

# Procedural learning in perceptual categorization

F. GREGORY ASHBY and SHAWN W. ELL  
*University of California, Santa Barbara, California*

and

ELLIOTT M. WALDRON  
*Sightward Inc., Seattle, Washington*

In two experiments, observers learned two types of category structures: those in which perfect accuracy could be achieved via some explicit rule-based strategy and those in which perfect accuracy required integrating information from separate perceptual dimensions at some predecisional stage. At the end of training, some observers were required to switch their hands on the response keys, whereas the assignment of categories to response keys was switched for other observers. With the rule-based category structures, neither change in response instructions interfered with categorization accuracy. However, with the information-integration structures, switching response key assignments interfered with categorization performance, but switching hands did not. These results are consistent with the hypothesis that abstract category labels are learned in rule-based categorization, whereas response positions are learned in information-integration categorization. The association to response positions also supports the hypothesis of a procedural-learning-based component to information integration categorization.

Categorization is a critical skill that every organism must possess in at least a rudimentary form, because it allows organisms to respond differently, for example, to nutrients and poisons and to predators and prey. There is much recent evidence that human category learning relies on multiple systems (e.g., Ashby, Alfonso-Reese, Turken, & Waldron, 1998; Ashby & Ell, 2001, 2002a, 2002b; Erickson & Kruschke, 1998; Pickering, 1997; Smith, Jonides, & Koeppel, 1996; Smith, Patalano, & Jonides, 1998; Waldron & Ashby, 2001). In all cases in which multiple systems have been proposed, it has been hypothesized that one system uses explicit (i.e., rule-based) reasoning and at least one other system involves some form of implicit learning.

Nevertheless, there is still much disagreement. First, the proposal that there are multiple category-learning systems is disputed. In particular, Nosofsky and his colleagues have argued that single-system models can account for many of the phenomena that have been used to support the notion of multiple systems (Nosofsky & Johansen, 2000; Nosofsky & Kruschke, 2002; Nosofsky & Zaki, 1998). Second, even among those researchers postulating separate explicit and implicit systems, there is disagreement about the nature of the implicit system. For example, Ashby et al. (1998) have proposed a procedural-learning-based implicit system (see also Ashby & Waldron, 1999; Ashby, Waldron, Lee, & Berkman, 2001; Waldron & Ashby, 2001). In contrast, several researchers have proposed that the implicit system is

exemplar memory based (Erickson & Kruschke, 1998; Pickering, 1997), and there have also been proposals that the perceptual representation memory system participates in implicit category learning (Ashby & Ell, 2001; Knowlton, Squire, et al., 1996; Reber, Stark, & Squire, 1998).

One distinctive feature of procedural learning, which sets it apart from exemplar-based learning and from perceptual learning, is its association with motor performance (e.g., Squire, Knowlton, & Musen, 1993; Willingham, Nissen, & Bullemer, 1989). Whereas there is no reason to expect that categories learned via an exemplar-based or a perceptual-learning-based process would have any close association with a motor response, the procedural-learning literature suggests that categories acquired via procedural learning should be more closely associated with motor responses than are categories acquired via other forms of learning. This article reports results supporting this prediction. In particular, we describe results from two experiments in which participants were required to learn to assign stimuli to one of two contrasting categories. In brief, our results support the hypothesis that abstract category labels were learned when perfect accuracy could be achieved via some explicit reasoning strategy but that response positions were learned when perfect accuracy required integrating information from separate perceptual dimensions at some predecisional stage. As such, to our knowledge, these are the first results to provide direct evidence of a procedural-learning-based component of perceptual categorization.

## Two Different Category-Learning Tasks

Much of the evidence for multiple category-learning systems comes from two different kinds of tasks (Ashby

---

This research was supported in part by National Science Foundation Grant BCS99-75037 and by Public Health Service Grant MH3760. Correspondence concerning this article should be addressed to F. G. Ashby, Department of Psychology, University of California, Santa Barbara, CA 93106 (e-mail: ashby@psych.ucsb.edu).

& Ell, 2001). *Rule-based category-learning tasks* are those in which the categories can be learned via some explicit reasoning process. Frequently, the rule that maximizes accuracy (i.e., the optimal strategy) is easy to describe verbally (Ashby et al., 1998). In the most common applications, only one stimulus dimension is relevant, and the observer's task is to discover this relevant dimension and then to map the different dimensional values to the relevant categories.

*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 strategy in information-integration tasks is difficult or impossible to describe verbally (Ashby et al., 1998).

The two category structures used in the present study are described in Figure 1. Every stimulus was a single line that varied across trials in length and orientation. Each symbol in Figure 1 denotes the length and orientation of a single stimulus. Category A exemplars are denoted by plus signs, and Category B exemplars are denoted by circles. In each condition, there were two distinct categories that did not overlap, so perfect accuracy was always possible. Also shown in Figure 1 are the decision bounds that maximize categorization accuracy. In Figure 1A, the optimal bound requires observers to attend to line length and ignore orientation, so we call these unidimensional categories. With the diagonal categories of Figure 1B, 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.<sup>1</sup>

The unidimensional condition is a rule-based task because there is a simple explicit rule that separates the contrasting categories. Specifically, the vertical bound in Figure 1A corresponds to the rule “respond A if the line is short and B if it is long.” In contrast, the diagonal condition is an information-integration task. In this condition, perfect accuracy requires integrating length and orientation information, and there is no simple verbal description of the optimal decision bound.<sup>2</sup>

Note that we use the word *rule* more narrowly than is common in the psychological literature, where it is often used to refer to any strategy from an explicit reasoning

