Journal of Experimental Psychology:
Learning, Memory, and Cognition
2003, Vol. 29, No. 4, 650–662

0278-7393/03/$12.00 DOI: 10.1037/0278-7393.29.4.650

Delayed Feedback Effects on Rule-Based and Information-Integration Category Learning

W. Todd Maddox
University of Texas at Austin

F. Gregory Ashby
University of California, Santa Barbara

Corey J. Bohil
University of Texas at Austin

The effect of immediate versus delayed feedback on rule-based and information-integration category learning was investigated. Accuracy rates were examined to isolate global performance deficits, and model-based analyses were performed to identify the types of response strategies used by observers. Feedback delay had no effect on the accuracy of responding or on the distribution of best fitting models in the rule-based category-learning task. However, delayed feedback led to less accurate responding in the information-integration category-learning task. Model-based analyses indicated that the decline in accuracy with delayed feedback was due to an increase in the use of rule-based strategies to solve the information-integration task. These results provide support for a multiple-systems approach to category learning and argue against the validity of single-system approaches.

Categorization is fundamental to the survival of all organisms (Ashby & Maddox, 1998). Every day people make thousands of categorization judgments and are often remarkably accurate. An understanding of the psychological processes involved in category learning, and whether different processes are involved in learning different types of category structures, is of critical importance. We learn about categories in a number of different ways (e.g., Ashby, Maddox, & Bohil, 2002; Estes, 1994; Yamauchi & Markman, 2000a, 2000b). For example, we might be given the category label prior to viewing the object, as might be the case when a medical school professor informs the students that they will be shown a series of lung X-rays that are all indicative of cancer. Another possibility is that we might be presented with an object, be asked to generate a categorization response, and receive immediate feedback, as might be the case when a medical student intern examines an X-ray under the supervision of an MD. A third possibility is that we might be presented with an object, generate a response, and only later receive feedback, as might be the case when a medical student examines an X-ray, makes a determination, and is later informed of the accuracy of his or her diagnoses. Despite the fact that we learn about categories in a variety of ways, the focus of nearly all categorization research has been on learning situations in which the observer is presented with a stimulus, generates a response, and is given immediate feedback.

This article examines the effects of feedback delay on category learning. In each condition of all experiments, observers learned two categories of Gabor patches (sine-wave gratings in which contrast is modulated by a circular Gaussian filter) that varied across trials in spatial frequency and orientation. The basic design is illustrated in Figure 1. In the immediate feedback conditions, the stimulus was displayed on each trial until the observer responded with a category label (i.e., "A" or "B"). Next a mask (i.e., a different Gabor patch) was presented for 500 ms, followed immediately by corrective feedback. Finally, there was a 5-s delay before the start of the next trial. The analogous delay condition was identical except that the mask was displayed for 5 s and the intertrial interval (ITI) was 500 ms. We also included an identical pair of conditions in which the 5-s delays were replaced with 2.5- and 10-s delays.

A number of recent results have suggested that such feedback delays might have different effects on performance depending on the specific category structures that are chosen. In this article, we focus on two different types of category structures (Ashby & Ell, 2001). *Rule-based category-learning tasks* are those in which the category structures can be learned by means of some explicit reasoning process. Frequently, the rule that maximizes accuracy (i.e., *the optimal rule*) is easy to describe verbally (Ashby, Alfonso-Reese, Turken, & Waldron, 1998). In the most common applications, only one stimulus dimension is relevant, and the observer's task is to discover this relevant dimension and then to map the different dimensional values to the relevant categories.

W. Todd Maddox and Corey J. Bohil, Department of Psychology, University of Texas at Austin; F. Gregory Ashby, Department of Psychology, University of California, Santa Barbara.

This research was supported in part by National Science Foundation Grants SBR-9796206 and BCS99-75037 and National Institute of Health Grant R01 MH59196. We thank Amy Kaderka for help with data collection.

Correspondence concerning this article should be addressed to W. Todd Maddox, Department of Psychology, 1 University Station A8000, University of Texas, Austin, Texas 78712, or to F. Gregory Ashby, Department of Psychology, University of California, Santa Barbara, California 93106. E-mail: maddox@psy.utexas.edu or ashby@psych.ucsb.edu

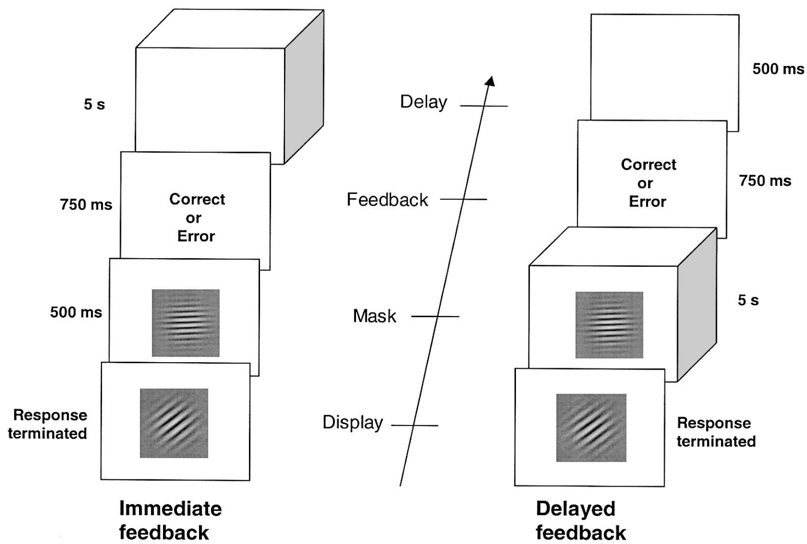

Figure 1. Basic design for the immediate and delayed feedback conditions.

Rule-based tasks have a long history in cognitive psychology, and not surprisingly, they have been popular with proponents of the so-called *classical theory of categorization,* which assumes category learning is the process of discovering the set of necessary and sufficient conditions that determine category membership (e.g., Smith & Medin, 1981).

Information-integration category-learning tasks are those in which accuracy is maximized only if information from two or more stimulus components (or dimensions) is integrated at some predecisional stage (Ashby & Gott, 1988). Perceptual integration could take many forms—from treating the stimulus as a Gestalt to computing a weighted linear combination of the dimensional values. However, a *conjunction rule* (e.g., "Respond 'A' if the stimulus is small on dimension x and small on dimension y") is a rule-based task rather than an information-integration task because separate decisions are first made about each dimension (e.g., small or large) and then the outcome of these decisions is combined (integration is not predecisional). In many cases, the optimal rule in information-integration tasks is difficult or impossible to describe verbally (Ashby et al., 1998). Information-integration tasks have been favorites of exemplar theorists, who argue that categorization requires accessing the memory representations of every previously seen exemplar from each relevant category (Estes, 1994; Medin & Schaffer, 1978; Nosofsky, 1986), and decision bound theorists, who argue that category learning is a process of associating category labels with regions of perceptual space (Ashby & Gott, 1988; Ashby & Maddox, 1992, 1993). In typical applications, however, exemplar theorists have used information-integration tasks with few exemplars per category, whereas decision bound theorists have used many exemplars per category.

