern of activation that should become associated with a category B response. In general, because proximal cells within inferotemporal cortex process perceptually similar stimuli (Tanaka, 2000), and these cells project to the same cell in the posterior caudate (Cheng et al., 1997), a reasonable assumption is that perceptually similar stimuli will elicit a similar pattern of activity within the posterior caudate. Note in Figure 5A that the two black rectangles are represented as being close together in inferotemporal cortex, and the two white rectangles are close together.[3] As learning occurs via feedback and the dopamine reward signal, this association becomes stronger and the reactivation of a pattern of activity in the posterior caudate associated with a particular category response becomes facilitated by proximally similar cells, or stimuli, represented in inferotemporal cortex. As such, information-integration category learning proceeds normally and perceptually similar stimuli processed at the level of the inferotemporal cortex become associated together at the level of the posterior caudate.

In contrast, Figure 5B depicts a situation where stimulus clusters are discontinuous and learning is restricted. Specifically, in this example, the two black rectangles are far apart denoting distinct subclusters of stimuli to be associated with category A, and the two white rectangles are far apart denoting distinct subclusters of stimuli to be associated with category B. In this case, distinct clusters of perceptually dissimilar stimuli must be associated with a single category and hence the same response in the posterior caudate, whereas the system would rather associate perceptually similar stimuli with the same response. Thus, in this scenario, the pattern of activity occurring in the posterior caudate is less strongly associated with a particular category or response (which is represented by the gray ovals), making it difficult for the procedural-based learning system to acquire the categories. This should lead to poorer learning, which was found in Experiment 1.

However, participants were able to learn in the discontinuous information-integration condition, and our model-based analyses are consistent with the notion that participants adopted a procedural-based approach to learning (other computational models will be described shortly). This raises the interesting question of how participants were able to acquire discontinuous categories. One potential solution to this problem is to devote more of the applicable cellular processing resources during learning. Although highly speculative at this point, the results of the model-based analysis suggest that participants may have done so while learning the discontinuous categories. Specifically, we found that in the discontinuous information-integration condition, accuracy rates were greater for participants whose data were best fit by an SPC model with more units. If these units represent activity in the posterior caudate, as we hypothesize, then these model results suggest that as participants increase the number of cells associated with the discontinuous categories, better learning occurs. Again, this explanation is highly tentative at this point, but these findings raise some interesting predictions for future studies. For example, it is reasonable to predict that when learning discontinuous categories, there will be greater functional activity in the posterior caudate in individuals whose data is best fit by an SPC model with more units. Similarly, we might also anticipate that individuals whose data is fit by an SPC model more units will also perform better on other measures of procedural-based learning, such as SRT tasks.

Model-Based Analyses of Verbal Decision Rule Complexity Participant's Strategies (Experiment 2)

There is no a priori reason to predict that participants in the complex rule-based decision rule condition would shift away from hypothesis-testing strategies toward information-integration strategies, and no such shift was observed. It was an open question, however, where the locus of the complex rule-based decision rule performance deficit might reside. One possibility is that decision

[3] Although we represent these cells as lying along the surface of inferotemporal cortex, they are actually arranged in columns perpendicular to the surface of the cortex (Cheng et al., 1997; Fujita et al., 1992; Tanaka, 2000).