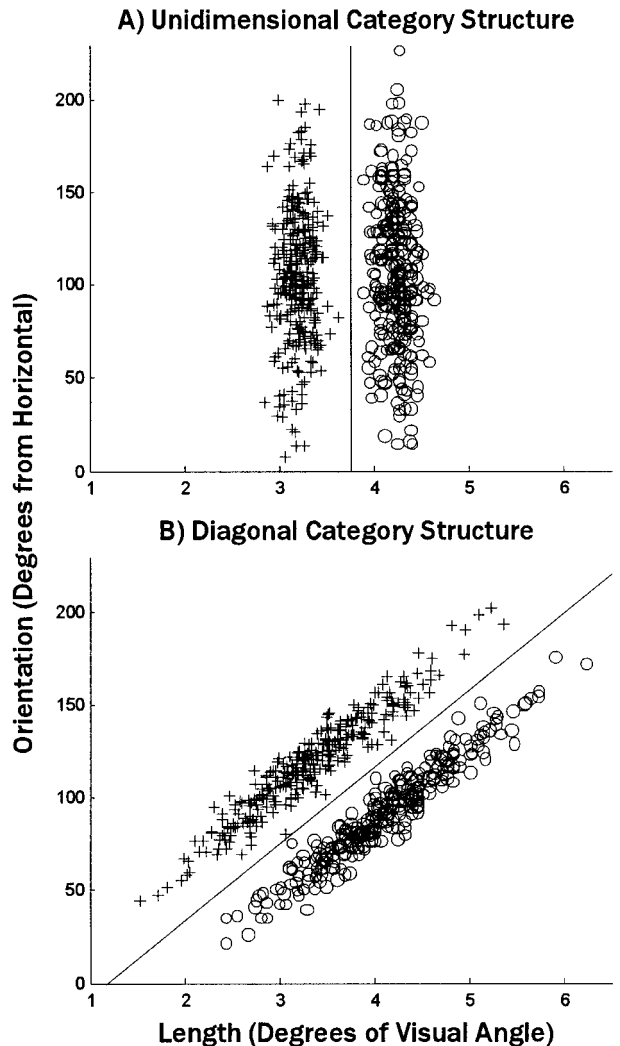


Figure 1. Category structures used in the unidimensional (A) and diagonal (B) conditions.

process to an algorithm that can be expressed formally. For example, the optimal decision strategy in the diagonal condition shown in Figure 1B could be said to use the following “rule”: “Respond A if orientation is greater than length; otherwise, respond B.” However, what does it mean to compare orientation and length? For example, which is greater, 43° of orientation or 2 cm of length? None of the many people we have run in conditions of this type (over the past decade) have described their decision strategy in this form, even when their data were well described by this “rule.” Thus, it is important to remember that we define *rule-based strategy* narrowly to refer specifically to an explicit reasoning process. Note that according to this criterion, there is no limit on the complexity of the optimal rule in rule-based tasks. However, as the complexity of the optimal rule increases, its salience decreases, and it becomes less likely that observers will learn the associated categories through an

explicit reasoning process. In fact, Alfonso-Reese (1997) found that even simple conjunction rules have far lower salience than do unidimensional rules. This does not mean that people cannot learn conjunction rules, but only that they are unlikely to experiment with such rules unless feedback compels them in this direction. This discussion should make it clear that the boundary on what constitutes a rule-based task is fuzzy. Tasks in which the optimal rule is unidimensional are unambiguously rule based (at least with separable stimulus dimensions), and tasks in which the optimal rule is significantly more complex than a conjunction rule are almost never rule based. In between, the classification is not so clear-cut. For this reason, the rule-based tasks used in this article all have a unidimensional optimal rule.

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

For each type of category structure shown in Figure 1, every observer in Experiment 1 completed 12 blocks of

50 trials each. The first 10 blocks were training trials. On each of these trials, the observers were given 5 sec to make their categorization responses, and they were given feedback about their accuracy after every response. The last 2 blocks of the session measured transfer under different response instructions. In all transfer conditions, the observers were required to respond within 1.5 sec. Three types of instructions were used: control, hand switch, and button switch. These are illustrated in Figure 2. In the control condition, the only difference from training was the reduced response deadline. In the hand-switch condition, the observers began the training blocks with their hands crossed on the response buttons (i.e., left hand on the right button and right hand on the left button). During the last 2 transfer blocks, the observers were instructed to “uncross” their hands (i.e., the right hand on the right button and the left hand on the left button). The button-switch condition was identical to the control condition, except that the buttons used to make the category responses were reversed during the transfer blocks. The observers were informed of the switch immediately before the switch occurred. Thus, in the button-switch condition, both the response locations and the motor responses were reversed, whereas in the hand-switch condition, only the motor responses were reversed. The response deadline was added to all transfer conditions to ensure that the observers could not overcome an interference from, say, a button switch by simply inhibiting their initial response. If there is a procedural-learning component in either rule-based or information-integration category learning, a hand or a button switch could cause an interference at the time of transfer.

Results from the procedural-learning literature allow us to sharpen these predictions considerably. The most




Figure 2. Design of the three conditions in Experiment 1. “A” and “B” denote the locations of the response keys assigned to Categories A and B, respectively. “L” and “R” denote the left and the right hands, respectively.

popular experimental paradigm in recent investigations of procedural learning uses Nissen and Bullemer's (1987) serial reaction time (SRT) task. In the prototypical application, one of  $n$  stimuli is presented on each trial, and each stimulus is associated with a different response key. The observer's task is to depress the correct key as quickly as possible. The greatest improvements in reaction time occur when the stimulus sequence has some repeating structure, even when observers display no explicit knowledge of the sequential structure (a defining characteristic of procedural learning) and even when observers suffer from medial temporal lobe amnesia (Nissen & Bullemer, 1987; Reber & Squire, 1994). Although there is not universal agreement, the prevailing view seems to be that SRT learning involves both implicit and explicit components (for a review, see Curran, 1995). Willingham (1998) has hypothesized that in SRT learning, an explicit system learns stimulus positions, whereas an implicit system learns response positions. In other words, the explicit system learns the spatial location of the next stimulus in the sequence, whereas the implicit system learns the position of the next motor response to be executed. As impressive evidence in support of this hypothesis, Willingham, Wells, Farrell, and Stemwedel (2000) have shown that changing response locations interferes with implicit SRT learning, even when the sequence of stimulus positions is unchanged. They also showed that implicit SRT learning is maintained if the response locations remain unchanged, even if the sequence of finger movements changes (i.e., in a hand-switch condition). Willingham et al.'s (2000) results provide a strong test of the controversial hypothesis that information-integration category learning is mediated by a procedural-learning-based categorization system. If this is the case, then with the diagonal categories, an interference should occur in the button-switch condition, but not in the hand-switch condition. Furthermore, if learning in the rule-based conditions is mediated by an explicit reasoning system, no selective interference should be observed.