The two category structures used in the present study are described in Figure 2. Each symbol in Figure 2 denotes the spatial frequency and orientation of a single Gabor patch. Category A exemplars are denoted by circles and Category B exemplars are denoted by squares. In each condition, there were two distinct categories that did not overlap, so perfect accuracy was always

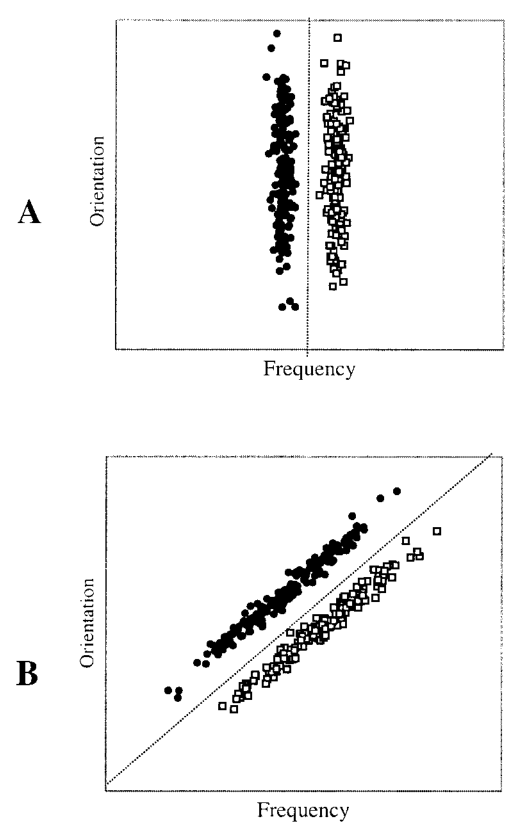

Figure 2. A: Rule-based category structure from Experiment 1. B: Information-integration category structure from Experiment 1. Each circle denotes the spatial frequency and spatial orientation of a Gabor pattern from Category A. Each square denotes the spatial frequency and spatial orientation of a Gabor pattern from Category B. The dotted line in each panel denotes the location of the optimal decision bound.

possible. Also shown in Figure 2 are the decision bounds that maximize categorization accuracy. In Figure 2A, the optimal bound requires observers to attend to spatial frequency and ignore orientation, so we call these *unidimensional categories.* With the diagonal categories of Figure 2B, which were generated by rotating the unidimensional categories by 45°, equal attention must be allocated to both stimulus dimensions. In this condition, the most accurate unidimensional rule yields a response accuracy of about 70%. In addition, because of the continuous-valued stimulus dimensions, it would be difficult or impossible to respond optimally in the diagonal condition by using a unidimensional rule and memorizing exceptions.

As shown in Figure 2, the stimulus dimensions of Gabor patches are separable, have simple verbal labels (bar width and orientation), and have no emergent (or configural) features (e.g., all patches are exactly the same size and shape, regardless of frequency or orientation). For these reasons, the unidimensional condition is a rule-based task because there is a simple explicit rule that separates the contrasting categories. In particular, the vertical bound in Figure 2A corresponds to the rule: "Respond 'A' if the bars are thick and 'B' if they are thin." In contrast, the diagonal condition is an information-integration task because perfect accuracy requires integrating spatial frequency and orientation information, and there is no simple verbal description of the optimal decision bound.

The category structures shown in Figure 2 were used by Ashby, Queller, and Berretty (1999) in a study of unsupervised categorization. In a series of experiments, observers were told the number of categories (i.e., two) and that perfect accuracy was possible, and they were given extensive experience in the task (i.e., 800 trials), but they were never given any feedback about the accuracy of their responses. Thus, participants in the Ashby, Queller, and Berretty study received no feedback of any kind. Despite their lack of feedback, however, by the end of the single experimental session, every observer in the rule-based conditions was responding with near perfect accuracy, whereas every observer in the two information-integration conditions was responding with accuracy that was far below optimal. Instead they all used some sort of unidimensional rule, even when explicitly encouraged to use both stimulus dimensions. When trial-by-trial feedback was provided in a separate experiment, all participants responded optimally in the diagonal conditions. These results provide clear evidence that feedback is much more important in information-integration category learning than in rule-based category learning.

There have been a number of recent proposals that category learning in rule-based and information-integration tasks is mediated by separate systems (Ashby et al., 1998; Ashby & Ell, 2002b; Erickson & Kruschke, 1998). Ashby and his colleagues have presented evidence that the learning of rule-based categories is mediated by an explicit-reasoning/hypothesis-testing system, whereas the learning of information-integration categories (like those shown in Figure 2B) is mediated by a procedural-learning based system (Ashby et al., 1998; Ashby, Isen, & Turken, 1999; Ashby & Waldron, 1999; Ashby, Waldron, Lee, & Berkman, 2001; Waldron & Ashby, 2001). If learning in rule-based and information-integration categorization tasks is mediated by separate systems, then it should be possible to find other dissociations, besides the one reported by Ashby, Queller, and Berretty (1999). In fact, we have recently found evidence of a number of such dissociations. Collectively, these data provide strong evidence that learning in these two types of tasks is mediated by separate systems.

First, Ashby, Maddox, and Bohil (2002) trained observers on rule-based and information-integration categories using an observational training paradigm in which participants were informed before stimulus presentation of what category the ensuing stimulus was from. Following stimulus presentation, participants then pressed the appropriate response key. Traditional feedback training was as effective as observational training with rule-based categories, but with information-integration categories, feedback training was significantly more effective than observational training. This result, together with that of Ashby, Queller, and Berretty (1999), suggests that with information-integration categories, learning is most effective when feedback is given after the response, whereas rule-based category learning is quite flexible with respect to the nature of the training signal. As a result, we might predict that delayed feedback might be more disruptive in information-integration category learning than with rule-based categories.

Another qualitative difference between these two tasks is that information-integration category learning is more closely tied to motor outputs than rule-based category learning. Ashby, Ell, and Waldron (in press) had observers learn either rule-based or information-integration categories using traditional feedback training. Next, some observers continued as before, some switched their hands on the response keys, and for some the location of the response keys was switched (so the Category A key was assigned to Category B and vice versa). For those observers learning rule-based categories, there was no difference among any of these transfer instructions, thereby suggesting that abstract category labels are learned in rule-based categorization. In contrast, for those observers learning information-integration categories, switching hands on the response keys caused no interference, but switching the locations of the response keys caused a significant decrease in accuracy. Thus, it appears that response locations are learned in information-integration categorization, but not specific motor programs.

One possible problem with these results is that information-integration tasks are usually more difficult than rule-based tasks in the sense that information-integration tasks usually require more training to reach the same level of expertise. Because of this difficulty difference, one concern is that, collectively, these studies might only show that it is relatively easy to disrupt learning in difficult tasks compared with simpler tasks. However, several results argue strongly against this hypothesis. First, Waldron and Ashby (2001) showed that a simple rule-based category-learning task (which required attending to a single stimulus dimension) was disrupted more by a simultaneous task that activates frontal cortex (a numerical Stroop task) than was a complex information-integration task (which required attending to three stimulus dimensions). If difficulty is the most important factor, then simultaneously performing a second task should interfere more strongly with learning the more difficult information-integration structures. Because Waldron and Ashby found the opposite pattern of results, it seems likely that factors other than difficulty are at work. Second, Ashby, Noble, Filoteo, Waldron, and Ell (2003) found that the same group of Parkinson's disease patients were much more impaired at rule-based category learning (one dimensional) than at information-integration category learning (three dimensional). If a single system mediates learning in these two types of categorization tasks, and if Parkinson's disease damages this system, then we would expect the more serious deficits to occur in the more difficult information-integration tasks.

These dissociations between rule-based and information-integration category learning are important for several reasons. First, they demonstrate that it is invalid to predict that learning in the easier of the two conditions shown in Figure 2 (i.e., the unidimensional condition) is necessarily more resistant to interference than learning in the more difficult (diagonal) condition. In the present study, this means that the difficulty difference alone is not enough to allow us to predict that delaying feedback will be more detrimental in the diagonal conditions. Instead, our previous results show that some kinds of interference more strongly affect the difficult information-integration task and some kinds more strongly affect the simpler rule-based task. Second, these dissociations strongly argue that people learn rule-based and information-integration categories using separate systems. For example, consider just the single Waldron and Ashby (2001) dual-task experiment. Arguably the most successful existing single-process model of category learning is Kruschke's (1992) attentional learning covering map (ALCOVE) model. Ashby and Ell (2002a) showed that the only versions of ALCOVE that could fit Waldron and Ashby's data made the strong prediction that after reaching criterion accuracy on the simple (unidimensional) rule-based structures, participants would have no idea that only one dimension was relevant in the dual-task conditions. Ashby and Ell reported new empirical evidence that strongly disconfirmed this prediction of ALCOVE. Thus, the best available single-system model failed to account even for the one dissociation reported by Waldron and Ashby (2001).