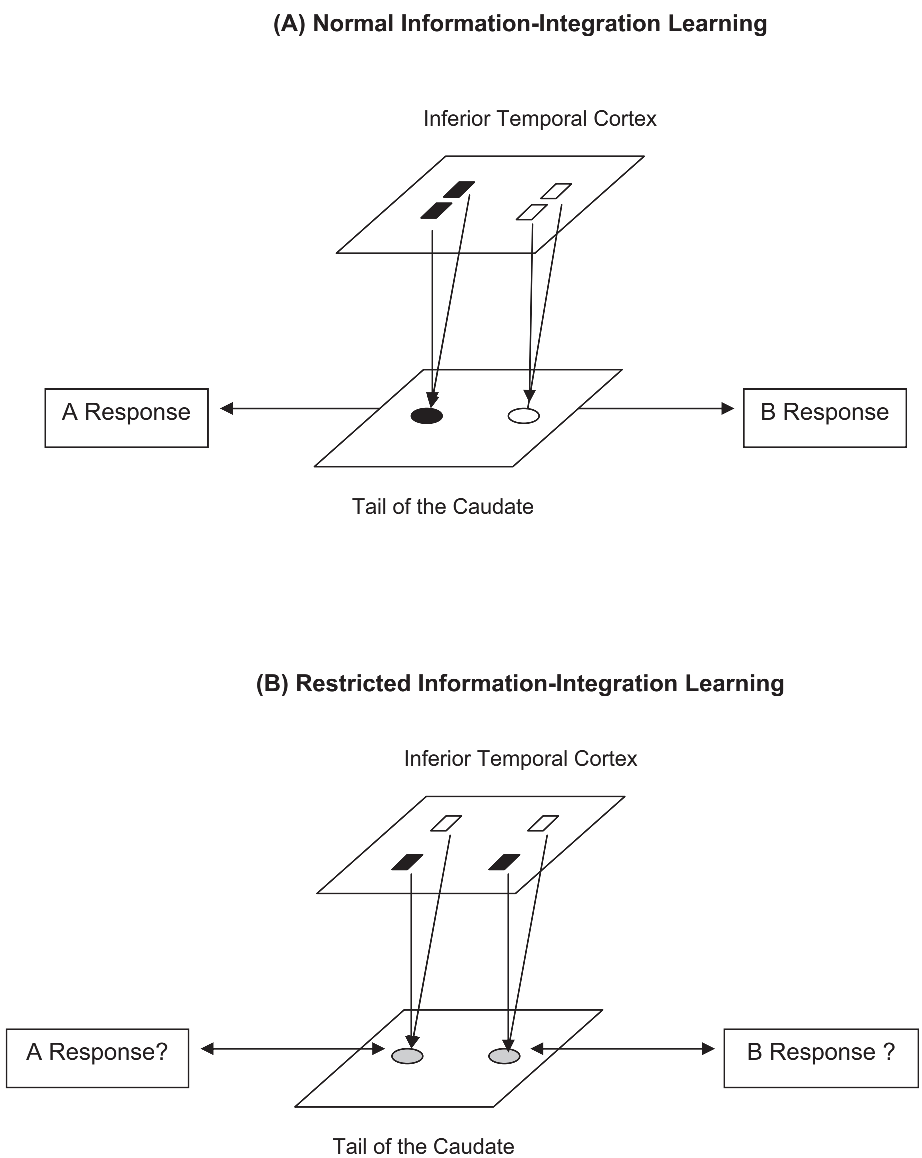

Figure 5. Schematic examples of (A) normal information-integration category learning with continuous stimulus clusters and (B) restricted information-integration category learning with discontinuous stimulus clusters.

rule complexity slows learning of the optimal decision criteria. A second possibility is that decision rule complexity leads to more noise in the memory for and trial-by-trial application of the decision rule. A third possibility is that decision rule complexity impacts both processes. The model-based analyses of Experiment 2 support the latter.

Two possible mechanisms that could account for the effect of rule complexity on rule-based category learning are the impact that this manipulation had on executive attention or working memory. In the present case, it is likely that both are relevant. The attentional demands are quite different across conditions with focused attention on the length dimension being relevant in the simple condition and a spreading of attention across both dimensions in the complex condition. In addition, the working memory load is likely quite different. In the simple condition the participant must store a very simple rule in working memory– "respond A to short lines and B to long lines." In the complex condition, on the other hand, the participant must store a more complex rule in working memory– "respond A if the line is short and shallow *or* long and steep, should respond B if the line is short and steep *or* long and shallow." This leads to a longer logical expression needed to correctly categorize the stimuli as compared to the simple condition. The impact of such manipulations has been reported in previous research (Feldman, 2001; Maddox, Filoteo, Hejl, & Ing, 2004; Salatas & Bourne, 1974; Shepard et al., 1961). The added length of the rule in the complex condition likely placed a greater emphasis on working memory and thus resulted in the decreased performance as compared to the simple condition. Overall, it is likely that both attentional demands and working memory requirements resulted in the differences we observed between the simple and the complex rule-based conditions.

Decision Rule "Complexity" Effects on Information-Integration Category Learning (Experiment 2)

In Experiment 2 we suggested that the terms simple versus complex verbal decision rule that applied to the rule-based category structures was a misnomer when applied to the information-integration category structures because the procedural-based learning system learns to assign category labels to regions of perceptual space and does not learn a "rule" in the same sense as that postulated for the hypothesis-testing system. However, we pointed out that the two information-integration conditions might differ in complexity because they differ somewhat in coherence. Specifically, the categories in the II-Simple condition are somewhat more coherent than those in the II-Complex condition. In light of this fact, we left open the possibility that a small effect might emerge because the sub-groups were less coherent in the "complex" information-integration condition. Despite this reasonable possibility, the small performance difference of .05 was nonsignificant and was much smaller than the .30 difference observed in the rule-based conditions. In Experiment 2 we speculated that the lack of any discontinuity (or distinctness) between the category subgroups might have led to the lack of an effect. The data from Experiment 2 certainly support this claim, but our sense is that conditions exist for which a small effect might be observed. For example, suppose that the basic nature of the "simple" and "complex" information-integration category structures was preserved but that the amount of stimulus space associated with each was expanded (i.e., more of the length-orientation space was sampled), or suppose that the "complex" condition included a discontinuity. It is possible that under these conditions, additional resources would have to be implemented in order to learn the task (which would theoretically be instantiated by the recruitment of additional striatal units in the posterior caudate) because more of the stimulus space is being sampled or because the categories are discontinuous. Under these conditions, some performance deficit may emerge in the complex condition relative to the simple condition. Importantly though, this effect would not be due to verbal rule complexity, but rather to something else, such as discontinuity.