## EXPERIMENT 1

### Method

**Observers and Design.** In this experiment, we used a  $2 \times 3$  factorial design, with two category types (unidimensional/diagonal) and three types of response instructions (control/hand-switch/button-switch). One hundred sixteen observers were solicited from the

University of California, Santa Barbara community and received course credit for participation. The number of observers who participated in each of six experimental conditions was as follows: unidimensional/control, 16; unidimensional/hand-switch, 13; unidimensional/button-switch, 13; diagonal/control, 22; diagonal/hand-switch, 21; and diagonal/button-switch, 38. No observer participated in more than one experimental condition. All the observers reported 20/20 vision or vision corrected to 20/20. Each observer completed one session of approximately 45-min duration. The criterion for learning was defined as 70% correct during the block preceding the change in response instructions (i.e., Block 10). The data from the observers who did not meet this criterion were excluded from all the subsequent analyses. This criterion resulted in the exclusion of 1 observer each from the control and the hand-switch conditions with the unidimensional category structures and 1, 2, and 3 observers, respectively, from the control, hand-switch, and button-switch conditions with the diagonal category structures.

**Stimuli and stimulus generation.** In the experiment, we used the randomization technique introduced by Ashby and Gott (1988), in which each category was defined as a bivariate normal distribution. Each category distribution was specified by a mean and a variance on each dimension and by a covariance between dimensions. The exact parameter values are displayed in Table 1. On each trial, a random sample  $(x, y)$  was drawn from the Category A or B distribution, and these values were used to construct a stimulus line of length  $x$  pixels and orientation  $y \times (\pi / 600)$  radians. The scale factor  $(\pi / 600)$  was chosen from past research to make changes in perceived length about equally discriminable to changes in perceived orientation. The complete set of stimuli used in the experiment are shown in Figure 1.<sup>3</sup> Each symbol in Figure 1 denotes the length and orientation of a single line stimulus. Category A stimuli are denoted by the + signs, and Category B stimuli are denoted by the circles. For both the unidimensional and the diagonal conditions, the two categories had identical length and orientation variances and an identical length-orientation covariance. Under these conditions, the optimal decision bounds are the vertical and diagonal lines shown in Figure 1.

First, we generated the stimuli for the unidimensional conditions (Figure 1A) by randomly sampling 300 stimuli from the Category A distribution and 300 stimuli from the Category B distribution. A linear transformation was then performed to ensure that the sample and the population means, variances, and covariances were identical. The order of the resulting 600 stimuli was then randomized separately for each observer and divided into 12 blocks of 50 trials each. The stimuli for the diagonal conditions (Figure 1B) were generated by rotating the unidimensional categories 45° clockwise around a central point located at 350 pixels in length (3.75° of visual angle) and 350 orientation units (i.e., 105° from horizontal). Thus, the optimal strategy in the diagonal condition required integrating information from both the length and the orientation of the stimulus. For these category structures, the best unidimensional rule would achieve approximately 70% accuracy.

The stimuli were computer generated and displayed on a 15-in. CRT with 832 × 624 pixel resolution in a dimly lit room. Each line was presented in white on a black background and subtended a visual angle between 1.0° and 6.5°. The midpoint of the line was fixed at the

**Table 1**  
**Parameter Values Used in the Generation of the Unidimensional and Diagonal Category Structures**

Category Structure	Means				Variances		Covariance (Both Categories)
	Length		Orientation		(Both Categories)		
	Category A	Category B	Category A	Category B	Length	Orientation	
Unidimensional	280	420	350	350	330	14,000	0
Diagonal	300	400	400	300	8,000	8,000	7,800

center of the monitor. To minimize line jaggedness, Alfonso-Reese's (1997) anti-aliasing routine, developed for use with Brainard's (1997) Psychophysics Toolbox, was applied.

**Procedure.** The observers were run individually in a dimly lit testing room. They were told that there were two equally likely categories and that perfect performance was possible.

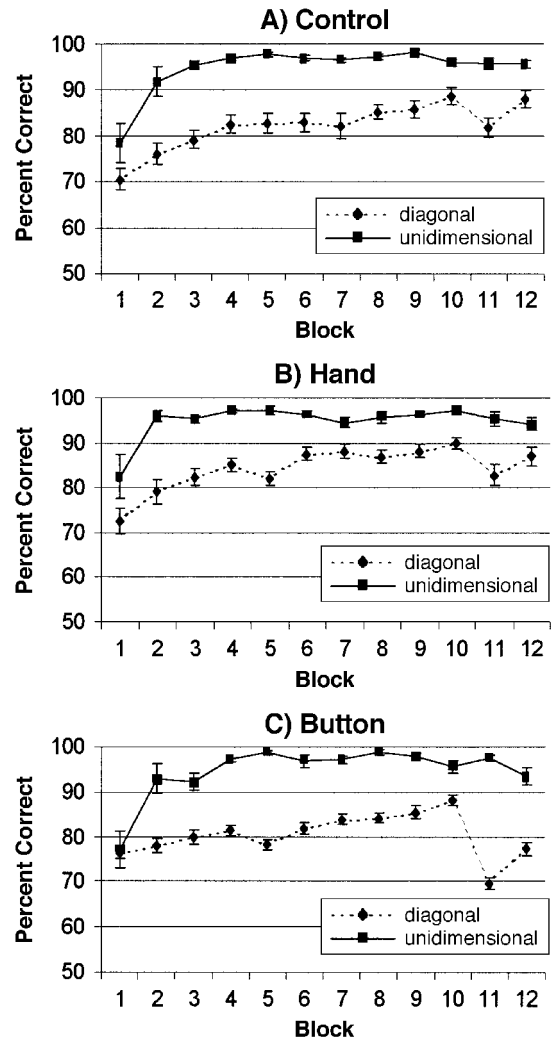
There were three experimental conditions: control, hand switch, and button switch. In all the conditions, the observers depressed the two response keys with their index fingers, and trial-by-trial feedback was provided. In the control condition, the observers were given 500 trials (10 blocks of 50 trials) of training, with 5 sec to make their response. They were then given 100 trials (2 blocks of 50 trials), with 1.5 sec to respond. If a response was not given in that time period, the observer was prompted to speed up his or her response, and that trial was discarded. A brief (1-sec) high-pitched tone (500 Hz) was presented if the response was correct, and a low-pitched tone (200 Hz) was presented if the response was incorrect. In addition, feedback was given at the end of each block of 50 trials regarding the participant's accuracy during that block. The response-stimulus interval was 1 sec.