Neurobiological Underpinnings of Rule-Based and Information-Integration Category Learning

A wide variety of evidence now indicates that the learning of rule-based and information-integration category structures might be mediated by different neural circuits, and this literature provides further reason to expect that the experimental manipulations studied in this article might have different effects on these two types of tasks (Ashby et al., 1998; Ashby & Ell, 2002a; Erickson & Kruschke, 1998; Smith, Patalano, & Jonides, 1998; Waldron & Ashby, 2001). In particular, Ashby and his colleagues have proposed that with rule-based structures, learning is mediated by a circuit that includes the anterior cingulate, the prefrontal cortex, and the head of the caudate nucleus, whereas in information-integration tasks, learning is mediated largely within the tail of the caudate nucleus (with visual stimuli; Ashby et al., 1998; Ashby & Ell, 2001; Ashby, Isen, & Turken, 1999; Ashby & Waldron, 1999). In primates, all of extrastriate visual cortex projects directly to the tail of the caudate nucleus, with about 10,000 visual cortical cells converging on each caudate cell (Wilson, 1995). Cells in the tail of the caudate (medium spiny cells) then project to prefrontal and premotor cortex (via the globus pallidus and thalamus; e.g., Alexander, DeLong, & Strick, 1986). Ashby and Waldron (1999) hypothesized that through a procedural learning process, each caudate unit learns to associate a category label, or perhaps an abstract motor program, with a large group of visual cortical cells (i.e., all that project to it).

A review of all the evidence supporting this hypothesis is beyond the scope of this article. Instead, we mention only a few of the more prominent supporting results. We begin with a few results supporting the role of the anterior cingulate, prefrontal cortex, and head of the caudate nucleus in rule-based category learning. First, patients with lesions of the prefrontal cortex are well known to be impaired in rule-based tasks (e.g., such as the Wisconsin Card Sorting Test; see Robinson, Heaton, Lehman, & Stilson, 1980), but not in information-integration tasks (Knowlton, Mangels, & Squire, 1996). Second, a functional magnetic resonance imaging (fMRI) study of a rule-based task showed activation in the right dorsal-lateral prefrontal cortex, the anterior cingulate, and the head of the right caudate nucleus (Rao, et al., 1997). Third, many studies implicate these structures as key components of executive attention (e.g., Posner & Petersen, 1990) and working memory (e.g., Goldman-Rakic, 1987, 1995), both of which are likely to be critically important to the explicit processes of rule formation and testing that are assumed to mediate rule-based category learning. Fourth, a recent neuroimaging study identified the (dorsal) anterior cingulate as the site of hypothesis generation in a rule-based category-learning task (Elliott & Dolan, 1998). Finally, lesion studies in rats implicate the dorsal caudate nucleus in rule switching (Winocur & Eskes, 1998).

A prominent role for the tail of the caudate nucleus in (visual) information-integration category learning was supported by several studies that have reported information-integration category-learning deficits in patients with diseases of the basal ganglia (i.e., Parkinson's or Huntington's disease; Filoteo, Maddox, & Davis, 2001b; Knowlton et al., 1996; Maddox & Filoteo, 2001). In contrast, medial temporal lobe amnesic patients were normal (Filoteo, Maddox, & Davis, 2001a). Other evidence for a basal ganglia contribution to category learning comes from a long series of lesion studies in rats and monkeys that showed that the tail of the caudate nucleus[1] is both necessary *and* sufficient for normal visual discrimination learning. In primates, all of extrastriate visual cortex projects directly to the tail of the caudate nucleus, and the cells in this area then project, via the globus pallidus (the output portion of the basal ganglia) and thalamus, to the prefrontal and premotor cortices. These projections place the caudate in an ideal position to link percepts and actions, and many researchers have hypothesized that this is its primary role (e.g., Rolls, 1994; J. Wickens, 1993). Many studies have shown that lesions of the tail of the caudate nucleus impair the ability of animals to learn visual discriminations that require one response to one stimulus and a different response to some other stimulus (e.g., McDonald & White, 1993, 1994; Packard, Hirsh, & White, 1989; Packard & McGaugh, 1992). Since the visual cortex is intact in these animals, it is unlikely that their difficulty is in perceiving the stimuli. Rather, it appears that their difficulty is in learning to associate an appropriate response with each stimulus alternative. Technically, such studies were categorization tasks with one exemplar per category.

[1] In the rat, the caudate and putamen merge into a single entity, so for the rat studies, it is more proper to refer to the dorsal striatum.

It is difficult to imagine how adding more exemplars to each category could alleviate the deficits caused by caudate lesions, and it is for this reason that the caudate lesion studies support the hypothesis that the caudate contributes to normal category learning. The sufficiency of the caudate nucleus for visual discrimination learning was shown in a series of studies by Gaffan and colleagues that lesioned all pathways out of visual cortex except into the tail of the caudate (e.g., projections into prefrontal cortex were lesioned by Eacott & Gaffan, 1992, and Gaffan & Eacott, 1995; projections to the hippocampus and amygdala were lesioned by Gaffan & Harrison, 1987). None of these lesions prevented visual-discrimination learning.

These results are important because the mechanisms that mediate learning-related changes in synaptic efficacy within these two neural circuits are qualitatively different, and such differences suggest that delaying the feedback may have different effects on rule-based and information-integration tasks. The rule-based category-learning system proposed above is under conscious control and has full access to working memory and executive attention. As a result, the timing of the feedback signal should not matter in rule-based tasks. In contrast, an information-integration category-learning system that is mediated within the tail of the caudate nucleus would not be accessible to conscious awareness and is far removed from working memory. As a result, it would depend more heavily on local learning mechanisms.

Within the tail of the caudate nucleus, a reward-mediated feedback signal is thought to be provided by dopamine released from the substantia nigra (e.g., J. Wickens, 1993). Specifically, dopamine is released into the tail of the caudate (among other regions) from the substantia nigra (*pars compacta*) shortly after the animal receives an unexpected reward (Hollerman & Schultz, 1997; Schultz, 1992), and the presence of this dopamine is widely thought to strengthen recently active synapses (which presumably are responsible for the animal obtaining the reward; e.g., Arbuthnott, Ingham, & Wickens, 2000; Calabresi, Pisani, Centonze, & Bernardi, 1996). In reward-mediated learning, it is essential to strengthen those (and only those) synapses that actively participated in the response that elicited the reward. Because there is necessarily some delay between response and reward delivery, this means, therefore, that some trace must be maintained that signals which synapses were recently active. In the case of the medium spiny cells in the caudate nucleus, the morphology of the dendritic spines allows this trace to exist for several seconds after the response is initiated (Gamble & Koch, 1987; MacDermott, Mayer, Westbrook, Smith, & Barker, 1986). If the reward is delayed by more than this amount, then the ensuing dopamine release will strengthen inappropriate synapses and learning will be adversely affected.

Experiment 1

Method

Observers and design. One hundred twenty observers were solicited from the University of Texas community and received course credit for participation. Ten observers participated in each of 12 experimental conditions constructed from the factorial combination of 2 (feedback conditions: delay vs. immediate) × 3 (feedback delay or ITI durations: 10 s vs. 5 s vs. 2.5 s) × 2 (category structures: rule-based vs. information-integration). No observer participated in more than one experimental condition. All observers reported 20/20 vision or vision corrected to 20/20. Each observer completed one session of approximately 60-min duration.