Generalization Profiles for the Two Systems

The transfer results suggest that the nature of performance generalization might be very different in the two category learning systems. An abstract "rule" appears to be learned in the rule-based conditions that leads to a distance-to-the-bound effect whereby stimuli that are farther from the bound yield higher accuracy rates, even though stimuli from this region were never presented during training and might be quite distant from trained items. This rule is abstract in the sense that it can be applied to novel items and is not tied directly to the trained items. This result is very much in line with that predicted from an explicit, hypothesis-testing system. Information-integration learning is also partially affected by distance to the bound, but in contrast to the rule-based system, it appears to be more closely linked to the regions of the stimulus space that are trained, suggesting that generalization in information-integration tasks is strongly affected by distance-to-the-trained response region with more distant items yielding worse performance, a finding that is consistent with the known neurobiology of the striatum (i.e., convergence and lack of interneurons within the striatum).

Neurobiological Plausibility and Computational Alternatives

The approach taken in the current work was to consider specific aspects of the neuroanatomy that have been implicated in information-integration and rule-based classification and to use what is known about this neural circuitry (fully aware that the neuroanatomy is more complex than we assume and that much is unknown) to generate predictions that were then tested in two experiments. With respect to the proposed procedural-based learning system, two aspects of the neuroanatomy are generally agreed upon. First, it is well known that there is a many-to-one convergence of cells in inferotemporal cortex onto cells in the posterior caudate (Wilson, 1995). Second, only about 5% of caudate cells are interneurons and these are sparsely distributed throughout this nucleus (Wilson, 1995). The computational implications of this architecture are beginning to emerge. One possibility that we put forth in this article is that the funnel-like convergence of cells in inferotemporal cortex onto caudate cells leads to a low-resolution map of the perceptual space being represented in the posterior caudate. The results from Experiment 1 support this hypothesis, but other approaches have been offered in the literature. For example, recent work by Bar-Gad et al. (2003) suggests that the striatum might be better understood in terms of dimensionality reduction. Although these two accounts of the nature of the representations in the striatum are equally plausible based on the known biology of this brain region, the important assumption they both make is that there is little redundancy in the striatum (Kincaid, Zheng, & Wilson, 1998), supporting the role that this structure could play in category learning.

In general, the predictions we made regarding the neurobiological constraints of the striatum were supported, lending some insight into the processing characteristics of the procedural-based learning system. Nevertheless, it is important to consider some relevant issues regarding this line of research, including (1) whether enough is known about the neurobiology of the proposed structures to help guide hypotheses generation, and (2) whether one needs to appeal to the neurobiology at all, focusing instead on computational modeling approaches.

First, with respect to our neuroanatomical knowledge base, we acknowledge that we focus on only certain aspects of what is already known about the neurobiology of the striatum and do not consider other known properties. However, given the complexity of the striatum, this is a reasonable approach that has been taken by most theorists of striatal functioning. We also acknowledge that the there is much to be learned about the neuroanatomy and that future advances will likely change the nature of any neurobiologically inspired theory. Even so, in our opinion enough is known that reasonable predictions can be generated and tested. Over the past decade or so many have taken this approach and have made some striking discoveries. Importantly, a consideration of the neurobiology often leads to nonintuitive predictions in the sense that it is unlikely that a researcher would consider them had they not followed from an examination of the underlying neurobiology. Thus, this approach significantly expands the range of experimental manipulations that are examined and increases the size of the empirical testbed for cognitive theories. As an example, previous studies (Ashby et al., 2002; Maddox et al., 2003; Maddox & Ing, 2005) that examined whether manipulations of the nature and timing of the feedback would differentially impact procedural-based category learning as opposed to hypothesis-testing category learning were based directly on what is known about the nature and timing of the dopamine mediated reward signal (Beninger, 1983; Miller et al., 1981; Montague et al., 1996; J. Wickens, 1993). In the same spirit, the prediction from Experiment 1 that discontinuous categories would differentially impact procedural-based category learning as compared to hypothesis-testing category learning followed from the known neurobiology of the striatum.