In the hand-switch condition, the observers began the first 500 trials with their hands crossed on the response buttons (i.e., left hand on the right button and right hand on the left button). For the last 100 trials, the observers were instructed to "uncross" their hands (i.e., the right hand on the right button and the left hand on the left button). The button-switch condition was identical to the control condition, except that the buttons used to make the category response were reversed for the last 100 trials. Thus, in the button-switch condition, both the response location and the motor response were reversed, whereas in the hand-switch condition, only the motor responses were reversed. The observers were instructed at the beginning of the experiment that the response instructions would change following Block 10 and that they would be informed of the switch immediately before the switch occurred. The words "respond quickly," "respond quickly and switch your hands now," and "respond quickly and the response buttons will now be switched" were displayed on the monitor immediately prior to the change in the response instructions in the control, hand-switch, and button-switch conditions, respectively.

## Results

**Accuracy-based analyses.** The learning curves for each of the three experimental conditions are shown in Figure 3. Several important trends are evident. First, as was expected, accuracy was consistently higher with the unidimensional structures than with the diagonal structures. Second, learning was abrupt with the unidimensional structures but was gradual with the diagonal structures. Third, during the critical Transfer Blocks 11 and 12, there was no decrement in performance with the unidimensional structures in any of the response conditions. Fourth, with the diagonal structures, accuracy declined in Block 11 in all three response conditions, and this loss was greatest in the button-switch condition. Finally, during the last block, there was an accuracy recovery, in all three response conditions, of about the same magnitude. However, in the control and the hand-switch conditions, this recovery completely erased the losses during Block 11, but in the button-switch condition, accuracy was still much lower in Block 12 than at the end of the training period.

The most important question for statistical analysis is whether a change in response instructions caused an interference. For example, suppose that switching the response



**Figure 3.** Unidimensional and diagonal category-learning curves for the control (A), the hand-switch (B), and the button-switch (C) conditions in Experiment 1.

buttons did interfere with the expression of category learning. A problem in documenting such interference, however, is that no matter how great this interference, eventually the observers would adapt to the new button locations and the interference would disappear. So if too many transfer trials are included in the analysis, no interference could ever be discovered, no matter how large that interference was initially. Ideally, then, only one transfer block would be included in the analyses. The complication is that there were two types of changes in response instructions in Experiment 1: (1) hand and button switches and (2) the introduction of a response deadline. Therefore, to verify a possible hand- or button-switch interference, we must first account for any possible response deadline interference. In fact, Figure 3 shows that with the diagonal category structures, even the control group experienced a significant drop in accuracy during

Block 11, which disappeared by Block 12 (a statistical analysis is presented below). Thus, in addition to any interference that might have been caused by switching hands or response buttons, with the diagonal categories there is strong evidence that mild speed stress also caused a temporary interference that disappeared after one block. For this reason, it is important to include Block 12 in any further statistical analyses.

An analysis of the accuracy rates prior to (Block 10) and following (Blocks 11 and 12) the change in response instructions largely supported the conclusions drawn from a visual analysis of the Figure 3 learning curves. To begin with, we investigated the effect of the change in response instructions by performing a 2 (category type: unidimensional vs. diagonal)  $\times$  3 (response instructions: control vs. hand switch vs. button switch)  $\times$  3 (blocks: 10–12) mixed design analysis of variance (ANOVA, with block as the within-subjects factor). The main effects of category type [ $F(1,109) = 105.1, MS_e = 100.78, p < .01$ ], response instructions [ $F(2,109) = 6.16, MS_e = 100.78, p < .01$ ], and block [ $F(2,218) = 23.59, MS_e = 31.98, p < .01$ ] were all significant. More important, however, these main effects were qualified by a significant three-way interaction [ $F(4,218) = 5.16, MS_e = 31.98, p < .01$ ]. Because we hypothesized that the effect of transfer would differ for the unidimensional and the diagonal category structures, we chose to examine the three-way interaction by performing separate response instructions  $\times$  block ANOVAs for the unidimensional and the diagonal categories. For the unidimensional category structures, neither the main effect of response instructions [ $F(2,37) = 0.03, MS_e = 27.12, p = .97$ ], nor the main effect of block [ $F(2,37) = 3.07, MS_e = 29.93, p = .06$ ], nor the response instruction  $\times$  block interaction [ $F(4,37) = 1.49, MS_e = 29.93, p = .21$ ] was significant. This analysis indicates that neither the sudden introduction of mild speed stress nor switching hands or response buttons interfered with categorization accuracy in the unidimensional conditions. This is strong evidence that the observers in these conditions were learning abstract category labels that were not tied to specific motor programs or response positions.