Stimuli and stimulus generation. The experiment used the randomization technique introduced by Ashby and Gott (1988). The rule-based and information-integration category structures are displayed in Figures 2A and 2B, respectively. The stimuli for the rule-based categories were generated by sampling randomly from two bivariate normal distributions. The stimuli for the information-integration categories were generated by rotating the rule-based stimuli clockwise by 45°. Thus, optimal accuracy, within-category scatter, and category coherence are identical in both conditions. Each category distribution is specified by a mean and a variance on each dimension, and by a covariance between dimensions. For both category structures it was always the case that the covariance matrix for Category A was identical to the covariance matrix for Category B. The categories differed only in the location of the means. Under these conditions the optimal decision bound will be linear. The optimal decision bound is displayed for each condition in Figure 1. The exact parameter values are listed in Table 1.

To begin, we generated 40 Category A stimuli and 40 Category B stimuli from the rule-based category structures by randomly sampling from the appropriate category distribution. The order of these 80 stimuli was randomized separately for each observer and was presented in Block 1. Three additional random orders were generated for each observer and were presented in Blocks 2–4. The stimuli for the information-integration category structures were generated using the algorithm described above and these were randomized separately for each block and observer.

The stimuli were computer generated and displayed on a 21-in. monitor with 1360 × 1024 resolution in a dimly lit room. Each Gabor patch was generated using MATLAB routines from Brainard's (1997) Psychophysics Toolbox. Each random sample (x_1, x_2) was converted to a stimulus by

Table 1
Category Distribution Parameters for Experiments 1 and 2

	Category A					Category B				
Condition	μ_f	μ_o	σ_f^2	σ_o^2	$\mathrm{cov}_{f,o}$	μ_f	μ_o	σ_f^2	σ_o^2	$\mathrm{cov}_{f,o}$
Experiment 1										
II	272	153	4,538	4,538	4,463	327	97	4,538	4,538	4,463
RB	260	125	75	9,000	0	340	125	75	9,000	0
Experiment 2										
RB-4.6	280	125	75	9,000	0	320	125	75	9,000	0
RB-3.5	285	125	75	9,000	0	315	125	75	9,000	0

Note. II = information integration; RB = rule based.

deriving the frequency, $f = .25 + (x_1/50)$, and orientation, $o = x_2\ (\pi/500)$. For example, the Category A mean for the rule-based category structure was converted to a Gabor pattern with frequency $f = .25 + (260/50) = 5.45$ cycles/degree and orientation, $o = 125\ (\pi/500) = 0.785$ radians counterclockwise from horizontal. The scaling factors were chosen in an attempt to equate the salience of frequency and orientation.

Procedure. Each observer was tested individually in a dimly lit room. The observers were informed that there were two categories and that each category was equally likely. They were informed that perfect performance was possible and were instructed to learn about the two categories. They were told to be as accurate as possible and not to worry about speed of responding. The procedure for a typical trial in the four feedback conditions was as follows:

10-s delay: stimulus—response terminates display—10 s. Mask—750 ms feedback—500 ms blank screen ITI.

5-s delay: stimulus—response terminates display—5 s. Mask—750 ms feedback—500 ms blank screen ITI.

2.5-s delay: stimulus—response terminates display—2.5 s. Mask—750 ms feedback—500 ms blank screen ITI.

Immediate, 10-s ITI: stimulus—response terminates display—500 ms. Mask—750 ms feedback—10 s blank screen ITI.

Immediate, 5-s ITI: stimulus—response terminates display—500 ms. Mask—750 ms feedback—5 s blank screen ITI.

Immediate, 2.5-s ITI: stimulus—response terminates display—500 ms. Mask—750 ms feedback—2.5 s blank screen ITI.

The mask was a Gabor pattern that subtended approximately 11° of visual angle and was of a random frequency and orientation from within the range of stimulus values. Notice that the 5-s delay and immediate/5-s ITI conditions are identical except that the duration of the mask and ITI are reversed with a similar relationship holding between the 2.5-s delay and immediate/2.5-s ITI conditions, and between the 10-s delay and immediate/10-s ITI conditions. This was necessary to equate the trial-to-trial timing across immediate and delayed feedback conditions.

Results and Theoretical Analyses

Analyses were performed separately on each block of data. In the first section we analyze the accuracy rates using an analysis of variance (ANOVA). In the second section we introduce the model-based analyses.

ANOVA Results

A 2 (feedback condition: delay vs. immediate) × 3 (feedback delay/ITI duration: 10 s vs. 5 s vs. 2.5 s) × 2 (category structure: rule-based vs. information-integration) × 4 (block) mixed-design ANOVA was conducted on the accuracy rates. The accuracy rates averaged across observers are presented in Table 2. The two most important results were a significant Feedback Condition × Category Structure × Block three-way interaction, $F(3, 324) = 5.29$,

Table 2
Averaged Accuracy Rates and Standard Errors of the Means for the Feedback Delay and Category Structure Conditions by Block From Experiment 1

	Block								
	1		2		3		4		
Feedback condition	*M*	*SE*	*M*	*SE*	*M*	*SE*	*M*	*SE*	Average
II category structure									
Delay									
10 s	59.5	2.6	61.4	3.3	58.5	3.7	59.3	2.0	59.7
5 s	55.3	3.4	50.3	1.5	54.4	2.9	58.0	3.2	54.5
2.5 s	59.1	3.3	56.1	3.4	57.4	4.2	59.0	4.5	57.9
Immediate									
10 s	61.3	3.8	68.3	4.5	71.0	4.0	77.3	2.7	69.4
5 s	59.1	3.8	69.1	4.0	71.4	4.9	73.4	3.9	68.3
2.5 s	61.1	3.0	68.3	2.7	67.0	2.9	73.0	3.8	67.3
RB category structure									
Delay									
10 s	83.8	5.1	92.1	2.3	95.4	1.3	97.5	0.6	92.2
5 s	75.3	4.9	92.9	3.4	96.8	0.6	96.1	1.1	90.3
2.5 s	73.1	4.8	92.9	1.8	97.6	0.9	96.5	0.7	90.0
Immediate									
10 s	75.9	5.1	90.4	3.3	94.0	1.5	94.8	1.1	88.8
5 s	74.8	5.0	98.0	0.6	96.3	1.0	97.4	0.7	91.6
2.5 s	83.4	2.2	97.6	0.9	97.6	0.9	98.4	0.3	94.3

Note. II = information integration; RB = rule based.

$p < .01$; and a significant Feedback Condition × Category Structure two-way interaction, $F(1, 108) = 13.71$, $p < .01$. The main effect of delay, $F(1, 108) = 17.74$, $p < .01$, was significant and suggested better performance in the immediate feedback conditions (80%) than in the delayed feedback conditions (74%). The main effect of category structure, $F(1, 108) = 415.50$, $p < .01$, was significant and suggested better performance in the rule-based task (91%) than in the information-integration task (63%). The main effect of block, $F(3, 324) = 79.59$, $p < .01$, was also significant and suggested consistent learning over blocks. Interestingly, the feedback delay/ITI duration manipulation was nonsignificant, $F(2, 108) = 0.39$, $p > .50$, and did not interact with any other factors.

The three-way interaction between feedback condition, category structure, and block is displayed graphically in Figure 3A. Several comments are in order: First, equivalent rule-based category learning was observed in the delayed and immediate feedback conditions. This was confirmed by a nonsignificant main effect of feedback type for the rule-based category structures, $F(1, 58) = 0.28$, $p > .50$, suggesting no effect of delay on rule-based category learning. Second, there was significant information-integration category learning under immediate feedback conditions, but essentially no learning under delayed feedback conditions. This was confirmed by a significant main effect of feedback type for the information-integration category structures, $F(1, 58) = 20.83$, $p < .01$, suggesting a strong detrimental effect of delayed feedback on information-integration category learning.