Second, with respect to the notion that computational modeling approaches will suffice without reference to brain circuitry, we also acknowledge that computational modeling approaches are extremely useful and have advanced our understanding in many domains, especially classification learning (for a review see Ashby & Maddox, 1993, 1998; Love, Medin, & Gurekis, 2004; Nosofsky, 1992). Even so, there are a number of reasons to take a cognitive neuroscience approach and to make central issues of neurobiological plausibility. For one, and as suggested by Marr (1982), a complete theory should be formulated at many levels. Although the computational level is clearly important, formulations at the algorithmic and implementational (neurobiological) level are also important. An examination of the neurobiology helps constrain the quantitative modeling techniques that have become increasingly popular in cognitive research. By developing a model that considers what is currently known about the underlying neurobiology, the scientist minimizes the risk of building computational components into the model that are not supported by the neurobiology. In addition, it is well known in cognitive psychology, in general, and category learning, in particular, that models that make very different processing assumptions are often difficult or impossible to dissociate at the computational level. In fact, several articles have been published that derive direct equivalence mappings across parameters from disparate models that make very different processing assumptions (e.g., Ashby & Maddox, 1993; Nosofsky, 1990). Thus, whereas a number of computational models may be able to explain a pattern of results, a consideration of the neurobiological plausibility will help to differentiate whether the brain can support the processing assumptions of a given model. In addition, a number of other important byproducts of a neurobiological perspective can emerge. For example, by considering a neurobiological perspective we are not constrained to focus entirely on goodness-of-fit at the expense of generalizing across multiple levels of explanation. Similarly, by attending to the neurobiology, computational approaches that are biologically plausible can also help to direct studies attempting to understand how these brain regions are involved in various cognitive processes (e.g., Maddox & Filoteo, in press). The results of the present study, for example, might lead to a neurophysiological investigation of the role of the posterior caudate in category learning and help determine whether different regions in this structure are in fact tuned to different categories.

Although we have argued for the importance of neurobiological plausibility of category learning models, it is very important to acknowledge that computational models that have been developed without reference to neurobiology might also be able to account for the pattern of results in the present study. Some of the most successful computational models are RULEX (Nosofsky et al., 1994), ATRIUM (Erickson & Kruschke, 1998), and ALCOVE (Kruschke, 1992). Importantly, RULEX and ATRIUM posit a rule-based representation system that is very similar computationally to the hypothesis-testing system that we argue for in the present study. As such, it is likely that both of these models could provide a very good quantitative fit of the data from Experiment 2.[4] Because the rule-based component of both of these models appears to place a premium on working memory and attention, both of which are well known to be subserved by frontal-striatal circuitry, we take the liberty of arguing that neurobiologically and computationally the rule-based system in RULEX and ATRIUM is very similar to the hypothesis-testing system proposed here.

Where these other models might differ from ours is with respect to the category discontinuity effects observed in the information-integration conditions from Experiment 1. Although the increased within-category dissimilarity in the discontinuous condition would be predicted to lead to a performance deficit and thus may be accounted for by these models, the question remains regarding neurobiological plausibility. Whereas we postulate a procedural-based system that is mediated within the striatum, all three of these models (RULEX, ATRIUM, and ALCOVE) assume an exemplar-based system. To date, no assumptions regarding possible neurobiological underpinnings for these models have been proposed. As outlined earlier, fMRI studies implicate the posterior caudate in information-integration category learning (Nomura et al., in press; Seger & Cincotta, 2005) and several neuropsychological studies implicate the striatum more generally (Ashby et al., 2003; Filoteo et al., 2001a, 2005; Maddox & Filoteo, 2001). Thus, the current evidence supports the conclusion that the striatum, in general, and the posterior caudate, in particular, mediates information-integration category learning. One possibility is that the exemplar-based system in RULEX, ATRIUM, and ALCOVE provides a computational implementation of processing in the posterior caudate. However, this seems unlikely. Recall that one critical aspect of the posterior caudate is the funnel-like nature of the cortico-caudate connections. It is this funnel-like property, along with the fact that there is little information transfer across caudate cells due to the relative scarcity of interneurons, that leads to the prediction that discontinuous categories will be more difficult to learn. This funnel-like property is absent in the exemplar-based representations proposed in RULEX, ATRIUM, and ALCOVE. In these models, each stimulus or "exemplar" has its own representation. Thus, given our current understanding of the neurobiology, it appears that the exemplar-based models do not provide a good representation of the posterior caudate that is believed to mediate information-integration category learning. In addition, it is unclear what foundational aspect of the model would lead to predictions regarding the nature and timing of the feedback that follow naturally from the nature of the dopamine-mediated reward signal in the striatum (Ashby et al., 2002; Maddox et al., 2003; Maddox & Ing, 2005), or to predictions regarding manipulations of the cate-

[4] ALCOVE does not have a rule-base component, and so it would be difficult for this model to account for the data from Experiment 2.

gory response locations that follow from the link between the procedural-based learning system and procedural learning (Ashby et al., 2003; Maddox et al., 2004; however see Nosofsky et al., 2005).