A different pattern of results was observed for the diagonal category structures. The main effects of both response instruction [ $F(2,72) = 9.54, MS_e = 105.50, p < .01$ ] and block [ $F(1,72) = 116.63, MS_e = 34.85, p < .01$ ] were significant. However, these main effects were qualified by a significant response instruction  $\times$  block interaction [ $F(2,72) = 18.83, MS_e = 34.85, p < .01$ ], suggesting that the change in accuracy across response blocks depended on the response instructions. A simple main effects analysis revealed a pattern of results consistent with the visual analysis of Figure 3. First, as was indicated above, there was a significant decrease in accuracy from Block 10 to Block 11 for the control group ( $p < .01$ ), suggesting that introducing a mild speed stress significantly lowered accuracy for the first 50 trials after the response deadline was imposed.<sup>4</sup> Even so, performance completely recovered during the second block of

50 trials (as evidenced by the lack of a significant difference between Blocks 10 and 12;  $p = .93$ ). Second, this same pattern of results was observed in the hand-switch group—that is, a significant decrease in accuracy during Block 11 ( $p < .01$ ), followed by a complete recovery during Block 12 ( $p = .21$ ). Similarly, performance in the button-switch group also significantly decreased in Block 11 ( $p < .01$ ), but in contrast to the control and the hand-switch groups, performance did not recover during Block 12 ( $p < .01$ ). Recall that, in addition to experiencing the same sudden speed stress as the control group, the observers in the hand-switch group were also required to switch their hands on the response keys. Despite this extra burden, the drop in accuracy during the first block after these changes was no different than the reduction observed in the control group [ $t(38) = -0.09, SE = 2.61, p = .93$ ]. In contrast, the button-switch group showed a much larger initial performance decrement than did either the control group [ $t(54) = 5.17, SE = 2.32, p < .01$ ] or the hand-switch group [ $t(52) = 4.91, SE = 2.39, p < .01$ ]. Therefore, unlike switching the hands, reversing the button assignments caused a substantial interference (decrement in performance between Block 10 and Block 11) in the diagonal condition. In sum, the results of these analyses are consistent with the hypothesis that the observers learn response positions in information-integration category learning, rather than abstract category labels or specific motor programs.

**Model-based analyses.** The interference in the button-switch condition with the diagonal categories could have occurred because the change in response locations interfered with the categorization strategy used by the observers or because it induced the observers to switch to a less accurate strategy. To test between these two alternatives, we fit a number of different decision bound models (Ashby, 1992a; Maddox & Ashby, 1993) to each observer's responses. Decision bound models assume that 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. Two different types of models were fit to each observer's responses (see the Appendix for more details). One type was compatible with the assumption that the observers used an explicit rule-based strategy, and one type assumed an information-integration strategy.

Table 2 shows the percentage of the observers whose data were best fit by a rule-based model. Inspection of Table 2 indicates that with the unidimensional categories, most observers used a rule-based strategy both prior to and during transfer. Importantly, a comparison of the pre- and the post-transfer data in Table 2, using a normal approximation to the binomial distribution, suggested that the change in response instructions had no significant effect on the dominance of rule-based strategies in any of the conditions [control,  $z = 0.61, SD = 0.11, p > .05$ ; hand switch,  $z = 0, SD = 0, p > .05$ ; button switch,  $z = 1.04, SD = 0.07, p > .05$ ]. Similarly, most observers who learned the

**Table 2**  
**Percentages of Observers Whose Data Were Best Fit by a**  
**Rule-Based Model One Block Before and After the**  
**Change in Response Instructions in Experiment 1**

Category Structure	Response Instruction	Before	After
Unidimensional	control	86.7	93.3
	hand switch	100	100
	button switch	92.3	100
Diagonal	control	9.5	14.3
	hand switch	0	15.8
	button switch	17.1	34.3

diagonal categories used an information-integration, rather than a rule-based, strategy both before and during transfer. However, an analysis of the Table 2 data revealed that although the change in response instructions had no effect in the control condition [ $z = 0.48$ ,  $SD = 0.10$ ,  $p > .05$ ], it did cause an increase in the dominance of rule-based strategies in the hand-switch [ $z = 1.97$ ,  $SD = 0.08$ ,  $p < .05$ ] and button-switch [ $z = 2.78$ ,  $SD = 0.06$ ,  $p < .01$ ] conditions. Thus, the instruction change disrupted the categorization strategy used by the observers more strongly with the diagonal structures than with the unidimensional structures, at least when hands or response buttons were switched. It appears, then, that information-integration strategies may be more fragile than rule-based strategies, in the sense that they are more easily disrupted by changes in experimental procedures.

Table 2 indicates that with the diagonal categories, a large number of the observers switched strategies when the response instructions changed. Because it is possible that the interference observed with the diagonal categories occurred simply because the participants switched to a less accurate response strategy, the data were examined for the observers who did not switch strategies. Analyzing the accuracy for those observers who used a strategy of the same type as the optimal classifier both before and after the change in response instructions (i.e., rule based with the unidimensional categories and information integration with the diagonal categories) did not change the pattern of results observed in the analysis of the accuracy data shown in Figure 3.

## EXPERIMENT 2

In Experiment 1, the diagonal categories were created by rotating the unidimensional categories shown in Figure 1 by 45°. For this reason, an ideal observer would perform equally on the two category structures, even if susceptible to perceptual and criterial noise (as would any common cluster or discriminant analysis algorithm). Even so, our observers were consistently less accurate on the diagonal categories, a result that replicates many other studies. We would argue that this accuracy difference is found in humans, but not in ideal observers, because humans have a system that can quickly learn unidimensional categories but this same system is relatively useless at learning the diagonal categories. Even if this

were true, however, one possible complication that could cloud the interpretation of the Experiment 1 results is that, because of the accuracy difference, observers in the unidimensional conditions had achieved a higher level of expertise at the time of transfer than had observers in the diagonal conditions. That is, although the observers in the two conditions had the same number of training trials, the diagonal categories took longer to learn, so the observers in these conditions had developed less expertise at the time of transfer than had the observers learning the unidimensional categories (i.e., note the difference in the learning curves shown in Figure 3). Because of this difference, the results of Experiment 1 make it difficult to rule out the hypothesis that switching hands or response keys interferes with categorization only when the categories are not well learned. The idea here is that, early in learning, category representations are fragile, and any change in the manner in which category knowledge is expressed could be catastrophic. As greater expertise is developed, however, category representations become more stable and might, therefore, be used in more flexible ways.