In line with the results from Ashby, Queller, & Berretty (1999) and Ashby et al. (2002) these accuracy-based analyses suggest that the timing of the corrective feedback affects performance in the information-integration categorization conditions, but not in the rule-based conditions, even when all structural aspects of the categories (optimal accuracy, within-category scatter, and category coherence) are invariant. Even so, it was the case that performance was consistently superior for the rule-based category structures. It is possible that the lack of a delay effect in the rule-based task might reflect a *ceiling effect* rather than a true performance dissociation. We address this ceiling effect interpretation in Experiment 2, but we first turn to the model-based analyses.

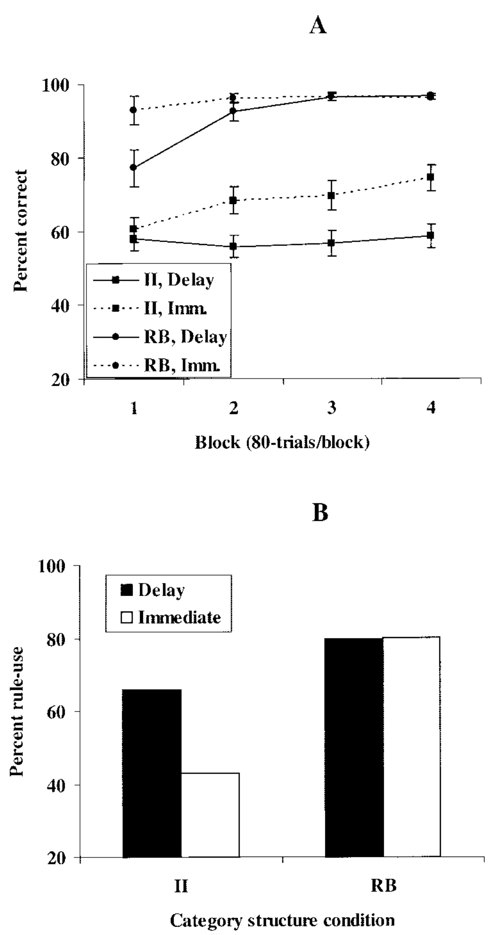

Figure 3. A: Percent correct for the Experiment 1 information-integration (II) and rule-based (RB) category structures under delayed and immediate feedback conditions for each block (standard error bars included). B: Percent of data sets for which a rule-based model provided the best account by feedback condition and category structure averaged across blocks. Imm. = Immediate.

Modeling Results

The accuracy-based analyses provide important information regarding overall performance, but they tell us nothing about the types of strategies that observers might use to solve these tasks. An understanding of strategy use, and how these strategies might be affected by the delay manipulation, is of critical importance to a complete understanding of category learning. One intriguing hypothesis suggested by the neuroscience literature and Ashby, Queller, & Berretty (1999) and the Ashby et al. (2002) studies is that observers in the information-integration condition will be forced to resort to rule-based categorization strategies when feedback is delayed because learning in the tail of the caudate nucleus will be impaired. To gain insight into these processes, we fit a number of different decision bound models (Ashby, 1992a; Maddox & Ashby, 1993) to the data separately by block and observer. Decision bound models are derived from general recognition theory (GRT; Ashby & Townsend, 1986), which is a multivariate generalization of signal-detection theory (e.g., Green & Swets, 1966). The fundamental assumption of GRT is that there is trial-by-trial variability in the perceptual information obtained from every stimulus, no matter what the viewing conditions (Ashby & Lee, 1993). On each trial, it is assumed that the percept can be represented as a point in a multidimensional psychological space. Decision bound theory assumes each observer partitions the perceptual space into response regions by constructing a decision bound. On each trial, the observer determines which region the percept is in and then emits the associated response. Despite this deterministic-decision rule, decision bound models predict probabilistic responding because of trial-by-trial perceptual and criterial noise.

Two different types of decision bound models were fit to each observer's responses (see Ashby, 1992a; Maddox & Ashby, 1993, for a more formal treatment of these models). One type is compatible with the assumption that observers used an explicit rule-based strategy and one type assumes an information-integration strategy. Even so, it is important to note that these models make no detailed process assumptions in the sense that a number of different process accounts are compatible with each of the models (e.g., Ashby, 1992a; Ashby & Waldron, 1999). For example, if an information-integration model fits significantly better than a rule-based model, then we can be reasonably confident that observers did not use a rule-based strategy, but we learn little about which information-integration strategy might have been used (e.g., decision bound, exemplar, or prototype interpretations would all be compatible with such results). In contrast, if a rule-based model fits significantly better than the information-integration models, then we gain confidence that observers used a rule-based strategy, but we cannot rule out all information-integration strategies because some of these can mimic rule-based responding. In summary, the modeling described in this section provides a powerful vehicle by which to test hypotheses about the decision strategies used by observers, but it has little to say about psychological processes. The following models were fit to each observer's responses (see Ashby, 1992a; Maddox & Ashby, 1993, for a more formal treatment of these models).

Rule-based models. Two models were compatible with the assumption that observers used an explicit rule-based strategy.

The *unidimensional model* assumes observers set a criterion on a single perceptual dimension and then make an explicit decision about the level of the stimulus on that dimension (Ashby & Gott, 1988; Shaw, 1982). For example, in the present experiments, observers might use the rule: "Respond 'A' if the spatial frequency is low and 'B' if it is high." A different version of this model assumes observers attend selectively to orientation, rather than to frequency. The unidimensional models have two free parameters: a decision criterion on the relevant perceptual dimension and the variance of internal (perceptual and criterial) noise (i.e., σ^2). In the unidimensional conditions, a special case of the unidimensional model assumes observers use the unidimensional decision bound that maximizes accuracy (i.e., the vertical bound shown in Figure 2A). This special case has only one free parameter (i.e., noise variance).

The *conjunction model* assumes observers use a conjunction rule in which they make separate decisions about the levels on the two dimensions and then select a response based on the outcome of these two decisions. Two conjunction rules were examined:

1. "Respond 'A' if spatial frequency is low and orientation is large, otherwise respond 'B,'" and

2. "Respond 'B' if spatial frequency is high and orientation is small, otherwise respond 'A'."

Both rules partition the perceptual space into four regions. The first assigns one to Category A and three to Category B, and the second assigns three to Category A and one to Category B. As mentioned in the introduction, such a strategy is rule-based because it is easy to describe verbally, and it does not require perceptual integration of spatial frequency and orientation. Conjunction models have three parameters (a criterion on each dimension, and σ^2).

Information–Integration models. The *general linear classifier* (GLC) assumes that the decision bound between each pair of categories is linear. This produces an information-integration decision strategy because it requires linear integration of perceived frequency and orientation. The GLC has three parameters (slope and intercept of the linear bound and σ^2). In the diagonal conditions, a special case of the GLC assumes observers use the linear bound that maximizes accuracy (i.e., the diagonal bound shown in Figure 2B). This model has only one free parameter (noise variance).

Model fits. Each of these models was fit separately to the data from each of the four blocks of trials for every observer. The model parameters were estimated using maximum likelihood (Ashby, 1992b; T. D. Wickens, 1982), and the goodness-of-fit statistic was $AIC = 2r - 2lnL$, 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 A information criterion (AIC; Akaike, 1974) 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 each of 480 data sets [2 (feedback conditions) × 3 (feedback delay/ITI durations) × 2 (category structures) × 10 (observers) × 4 (blocks)], we determined which model type (i.e., rule-based or information-integration) provided the best account of the data. The percentages of data sets for which a rule-based model provided the best account by feedback condition, feedback delay/ITI duration, category structure, and block are presented in Table 3. The percentages of data sets for which a rule-based model provided the best account by feedback condition and category structure averaged over feedback delay, ITI duration, and block are displayed graphically in Figure 3B.