Although exemplar-based models to not appear to provide a good accounting of information-integration category learning, it is nevertheless worth exploring its possible neurobiological underpinnings given this model's past computational success. Some have argued for a connection between exemplar-based processing and declarative memory. In fact, one success of exemplar models in the 1980s was in their ability to use the same representation to account for similarity, recognition, and categorization (see Nosofsky, 1992, for a review). Because declarative memory is mediated by medial temporal lobe structures, one possibility would be to tie the exemplar-based component with medial temporal lobe structures (e.g., Love & Gurekis, 2004). Indeed, Ashby and O'Brien (2005) suggested that declarative memory systems are used during explicit memorization. Thus, category structures that encourage memorization processes are especially likely to be learned by this system, and it is possible that an exemplar-based process is primarily responsible for learning under these conditions. Nevertheless, it is also important to point out that medial temporal lobe amnesiacs show normal information-integration category learning and retention (Filoteo, Maddox, & Davis, 2001b) as well as normal or near normal learning of other categories (Janowsky, Shimamura, Kritchevsky, & Squire, 1989; Leng & Parkin, 1988; Knowlton & Squire, 1993; Squire & Knowlton, 1995; however, see Zaki, Nosofsky, Jessup, & Unverzagt, 2003). These observations, along with the findings (reviewed above) implicating frontostriatal circuits in rule-based learning, and cortico-striatal circuits in information-integration learning, suggest that it would be incorrect to assume that all category learning is mediated by this medial temporal lobe based system. Even so, it is likely that this system mediates the learning of some category structures. In our opinion this type of theorizing regarding the neurobiological underpinnings of currently successful computational models is useful. Some computational modelers are beginning to take this approach (Love & Gurekis, 2004) and we urge others to follow.

Summary

Two experiments were conducted that provide a test of the predicted interaction between within-category discontinuity and verbal rule complexity on the efficiency of the neurobiologically motivated procedural-based and hypothesis-testing category learning systems, and by extension, on information-integration and rule-based category learning. As predicted, within-category discontinuity adversely affected information-integration category learning but not rule-based category learning, whereas verbal rule complexity adversely affected rule-based category learning but not information-integration category learning.

References

Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control, 19,* 716–723.

Alexander, G. E., DeLong, M. R., & Strick, P. L. (1986). Parallel organization of functionally segregated circuits linking basal ganglia and cortex. *Annual Review of Neuroscience, 9,* 357–381.

Aron, A. R., Shohamy, D., Clark, J., Myers, C., Gluck, M. A., & Poldrack, R. A. (2004). Human midbrain sensitivity to cognitive feedback and uncertainty during classification learning. *Journal of Neurophysiology, 92,* 1144–1152.

Ashby, F. G. (1992). Multivariate probability distributions. In F. G. Ashby (Ed.), *Multidimensional models of perception and cognition* (pp. 1–34). Mahwah, NJ: Erlbaum.

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

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

Ashby, F. G., & Gott, R. E. (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., & Maddox, W. T. (1993). Relations between prototype, exemplar, and decision-bound models of categorization. *Journal of Mathematical Psychology, 37,* 372–400.

Ashby, F. G., & Maddox, W. T. (2005). Human category learning. *Annual Review of Psychology, 56,* 149–178.

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

Ashby, F. G., Maddox, W. T., & Lee, W. W. (1994). On the dangers of averaging across subjects when using multidimensional scaling or the similarity-choice model. *Psychological Science, 5,* 144–151.

Ashby, F. G., & O'Brien, J. B. (2005). Category learning and multiple memory systems. *Trends in Cognitive Sciences, 9,* 83–89.

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

Ashby, F. G., Waldron, E. M., Lee, W. W., & Berkman, A. (2001). Suboptimality in human categorization and identification. *Journal of Experimental Psychology: General, 130,* 77–96.

Bar-Gad, I., Morris, G., & Bergman, H. (2003). Information processing, dimensionality reduction and reinforcement learning in the basal ganglia. *Progress in Neurobiology, 71,* 439–473.

Beninger, R. J. (1983). The role of dopamine in locomotor activity and learning. *Brain Research, 287,* 173–196.

Bruner, J. S., Goodnow, J. J., & Austin, G. A. (1956). *A study of thinking.* Oxford, England: Wiley.

Cheng, K., Saleem, K. S., & Tanaka, K. (1997). Organization of corticostriatal and corticoamygdalar projections arising from the anterior inferotemporal area TE of the macaque monkey: A phaseolus vulgaris leucoagglutinin study. *Journal of Neuroscience, 17,* 7902–7925.