Experiment 2 tested this hypothesis. The conditions of Experiment 1 were replicated, except that the observers learning the diagonal categories were given twice as much training as they had received in Experiment 1, whereas the observers learning the unidimensional categories were given only 30% as much training as in Experiment 1. Thus, in Experiment 2, at the time of transfer, the observers learning the unidimensional categories were less expert than the observers learning the diagonal categories. As a result, if the expertise hypothesis is correct, some sort of interference should occur in the unidimensional conditions, but no interference should be observed in the diagonal conditions.<sup>5</sup> However, if rule-based categorization is mediated by a system that learns abstract category labels, whereas information-integration categorization is mediated by a procedural-learning-based system that learns response positions, changing the amount of training the observers receive should have relatively little effect on the amount of interference that occurs when the hands or the response keys are switched.

## Method

**Observers and Design.** Forty-two observers from the University of California, Santa Barbara community either received course credit for participation in this experiment (those learning the unidimensional categories) or were paid \$30 (those learning the diagonal categories). The number of participants in each of the experimental conditions were as follows: unidimensional/control, 10; unidimensional/hand-switch, 11; unidimensional/button-switch, 11; diagonal/control, 9; and diagonal/button-switch, 7. No observer participated in more than one experimental condition. All the observers reported 20/20 vision or vision corrected to 20/20. Each observer in the unidimensional conditions completed one session of approximately 15-min duration. Each observer in the diagonal conditions completed two sessions of approximately 45-min duration that were separated, on average, by 24 h. The criterion for learning was defined as 70% correct during the block preceding the change in response instructions. The data from the observers who did not meet

this criterion were excluded from all the subsequent analyses. This criterion resulted in the exclusion of 2 observers from the diagonal/control condition.

**Stimuli and stimulus generation.** The stimuli and stimulus generation procedures were identical to those in Experiment 1.

**Procedure.** The experimental procedures were identical to those in Experiment 1, except for the following. Each participant in the unidimensional conditions completed 5 blocks of trials (with 50 trials per block), and the change in response instructions occurred after Block 3 for all the conditions. Each participant in the diagonal conditions completed 12 blocks of training (50 trials per block) during the first experimental session. The second session occurred approximately 24 h later. The procedure during the second session was identical to that of the single sessions used in Experiment 1. Finally, the diagonal/hand-switch condition was omitted (see note 5).

**Results**

The learning curves for each of the three unidimensional conditions are shown in Figure 4. Note that the observers learned this category structure within two blocks. Also note the apparent lack of any effect of speed stress. In fact, for the control group, accuracy in the first transfer block was not significantly different from accuracy in the final training block [ $t(10) = 1.31, SE = 1.12, p = .22$ ]. For this reason, our statistical analyses can be focused on the more important first transfer block.

A 3 (response instructions)  $\times$  2 (block: 3 vs. 4) mixed ANOVA was performed on the accuracy data to investigate performance at transfer. Neither the main effect of response instructions [ $F(2,29) = 0.51, MS_e = 17.62, p = .61$ ] nor the response instructions  $\times$  block interaction [ $F(2,29) = 0.10, MS_e = 7.51, p = .90$ ] was significant. There was a small but significant decrease in accuracy across response instructions from Block 3 to Block 4 (the average was 2%;  $p < .05$ ). However, the magnitude of this decrease did not differ across conditions, so it does

not support the expertise hypothesis. More important, the absence of a main effect of response instructions or of an interaction between response instructions and block argues strongly against the expertise hypothesis.

In the case of the diagonal categories, the learning curves from the control and the button-switch conditions are shown in Figure 5. In both conditions, the observers steadily improved their accuracy throughout the training blocks, and there was no effect of speed stress during the first transfer block in the control condition [ $t(6) = 0.44, SE = 2.06, p = .68$ ]. The latter result suggests that the extra training (i.e., as compared with Experiment 1) was effective. In particular, at the time of transfer, the observers in the diagonal conditions of Experiment 2 were practiced enough that the sudden introduction of a response deadline had no effect on their performance. For this reason, the analyses that follow can focus on the critical first transfer block.

A comparison of the last block before transfer and the first block after is quite clear. Despite 1,100 trials of training, switching the response keys again caused a massive decrement in performance ( $M = 9.0\%, SD = 7.7\%$ ), whereas a simple introduction of mild speed stress had virtually no effect on performance ( $M = 0.9\%, SD = 5.45\%$ ). The results of a 2 response instructions  $\times$  2 block (Block 22 vs. Block 23) mixed ANOVA provided support for these conclusions. The main effect of response instructions was not significant [ $F(1,12) = 0.92, MS_e = 50.98, p = .36$ ], but there was a significant main effect of block [ $F(1,12) = 7.58, MS_e = 22.4, p < .05$ ] that was qualified by a significant response instructions  $\times$  block interaction [ $F(1,12) = 5.05, MS_e = 22.4, p < .05$ ]. A simple main effects analysis suggests that the interaction was driven by the presence of interference in the




Figure 4. Unidimensional learning curves in Experiment 2.



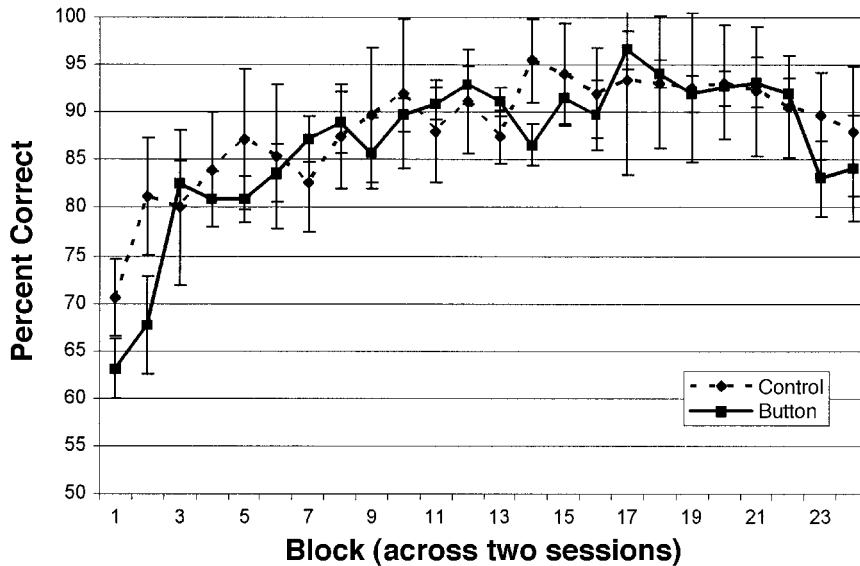


Figure 5. Diagonal learning curves for the control and the button-switch conditions over 2 days of training in Experiment 2.

button-switch condition ( $p < .01$ ) and the lack of interference in the control condition ( $p = .73$ ).