Table 3
Percentage of Data Sets (by Feedback Condition, Category Structure, Feedback Delay/Intertrial Interval Duration, and Block) Best Fit by a Rule-Based Model From Experiment 1

	Block				
Feedback condition	1	2	3	4	Average
II category structure					
Delay					
10 s	90.0	60.0	60.0	80.0	72.5
5 s	30.0	70.0	60.0	55.6	53.9
2.5 s	70.0	60.0	60.0	70.0	65.0
Immediate					
10 s	50.0	30.0	66.7	33.3	45.0
5 s	33.3	30.0	44.4	40.0	36.9
2.5 s	70.0	40.0	50.0	30.0	47.5
RB category structure					
Delay					
10 s	80.0	70.0	70.0	70.0	72.5
5 s	100.0	90.0	80.0	60.0	82.5
2.5 s	90.0	90.0	80.0	77.8	84.4
Immediate					
10 s	80.0	60.0	77.8	77.8	73.9
5 s	80.0	90.0	90.0	70.0	82.5
2.5 s	90.0	88.9	90.0	66.7	83.9

Note. II = information integration; RB = rule based.

Two findings are of interest: First, the delay manipulation had no effect on the distribution of best fitting models in the rule-based condition, $F(1, 18) = 0.004$, $p > .50$. During the final block of training, approximately 70% of the data sets for both the immediate and delay conditions were best fit by a rule-based model. Second, the delay manipulation had a large effect on the distribution of best fitting models in the information-integration categorization condition. Specifically, the use of a rule-based strategy was larger in the delayed than in the immediate feedback conditions, $F(1, 18) = 13.15$, $p < .01$. During the final block of training, approximately 69% of the data sets in the delayed feedback condition were best fit by a rule-based model, whereas only 34% of the data sets in the immediate feedback condition were best fit by a rule-based model.

Discussion

Taken together, the model-based results converge with those from the accuracy-based analyses in suggesting a large effect of delayed feedback on information-integration category learning, and little if any effect of delayed feedback on rule-based category learning. In addition, the model-based analyses suggest that in the rule-based condition, the delayed feedback had little effect on the type of decision strategy that observers used, but in the information-integration condition delaying the feedback caused many observers to adopt qualitatively different strategies than they would use when feedback was immediate (i.e., rule-based strategies rather than an information-integration strategy).

Experiment 2

Although the data from Experiment 1 converge with previous research (Ashby et al., 2002; Ashby, Queller, & Berretty, 1999) and with the neuroscience literature, one weakness is that performance is quite good in the rule-based category-learning task, and thus might reflect a ceiling effect. A stronger claim could be made regarding a dissociation between rule-based and information-integration category learning and the effects of delayed feedback if the same pattern of delayed feedback results held when immediate feedback performance was equated across rule-based and information-integration conditions. To address this possibility, we collected data from two additional rule-based category structures in which the category mean separation was reduced. In one condition, category discriminability, $d' = 4.6$, and in the other $d' = 3.5$. We refer to these as the RB-4.6 and RB-3.5 conditions. Most important, as in Experiment 1, the categories did not overlap and so perfect performance was possible. Observers were run in the 5-s delay and immediate feedback conditions.

Method

Observers and design. Forty-eight observers were solicited from the University of Texas community and received course credit for participation. Twelve observers participated in each of the four experimental conditions. All other aspects of the observers and design were identical to those in Experiment 1.

Stimuli and procedures. The stimuli, stimulus generation, and experimental procedures were identical to those in Experiment 1. Only the parameters of the category distributions were changed (i.e., shown in Table 1).

Results and Theoretical Analyses

To facilitate a comparison of performance with the information-integration condition from Experiment 1, we included the 5-s delay and immediate feedback information-integration data from Experiment 1 in our analysis.

ANOVA Results

A 2 (feedback condition: delay vs. immediate) × 3 (category structure: information-integration vs. RB-4.6 vs. RB-3.5) × 4 (block) mixed-design ANOVA was performed on the accuracy rates. The accuracy rates are displayed in Figure 4A. The most important finding was a Feedback Condition × Category Structure interaction, $F(2, 62) = 3.15$, $p = .05$. To isolate the locus of this interaction, we conducted ANOVAs separately on the data from

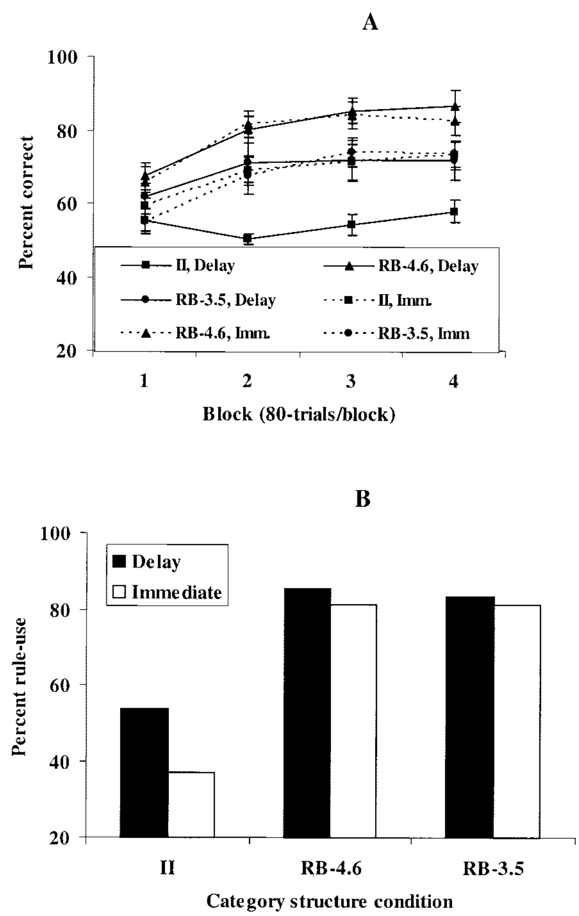

Figure 4. A: Percent correct for the Experiment 2 rule-based (RB) category structure and Experiment 1 information-integration (II) category structure under delayed and immediate feedback conditions for each block (standard error bars included). B: Percent of data sets for which a rule-based model provided the best account by feedback condition and category structure averaged across blocks. Imm. = Immediate.

the three category structures. The results were clear. There was a nonsignificant feedback effect for both the RB-3.5, $F(1, 22) = 0.09$, $p > .50$, and the RB-4.6 conditions, $F(1, 22) = 0.10$, $p > .50$, whereas the feedback effect was highly significant for the information-integration condition, $F(1, 18) = 10.16$, $p < .01$. Two additional results should be discussed: First, performance in the RB-4.6 condition reached asymptote at 86% (delay) and 83% (immediate), well below the possible 100%, and well below the asymptotes of 97% for the rule-based categories (delayed and immediate feedback) reached in Experiment 1. Second, performance was indistinguishable in the immediate feedback RB-3.5 condition and in the immediate feedback information-integration condition, $F(1, 20) = 0.26$, $p > .50$, whereas performance in the delayed feedback information-integration condition was much worse than in the delayed feedback RB-3.5 condition, $F(1, 20) = 7.13$, $p < .50$. Taken together, these results provide strong evidence against a ceiling effect explanation of the results from Experiment 1, and instead provide strong support for the predicted dissociation between rule-based and information-integration category learning under delayed and immediate feedback conditions.

Modeling Results

The models applied in Experiment 1 were applied to the Experiment 2 data. The percentage of data sets for which one of the rule-based models provided the best account by delay condition and category structure are displayed in Figure 4B. In line with the results from Experiment 1, the use of a rule-based strategy to learn the rule-based category structure was high (80%), and was unaffected by the delay manipulation and category discriminability.