Ell, S., Ing, A. D., & Maddox, W. T. (2006). *Working memory mediates the effects of delayed feedback on decision criterion learning in perceptual categorization*. Manuscript submitted for publication.

Erickson, M. A., & Kruschke, J. K. (1998). Rules and exemplars in category learning. *Journal of Experimental Psychology: General, 127*(2), 107–140.

Estes, W. K. (1956). The problem of inference from curves based on group data. *Psychological Bulletin, 53,* 134–140.

Feldman, J. (2001). Minimization of Boolean complexity in human concept learning. *Nature, 407,* 630–633.

Fernandez-Ruiz, J., Wang, J., Aigner, T. G., & Mishkin, M. (2001). Visual habit formation in monkeys with neurotoxic lesions of the ventrocaudal neostriatum. *Proceedings of the National Acadamy of Sciences of the United States of America, 98*(7), 4196–4201.

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

Filoteo, J. V., Maddox, W. T., Simmons, A. N., Ing, A. D., Cagigas, X. E., Matthews, S., & Paulus, M. P. (2005). Cortical and Subcortical Brain Regions Involved in Rule-based Category Learning. *NeuroReport, 16(2)*, 111–115.

Freedman, D. J., Riesenhuber, M., Poggio, T., & Miller, E. K. (2003). A comparison of primate prefrontal and inferior temporal cortices during visual categorization. *Journal of Neuroscience, 23*(12), 5235–5246.

Fujita, I., Tanaka, K., Ito, M., & Cheng, K. (1992). Columns for visual features of objects in monkey inferotemporal cortex. *Nature, 360*(6402), 343–346.

Heimer, L. (1995). *The human brain and spinal cord: Functional neuroanatomy and dissection guide (2nd ed.)*. New York, NY: Springer-Verlag Publishing.

Jahanshahi, M., Brown, R. G., & Marsden, C. D. (1992). The effect of withdrawal of dopaminergic medication on simple and choice reaction time and the use of advance information in Parkinson's disease. *Journal of Neurology, Neurosurgery and Psychiatry, 55,* 1168–1176.

Janowsky, J. S., Shimamura, A. P., Kritchevsky, M., & Squire, L. R. (1989). Cognitive impairment following frontal lobe damage and its relevance to human amnesia. *Behavioral Neuroscience, 103,* 548–560.

Keri, S. (2003). The cognitive neuroscience of category learning. *Brain Research Reviews, 43,* 85–109.

Kincaid, A. E., Zheng, T., & Wilson, C. J. (1998). Connectivity and convergence of single corticostriatal axons. *Journal of Neuroscience, 18*(12), 4722–4731.

Knowlton, B. J., & Squire, L. R. (1993). The learning of categories: Parallel brain systems for item memory and category knowledge. *Science, 262*(5140), 1747–1749.

Kruschke, J. K. (1992). ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review, 99*(1), 22–44.

Leng, N. R., & Parkin, A. J. (1988). Double dissociation of frontal dysfunction in organic amnesia. *British Journal of Clinical Psychology, 27*(4), 359–362.

Love, B. C., & Gureckis, T. M. (2004). The Hippocampus: Where a Cognitive Model meets Cognitive Neuroscience. *Proceedings of the 26th Annual Conference of Cognitive Science Society.*

Love, B. C., Medin, D. L., & Gureckis, T. M. (2004). SUSTAIN: A network model of category learning. *Psychological Review, 111*(2), 309–332.

Maddox, W. T. (1999). On the dangers of averaging across observers when comparing decision bound and generalized context models of categorization. *Perception & Psychophysics, 61,* 354–374.

Maddox, W. T. (2001). Separating perceptual processes from decisional processes in identification and categorization. *Perception & Psychophysics, 63,* 1183–1200.

Maddox, W. T. (2002). Learning and attention in multidimensional identification, and categorization: Separating low-level perceptual processes and high-level decisional processes, *Journal of Experimental Psychology: Learning, Memory, & Cognition, 28,* 99–115.

Maddox, W. T., & Ashby, F. G. (1998). Selective attention and the formation of linear decision bounds. Commentary on McKinley and Nosofsky (1996). *Journal of Experimental Psychology: Human Perception and Performance, 24*(1), 301–321.

Maddox, W. T., & Ashby, F. G. (2004). Dissociating Explicit and Procedural-Learning Based Systems of Perceptual Category Learning. *Behavioral Processes, 66,* 309–332.

Maddox, W. T., Ashby, F. G., & Bohil, C. J. (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., Ashby, F. G., Ing, A. D., & Pickering, A. D. (2004). Disrupting feedback processing interferes with rule-based, but not information-integration category learning. *Memory & Cognition, 32*(4), 582–591.