A comparison of Experiments 1 and 2 suggests that the biggest effect of extra training is to make categorization more immune to mild speed stress. In Experiment 1, accuracy of the control group dropped 6.7% when mild speed stress was introduced, whereas in Experiment 2, the control group showed no accuracy loss at all. It is well known that categorization response times decrease with practice (e.g., Nosofsky & Alfonso-Reese, 1999), so one possibility is that following the extra day of practice, the observers learning the diagonal categories in Experiment 2 were responding fast enough at the time of transfer so that the 1.5-sec response deadline was not experienced as stressful. Clearly, however, the accuracy loss exhibited by the button-switch group suggests that a similar phenomenon was not driving the accuracy loss that was seen in this group in Experiment 1.

Rule-based and information-integration models were fit to each observer's responses. Consistent with Experiment 1, all the observers in both the control and the button-switch conditions in Experiment 2 used information-integration strategies before the response instructions changed. The majority of the observers continued to use information-integration strategies following the change in response instructions (100% and 86% in the control and the button-switch conditions, respectively). This suggests that the additional training of Experiment 2 solidified the use of information-integration strategies.

## GENERAL DISCUSSION

The results of Experiments 1 and 2 lead to strikingly similar conclusions. When categories that can be sepa-

rated by a unidimensional rule are learned, there is no interference from the sudden introduction of mild speed stress, from requiring the observers to switch their hands on the response keys, or from reversing the button assignments. These results strongly suggest that with the unidimensional categories, the observers learned abstract category labels. As such, our results are consistent with the hypothesis that learning of the unidimensional categories is mediated by an explicit rule-based category-learning system in which the final product is a representation of an abstract category label. Under this scenario, the hand- and button-switch manipulations were accommodated by a straightforward remapping from the category label representations to motor planning and production systems.

The type of learning that occurred with the diagonal categories appears to be of a very different nature. First, there was a robust effect of speed stress on categorization accuracy in the diagonal condition that disappeared with practice. Second, there was no further interference caused by switching hands on the response buttons. In this condition, the position of the response buttons remained invariant, but the motor programs changed. Thus, the absence of any further interference in the hand-switch group (of Experiment 1) suggests that the category representation learned with the diagonal categories was more abstract than a simple motor command. Third, switching the motor response, as well as the position of the response buttons, produced significant additional interference that did not decrease with extended practice. This suggests that, in information-integration categorization, the observers learn what action to execute, but not a specific motor program with which to produce that action. In our experiments, the necessary action was to depress the re-

sponse button on the left in the presence of an exemplar from Category A and the response button on the right to an exemplar from Category B. As long as these actions did not change, there was no lasting interference caused by a sudden introduction of mild speed stress (e.g., in the hand-switch condition). Changing these actions, however, caused a profound and lasting interference.

All of these results are generally consistent with the many recent hypotheses that category learning is mediated by two or more separate systems (Ashby et al., 1998; Ashby & Ell, 2001, 2002b; Erickson & Kruschke, 1998; Pickering, 1997; Waldron & Ashby, 2001). For example, the dual-system COVIS model (Ashby et al., 1998) assumes separate explicit (rule-based) and implicit (procedural-learning-based) category-learning systems that compete throughout training. A system weight determines the relative contribution that each system makes to the observable categorization response. Initially, the system weight favors the explicit system, but it is then adjusted up and down depending on the relative success of the two systems. In rule-based tasks, the explicit system is successful, so the system weight strongly favors the explicit system throughout category learning. In information-integration tasks, however, the explicit system is less successful, so the weight given the procedural-learning system gradually grows. Even at asymptote, however, the weight given the explicit system in information-integration tasks remains reasonably large (because the explicit system is still suggesting many correct responses; Ashby et al., 1998; Ashby & Ell, 2002a). Thus, COVIS predicts that even highly practiced observers will use explicit rules on some trials in information-integration tasks (e.g., the rule that long lines are in Category B and short lines are in Category A works for the longest and the shortest lines in the diagonal categories of Figure 1). This could be why accuracy did not drop all the way to chance (or below) in the button-switch conditions and also why some observers were so adept at switching to a rule-based strategy during the transfer blocks.

Our results show that switching response locations selectively interferes with learning of the diagonal categories even when pretransfer accuracy is the same as with the unidimensional categories. Even so, one possible criticism of our conclusions is that even when accuracy is controlled in this way, the diagonal categories are still more difficult to learn than the unidimensional categories. If so, then perhaps our results show only that difficult tasks are more susceptible to interference than are simpler tasks.

This difficulty hypothesis is problematic for several reasons. First, there is no widely accepted measure of difficulty. Our experiments controlled difficulty according to two popular definitions (i.e., ideal observer performance and observed asymptotic accuracy), but many other definitions are possible (Alfonso-Reese, Ashby, & Brainard, 2002). Because there is such wide disagreement about how to define difficulty, it is not clear that

noncontroversial a priori predictions could be made from the difficulty hypothesis.

Second, several studies disconfirm the hypothesis that difficulty is the only difference between the unidimensional and the diagonal categories. If this were true, then unidimensional rule-based tasks should always be less susceptible to interference than are information-integration tasks. Evidence disconfirming this prediction was reported by Waldron and Ashby (2001), who showed that a simple rule-based category-learning task (which required attending to a single stimulus dimension) was disrupted more by a simultaneous (numerical Stroop) task than was a complex information-integration task (which required attending to three stimulus dimensions). Similarly, Ashby et al. (2003) found that Parkinson's disease patients were much more impaired at rule-based category learning (one-dimensional) than at information-integration category learning (three-dimensional).<sup>6</sup>

These studies are important because they show that it is invalid to predict a priori that learning in the easier unidimensional conditions is necessarily more resistant to interference than learning in the more difficult diagonal conditions. Therefore, the only way to know that switching response buttons will be more detrimental in the diagonal conditions is to run the experiment. Our data establish this result. We believe this selective interference occurred because information-integration category learning is largely procedural. Although we controlled for several measures of difficulty, it is important to acknowledge that our results do not completely rule out the possibility that difficulty played a role in our results. More work is needed on the role of difficulty in category learning.