General Discussion

This article reported the results from two experiments that examined the differential effects of delayed versus immediate feedback on the learning of rule-based and information-integration category structures that were equated on several measures of complexity (including optimal accuracy, within-category scatter, and category coherence). For the high and low discriminability rule-based category structures (Experiments 1 and 2, respectively), the delay manipulation had no effect on the accuracy of responding or on the type of decision strategy used by observers. For the information-integration category structure, a delay in corrective feedback led to a significant decrease in the accuracy of responding relative to immediate feedback situations and to an increase in the use of rule-based response strategies. This dissociation between rule-based and information-integration category learning under delayed feedback conditions held even when immediate feedback performance was equated across category structures.

Single Versus Multiple Systems

These findings add to the body of literature in support of multiple memory and category learning systems (e.g., Ashby & Ell, 2001, 2002b; Filoteo, Maddox, & Davis, 2001a, 2001b; Knowlton & Squire, 1993; Maddox & Filoteo, 2001; Pickering, 1997) and suggest that rule-based and information-integration category learning are mediated by separate systems. They are also consistent with the hypothesis that learning in rule-based tasks is dominated by an explicit system that uses working memory and executive attention. This system appears to learn through a conscious process of hypothesis generation and testing. Given this hypothesis, manipulations of the timing and placement of corrective feedback should, and do, have little effect on rule-based learning. In fact, as suggested by the Ashby, Queller, & Berretty (1999) results, people can learn some rule-based categories with *no* feedback of any kind.

In contrast, it has been hypothesized that learning in information-integration tasks is dominated by an implicit procedural-learning-based system, which is mediated largely within the tail of the caudate nucleus. It has been proposed that a dopamine-mediated reward signal is critical for learning in this system. The idea is that an unexpected reward causes dopamine to be released from the substantia nigra into the tail of the caudate nucleus and that the presence of this dopamine strengthens recently active synapses (e.g., Schultz, 1992; J. Wickens, 1993). Because this dopamine-mediated learning requires a close temporal correspondence between stimulus presentation and feedback, delayed feedback should, and does, adversely affect performance.

One might argue that the performance decrement observed for information-integration category learning under delayed feedback conditions is due to the fact that the feedback is being associated with the mask (a randomly generated Gabor patch), instead of the original training stimulus, and is not due to the delay in feedback. This interpretation would appear to predict poor learning in all experimental conditions, including the immediate feedback conditions because in all conditions the mask precedes the feedback. This pattern of results was not observed. Even if this association between mask and feedback occurred only in the delayed feedback conditions, it would still predict the same effect for information-integration and rule-based category learning, and so does not offer an alternative to the dual-systems view, but rather offers a different mechanism by which the dissociation emerges.

One possible argument against this dual-system interpretation takes the following form. First, although the rule-based and information-integration category structures were equated on several objective measures of difficulty (such as maximum attainable accuracy, within-category scatter, between-category separation, and category coherence; Fukunaga, 1990), no such manipulations can succeed because information-integration tasks are inherently more difficult than rule-based tasks. Second, category learning is mediated by a single system, so our observed dissociation is simply another example of the general principle that any factor interfering with learning will disrupt performance in difficult tasks more than in simple tasks.

There are at least two problems with this reasoning. First, one would need to offer an *a priori* reason why the information-integration task is more difficult. This is not as easy as it might first appear. For example, an ideal observer would perform equally on the two Experiment 1 category structures, even if susceptible to perceptual and criterial noise. So would any common cluster or discriminant analysis algorithm. In fact, pigeons have no more difficulty learning diagonal categories than unidimensional categories (Herbranson, Fremouw, & Shimp, 1999). Even more problematic is the fact that by all of these objective measures, the rule-based categories of Experiment 2 are more difficult than the information-integration categories of Experiment 1.

Second, and more important, this difficulty argument predicts that any experimental manipulation that disrupts category learning should have greater effects in information-integration tasks than in rule-based tasks. This prediction has been strongly disconfirmed in several recent studies. For example, Waldron and Ashby (2001) showed that a simple rule-based category-learning task (which required attending to a single stimulus dimension) was disrupted more by a simultaneous task that requires working memory and executive attention (a numerical Stroop task) than was a complex information-integration task (which required attending to three stimulus dimensions). If difficulty is the most important factor, then simultaneously performing a second task should interfere more strongly with learning the more difficult information-integration structures. Because Waldron and Ashby found the opposite pattern of results, it seems likely that factors other than difficulty are at work. In addition, Ashby et al. (2002) found that the same group of Parkinson's disease patients were much more impaired at rule-based category learning (one dimensional) than at information-integration category learning (three dimensional). If a single system mediates learning in these two types of categorization tasks, and if Parkinson's disease damages this system, then we would expect the more serious deficits to occur in the more difficult information-integration tasks.

It should be stressed that the dissociation reported in this article, between feedback delay and type of category structure, was predicted *a priori* from the dual-system model, and further that these predictions were parameter free. Since the difficulty hypothesis can be rejected, it is extremely difficult to see how a single-system model could predict our results in a similar *a priori* fashion. Of course it is more difficult to rule out the possibility that some single-system model might be able to account for our results by means of post hoc manipulation of its parameters. We would argue, however, that models making such post hoc accounts are less parsimonious than models making parameter-free *a priori* predictions, regardless of the number of systems assumed by the two models.

Magnitude of the Delay

Three temporal delays (2.5 s, 5 s, and 10 s) were included to provide some preliminary insights into the functional relationship between accuracy and the magnitude of the delay. Performance did not differ significantly across the 2.5-s, 5-s, and 10-s delay conditions in the information-integration tasks, suggesting that even the shortest 2.5-s duration was beyond that necessary to see significant synaptic efficacy change in the caudate. Future research should examine a number of additional delays in the 0 s to 2.5 s range to determine the functional relationship between accuracy and the magnitude of the delay.

Implications for Real-World Category Learning

The present results, along with those from Ashby, Queller, & Berretty (1999) and Ashby et al. (2002), have important implications for real-world categorization learning and training. When teaching rule-based categories that can be solved by applying an explicit (often verbalizable) rule, many training procedures will be adequate. However, when teaching information-integration categories that cannot be solved by applying verbalizable rules, care must be taken to ensure the most efficient training procedure. Our results suggest that it is critical that the feedback follows the presentation of the stimulus and the observer's response, and that it follow it in a timely manner. Although applicable in many domains, one in which most categories must be learned by means of information integration is medical diagnosis. For example, when training medical students to read X-rays, our results suggest that the most efficient method might be to have the students examine each X-ray, make a diagnosis, and receive immediate feedback.

In summary, the effect of immediate versus delayed feedback on rule-based and information-integration category learning was investigated. Feedback delay had no effect on the accuracy of responding or on the distribution of best fitting models in the rule-based category-learning task, but led to less accurate responding in the information-integration category-learning task and to an increase in the use of rule-based strategies to solve the information-integration task. These results provide support for a multiple-systems approach to category learning and argue against the validity of single-system approaches.

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.

Arbuthnott, G. W., Ingham, C. A., & Wickens, J. R. (2000). Dopamine and synaptic plasticity in the neostriatum. *Journal of Anatomy, 196,* 587–596.

Ashby, F. G. (1992a). Multidimensional models of categorization. In F. G. Ashby (Ed.), *Multidimensional Models of perception and cognition* (pp. 449–484). Hillsdale, NJ: Erlbaum.

Ashby, F. G. (1992b). Multivariate probability distributions. In F. G. Ashby (Ed.), *Multidimensional models of perception and cognition* (pp. 1–34). Hillsdale, 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. (2001). The neurobiological basis of category learning. *Trends in Cognitive Sciences, 5,* 204–210.

Ashby, F. G., & Ell, S. W. (2002a). Single versus multiple systems of category learning: Reply to Nosofsky and Kruschke (2001). *Psychonomic Bulletin & Review, 9,* 175–180.