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

Maddox, W. T., & Filoteo, J. V. (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., & Filoteo, J. V. (2005). The Neuropsychology of Perceptual Category Learning. In H. Cohen & C. Lefebre (Ed.), *Handbook of Categorization in Cognitive Science.* (pp. 573–599). Oxford, England: Elsevier, Ltd.

Maddox, W. T., & Filoteo, J. V. (in press). Modeling visual attention and category learning in amnesiacs, striatal-damaged patients, and normal aging. In R. W. J. Neufeld (Ed.), *Advances in clinical-cognitive science: Formal modeling and assessment of processes and symptoms.*

Maddox, W. T., Filoteo, J. V., Hejl, K. D., Ing, A. D. (2004). Category Numerosity 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.

Maddox, W. T., Filoteo, J. V., Lauritzen, J. S., Connally, E., & Hejl, K. D. (2005). Disontinuous Categories Affect Information-Integration, but not Rule-Based Category Learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 31*(4), 654–669.

Maddox, W. T., & Ing, A. D. (2005). Delayed feedback disrupts the information-integration learning system but not the hypothesis testing system in perceptual category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 31*(1), 100–107.

Maddox, W. T., Lauritzen, J. S., & Ing, A. D. (2006). Cognitive complexity effects in perceptual classification are dissociable. *Memory & Cognition.*

Malamut, B. L., Saunders, R. C., & Mishkin, M. (1984). Monkeys with combined amygdalo-hippocampal lesions succeed in object discrimination learning despite 24-hour intertrial intervals. *Behavioral Neuroscience, 98,* 759–769.

Marr, J. (1982). Determinism. *Behavior Analyst, 5,* 205–207.

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

McDonald, R. J., & White, N. M. (1994). Parallel information processing in the water maze: Evidence for independent memory systems involving dorsal striatum and hippocampus. *Behavioral Neural Biology, 61,* 260–270.

Merchant, H., Zainos, A., Hernandez, A., Salinas, E., & Romo, R. (1997). Functional properties of primate putamen neurons during the categorization of tactile stimuli. *Journal of Neurophysiology, 77,* 1132–1154.

Miller, J. D. Sanghera, M. K., & German, D. C. (1981). Mesencephalic dopaminergic unit activity in the behaviorally conditioned rat. *Life Sciences, 29,* 1255–1263.

Mishkin, M., Malamut, B., & Bachevalier, J. (1984). Memories and habits: Two neural systems. In G. Lynch, J. L. McGaugh, & N. M. Weinberger (Eds.), *Neurobiology of* human *learning and memory* (pp. 65–77). New York: Guilford.

Montague, P. R., Dayan, P., & Sejnowski, T. J. (1996). A framework for mesencephalic dopamine systems based on predictive Hebbian learning. *Journal of Neuroscience, 16,* 1936–1947.

Myers, C. E., Shohamy, D., Gluck, M. A., Grossman, S., Onlaor, S., & Kapur, N. (2003). Dissociating medial temporal and basal ganglia memory systems with a latent learning task. *Neuropsychologia, 41,* 1919–1928.

Myung, I. J. (2000). The importance of complexity in model selection. *Journal of Mathematical Psychology, 44,* 190–204.

Nomura, E. M., Maddox, W. T., Filoteo, J. V., Ing, A. D., Gitelman, D. R., Parrish, T. B., et al. (in press). Neural correlates of rule-based and information-integration visual category learning. *Cerebral Cortex.*

Nosofsky, R. M. (1990). Relations between exemplar-similarity and likelihood models of classification. *Journal of Mathematical Psychology, 34,* 393–418.

Nosofsky, R. M. (1992). Exemplar-based approach to relating categorization, identification, and recognition. In F. G. Ashby (Ed.), *Multiple dimensional models of perception and cognition* (pp. 363–393). Mahwah, NJ: Erlbaum.

Nosofsky, R. M., Palmeri, T. J., & McKinley, S. C. (1994). Rule-plus exception model of classification learning. *Psychological Review, 101,* 53–79.

Nosofsky, R. M., Stanton, R. D., & Zaki, S. R. (2005). Procedural interference in perceptual classification: Implicit learning or cognitive complexity? *Memory & Cognition, 33,* 1256–1271.

Packard, M. G., & McGaugh, J. L. (1992). 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.

Pitt, M. A., Myung, I. J., & Zhang, S. (2002). Toward a method of selecting among computational models of cognition. *Psychological Review, 109,* 472–491.

Poldrack, R. A., Clark, J., Paré-Blagoev, E. J., Shohamy, D., Moyano, J. C., Myers, C., & Gluck, M. A. (2001, November 29). Interactive memory systems in the human brain. *Nature, 414,* 546–550.

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

Poldrack, R. A., Prabhakaran, V., Seger, C. A., & Gabrieli, J. D. E. (1999). Striatal activation during acquisition of a cognitive skill. *Neuropsychology, 13,* 564–574.

Reber, P. J., Gitelman, D. R., Parrish, T. B., & Mesulam, M. M. (2003). Dissociating explicit and implicit category knowledge with fMRI. *Journal of Cognitive Neuroscience, 15,* 574–583.

Reber, P. J., Stark, C. E. L., & Squire, L. R. (1998). Cortical areas supporting category learning identified using functional magnetic resonance imaging. *Proceedings of the National Academy of Sciences USA, 95,* 747–750.

Romo, R., Merchant, H., Ruiz, S., Crespo, P., & Zainos, A. (1995). Neuronal activity of primate putamen during categorical perception of somaesthetic stimuli. *Neuroreport, 6,* 1013–1017.

Romo, R., Merchant, H., Zainos, A., & Hernandez, A. (1997). Categorical perception of somaesthetic stimuli: Psychophysical measurements correlated with neuronal events in primate medial premotor cortex. *Cerebral Cortex, 7,* 317–326.

Saint-Cyr, J. A., Ungerleider, L. G., & Desimone, R. (1990). Organization of visual cortical inputs to the striatum and subsequent outputs to the pallido-nigral complex in the monkey. *The Journal of Comparative Neurology, 298,* 129–156.

Salatas, H., & Bourne, L. E. (1974). Learning conceptual rules: III. Processes contributing to rule difficulty. *Memory & Cognition, 2,* 549–553.

Schultz, W., Apicella, P., & Ljungbert, T. (1993). Responses of monkey dopamine neurons to reward and conditioned stimuli during successive steps of learning a delayed rsponse task. *The Journal of Neuroscience, 13,* 900–913.

Seger, C. A., & Cincotta, C. M. (2002). Striatal activity in concept learning. *Cognitive, Affective, & Behavioral Neuroscience, 2,* 149–161.

Seger, C. A., & Cincotta, C. M. (2005). The roles of the caudate nucleus in human classification learning. *Journal of Neuroscience, 25,* 2941–2951.

Shepard, R. N., Hovland, C. I., & Jenkins, H. M. (1961). Learning and memorization of classifications. *Psychological Monographs, 75*(Whole No. 517, Pt. 13).

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

Smith, J. D., & Minda, J. P. (1998). Prototypes in the mist: The early epochs of category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 24,* 1411–1436.

Squire, L. R., & Knowlton, B. J. (1995). Memory, hippocampus, and brain systems. In M. S. Gazzanga (Ed.), *The cognitive neurosciences* (pp. 825–837). Cambridge, MA: MIT Press.

Takane, Y., & Shibayama, T. (1992). Structures in stimulus identification data. In F. G. Ashby (Ed.), *Multidimensional models of perception and cognition* (pp. 335–362). Mahwah, NJ: Erlbaum.

Tanaka, K. (2000). Mechanisms of visual object recognition studied in monkeys. *Spatial Vision, 13*(2–3), 147–163.

Teng, E., Stefanacci, L., Squire, L. R., & Zola, S. M. (2000). Contrasting effects on discrimination learning after hippocampal lesions and conjoint hippocampal-caudate lesions in monkeys. *Journal of Neuroscience, 20,* 3853–3863.

Van Hoesen, G. W., Yeterian, E. H., & Lavizzo-Mourey, R. (1981). Widespread corticostriate projections from temporal cortex of the rhesus monkey. *Journal of Comparative Neurology, 199,* 205–219.

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

Webster, M. J., Bachevalier, J., & Ungerleider, L. G. (1993). Subcortical connections of inferior temproral areas TE and TEO in macaque monkeys. *Journal of Comparative Neurology, 335,* 73–91.

Wickens, J. (1993). *A theory of the striatum.* New York: Pergamon Press.

Wickens, T. D. (1982). *Models for behavior: Stochastic processes in psychology.* San Francisco: Freeman.

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

Willingham, D. B., Nissen, M. J., & Bullemer, P. (1989). On the development of procedural knowledge. *Journal of Experimental Psychological Learning Memory and Cognition, 15,* 1047–1060.

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

Wilson, C. J. (1995). The contribution of cortical neurons to the firing pattern of striatal spiny neurons (pp. 29–50). Cambridge, MA: MIT Press.

Zaki, S. R., Nosofsky, R. M., Jessup, N. M. & Unverzagt, F. W. (2003). Categorization and recognition performance of a memory-impaired group: Evidence for single-system models. *Journal of the International Neuropsychological Society, 9,* 394–406.

Zeithamova, D., & Maddox, W. T. (2006). Dual-task interference in perceptual category learning. *Memory & Cognition, 34,* 387–398.

Received April 6, 2006
Revision received June 2, 2006
Accepted July 25, 2006 ■