Ashby, F. G., & Ell, S. W. (2002b). Single versus multiple systems of learning and memory. In J. Wixted & H. Pashler (Eds.), *Stevens' handbook of experimental psychology: Vol. 4. Methodology in experimental psychology* (3rd ed., pp. 655–692) New York: Wiley.

Ashby, F. G., Ell, S. W., & Waldron, E. M. (in press). Procedural learning in perceptual categorization. *Memory & Cognition.*

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., Isen, A. M., & Turken, A. U. (1999). A neuropsychological theory of positive affect and its influence on cognition. *Psychological Review, 106,* 529–550.

Ashby, F. G., & Lee, W. W. (1993). Perceptual variability as a fundamental axiom of perceptual science. In S. C. Masin (Ed.), *Foundations of perceptual theory* (pp. 369–399). Amsterdam: Elsevier.

Ashby, F. G., & Maddox, W. T. (1992). Complex decision rules in categorization: Contrasting novice and experienced performance. *Journal of Experimental Psychology: Human Perception and Performance, 18,* 50–71.

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. (1998). Stimulus categorization. In M. H. Birnbaum (Ed.), *Measurement, judgement, and decision making* (pp. 251–301). New York: Academic Press.

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,* 666–667.

Ashby, F. G., Noble, S., Filoteo, J. V., Waldron, E. M., & Ell, S. W. (2003). Category learning deficits in Parkinson's disease. *Neuropsychology, 17,* 115–124.

Ashby, F. G., Queller, S., & Berretty, P. T. (1999). On the dominance of unidimensional rules in unsupervised categorization. *Perception & Psychophysics, 61,* 1178–1199.

Ashby, F. G., & Townsend, J. T. (1986). Varieties of perceptual independence. *Psychological Review, 93,* 154–179.

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 categorization and identification. *Journal of Experimental Psychology: General, 130,* 77–96.

Brainard, D. H. (1997). Psychophysics software for use with MATLAB. *Spatial Vision, 10,* 433–436.

Calabresi, P., Pisani, A., Centonze, D., & Bernardi, G. (1996). Role of Ca^{2+} in striatal LTD and LTP. *Seminars in the Neurosciences, 8,* 321–328.

Eacott, M. J., & Gaffan, D. (1992). Inferotemporal-frontal disconnection: The uncinate fascicle and visual associative learning in monkeys. *European Journal of Neuroscience, 4,* 1320–1332.

Elliott, R., & Dolan, R. J. (1998). Activation of different anterior cingulate foci in association with hypothesis testing and response selection. *Neuroimage, 8,* 17–29.

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

Estes, W. K. (1994). *Classification and cognition.* Oxford, England: Oxford University Press.

Filoteo, J. V., Maddox, W. T., & Davis, J. D. (2001a). A possible role of the striatum in linear and nonlinear categorization rule 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.

Fukunaga, K. (1990). *Introduction to statistical pattern recognition.* New York: Academic Press.

Gaffan, D., & Eacott, M. J. (1995). Visual learning for an auditory secondary reinforcer by macaques is intact after uncinate fascicle section: Indirect evidence for the involvement of the corpus striatum. *European Journal of Neuroscience, 7,* 1866–1871.

Gaffan, D., & Harrison, S. (1987). Amygdalectomy and disconnection in visual learning for auditory secondary reinforcement by monkeys. *Journal of Neuroscience, 7,* 2285–2292.

Gamble, E., & Koch, C. (1987, June 5). The dynamics of free calcium in dendritic spines in response to repetitive synaptic input. *Science, 236,* 1311–1315.

Goldman-Rakic, P. S. (1987). Circuitry of the prefrontal cortex and the regulation of behavior by representational knowledge. In F. Plum & V. Mountcastle, Eds., *Handbook of Physiology Section 1: The Nervous System: Vol. 5, Parts 1 and 2; Higher functions of the brain.* (pp. 373–417), Oxford, England: Oxford University Press.

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

Green, D. M., & Swets, J. A. (1966). *Signal detection theory and psychophysics.* New York: Wiley.

Herbranson, W. T., Fremouw, T., & Shimp, C. P. (1999). The randomization procedure in the study of categorization of multidimensional stimuli by pigeons. *Journal of Experimental Psychology: Animal Behavior Processes, 25,* 113–134.

Hollerman, J. R., & Schultz, W. (1997). Dopamine neurons report an error in the temporal prediction of reward during learning. *Nature Neuroscience, 1,* 304–308.

Knowlton, B. J., Mangels, J. A., & Squire, L. R. (1996, September 6). A neostriatal habit learning system in humans. *Science, 273,* 1399–1402.

Knowlton, B. J., & Squire, L. R. (1993, December 10). The learning of categories: Parallel brain systems for item memory and category knowledge. *Science, 262,* 1747–1749.

Kruschke, J. K. (1992). ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review, 99,* 24–44.

MacDermott, A. B., Mayer, M. L., Westbrook, G. L., Smith, S. J., & Barker, J. L. (1986, May 29–June 4). NMDA-receptor activation increases cytoplasmic calcium concentration in cultured spinal cord neurones. *Nature, 321,* 519–522.

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

Maddox, W. T., & Filoteo, J. V. (2001). Striatal contribution to category learning: Quantitative modeling of simple linear and complex non-linear rule learning in patients with Parkinson's disease. *Journal of the International Neuropsychological Society, 7,* 710–727.

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 and Neural Biology, 61,* 260–270.

Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review, 85,* 207–238.

Nosofsky, R. M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General, 115,* 39–57.

Packard, M. G., Hirsh, R., & White, N. M. (1989). Differential effects of fornix and caudate nucleus lesions on two radial maze tasks: Evidence for multiple memory systems. *Journal of Neuroscience, 9,* 1465–1472.

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.

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

Posner, M. I., & Petersen, S. E. (1990). Attention systems in the human brain. *Annual Review of Neuroscience, 13,* 25–42.

Rao, S. M., Bobholz, J. A., Hammeke, T. A., Rosen, A. C., Woodley, S. J., Cunningham, J. M., et al. (1997). Functional MRI evidence for subcortical participation in conceptual reasoning skills. *Neuroreport, 8,* 1987–1993.

Robinson, A. L., Heaton, R. K., Lehman, R. A. W., & Stilson, D. W. (1980). The utility of the Wisconsin Card Sorting Test in detecting and localizing frontal lobe lesions. *Journal of Consulting and Clinical Psychology, 48,* 605–614.

Rolls, E. T. (1994). Neurophysiology and cognitive functions of the striatum. *Revue Neurologique, 150,* 648–660.

Schultz, W. (1992). Activity of dopamine neurons in the behaving primate. *Seminars in Neuroscience, 4,* 129–138.

Shaw, M. L. (1982). Attending to multiple sources of information: I. The integration of information in decision making. *Cognitive Psychology, 14,* 353–409.

Smith, E. E., & Medin, D. L. (1981). *Categories and concepts.* Cambridge, MA: Harvard University Press.

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

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

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

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.

Wilson, C. J. (1995). The contribution of cortical neurons to the firing pattern of striatal spiny neurons. In J. C. Houk, J. L. Davis, D. G. Beiser (Eds.), *Models of information processing in the basal ganglia* (pp. 29–50). Cambridge, MA: Bradford.

Winocur, G., & Eskes, G. (1998). Prefrontal cortex and caudate nucleus in conditional associative learning: Dissociated effects of selective brain lesions in rats. *Behavioral Neuroscience, 112,* 89–101.

Yamauchi, T., & Markman, A. B. (2000a). Inference using categories. *Journal of Experimental Psychology; Learning, Memory, and Cognition, 26,* 776–795.

Yamauchi, T., & Markman, A. B. (2000b). Learning categories composed of varying instances: The effect of classification, inference, and structural alignment. *Memory & Cognition, 28,* 64–78.

Received January 28, 2002
Revision received October 13, 2002
Accepted December 31, 2002 ■