## CONCLUSIONS

Some categories seem to be closely linked to a motor response. For example, when one suddenly sees a snake, it is natural to withdraw or retreat. When one plays tennis, certain shots hit by an opponent cause one to prepare a backhand return almost automatically, whereas other shots elicit forehand returns. Many other categories, however, seem completely divorced from any motor response. For example, categorizing an object as a square does not compel any specific action. The present results provide empirical evidence supporting this intuition.

Our results can be considered an extension of the results of Willingham et al. (2000) to the field of category learning. As was mentioned above, Willingham et al. (2000) showed that changing response locations interferes with implicit SRT learning but changing the sequence of finger movements does not. On the basis of these results, Willingham et al. (2000) argued that implicit SRT learning is the learning of response positions. Implicit SRT learning is perhaps the best-known example of procedural learning. As such, the similarity of our results to those of Willingham et al. (2000) supports the

hypothesis that information-integration category learning is largely mediated by a procedural-learning-based system.

## REFERENCES

- ALFONSO-REESE, L. A. (1997). *On the dangers of ignoring noise in high-level perception experiments*. Unpublished manuscript, Indiana University.
- ALFONSO-REESE, L. A., ASHBY, F. G., & BRAINARD, D. H. (2002). What makes a categorization task difficult? *Perception & Psychophysics*, **64**, 570-583.
- ASHBY, F. G. (1992a). Multidimensional models of categorization. In F. G. Ashby (Ed.), *Multidimensional models of perception and cognition* (pp. 449-483). 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 neurobiology of human 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 (2002). *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., & GOTT, R. E. (1988). Decision rules in the perception and categorization of multidimensional stimuli. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **14**, 33-53.
- 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., 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. M. (1999). On the dominance of unidimensional rules in unsupervised categorization. *Perception & Psychophysics*, **61**, 1178-1199.
- 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.
- BRAINARD, D. H. (1997). Psychophysics software for use with MATLAB. *Spatial Vision*, **10**, 433-436.
- CURRAN, T. (1995). On the neural mechanisms of sequence learning. *Psyche* [online], **2**(12).
- ERICKSON, M. A., & KRUSCHKE, J. K. (1998). Rules and exemplars in category learning. *Journal of Experimental Psychology: General*, **127**, 107-140.
- FILOTEO, J. V., MADDOX, W. T., & DAVIS, J. D. (2001). 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.
- 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.
- KNOWLTON, B. J., MANGELS, J. A., & SQUIRE, L. R. (1996). A neostriatal habit learning system in humans. *Science*, **273**, 1399-1402.
- KNOWLTON, B. J., SQUIRE, L. R., PAULSEN, J. S., SWERDLOW, N. R., SWENSON, M., & BUTTERS, N. (1996). Dissociations within non-declarative memory in Huntington's disease. *Neuropsychology*, **10**, 538-548.
- MADDOX, W. T., & ASHBY, F. G. (1993). Comparing decision bound and exemplar models of categorization. *Perception & Psychophysics*, **53**, 49-70.
- 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, & Cognition*, **29**, 650-662.
- 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.
- NISSEN, M. J., & BULLEMER, P. (1987). Attentional requirements of learning: Evidence from performance measures. *Cognitive Psychology*, **19**, 1-32.
- NOSOFSKY, R. M., & ALFONSO-REESE, L. A. (1999). Effects of similarity and practice on speeded classification response times and accuracies: Further tests of an exemplar-retrieval model. *Memory & Cognition*, **27**, 78-93.
- NOSOFSKY, R. M., & JOHANSEN, M. K. (2000). Exemplar-based accounts of "multiple-system" phenomena in perceptual categorization. *Psychonomic Bulletin & Review*, **7**, 375-402.
- NOSOFSKY, R. M., & KRUSCHKE, J. K. (2002). Single-system models and interference in category learning: Commentary on Waldron and Ashby (2001). *Psychonomic Bulletin & Review*, **9**, 169-174.
- NOSOFSKY, R. M., & PALMERI, T. J. (1998). A rule-plus-exception model for classifying objects in continuous-dimension spaces. *Psychonomic Bulletin & Review*, **5**, 345-369.
- NOSOFSKY, R. M., & ZAKI, S. R. (1998). Dissociations between categorization and recognition in amnesic and normal individuals: An exemplar-based interpretation. *Psychological Science*, **9**, 247-255.
- PICKERING, A. D. (1997). New approaches to the study of amnesic patients: What can a neurofunctional philosophy and neural network methods offer? In A. R. Mayes & J. J. Downes (Eds.), *Theories of organic amnesia* (pp. 255-300). Hove, U.K.: Psychology Press.
- REBER, P. J., & SQUIRE, L. R. (1994). Parallel brain systems for learning with and without awareness. *Learning & Memory*, **1**, 217-229.
- REBER, P. J., STARK, C. E. L., & SQUIRE, L. R. (1998). Contrasting cortical activity associated with category memory and recognition memory. *Learning & Memory*, **5**, 420-428.
- SCHWARZ, G. (1978). Estimating the dimension of a model. *Annals of Statistics*, **6**, 461-464.
- 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., JONIDES, J., & KOEPE, R. A. (1996). Dissociating verbal and spatial working memory using PET. *Cerebral Cortex*, **6**, 11-20.
- SMITH, E. E., PATALANO, A. L., & JONIDES, J. (1998). Alternative strategies of categorization. *Cognition*, **65**, 167-196.
- SOLIVERI, P., BROWN, R. G., JAHANSAHI, M., CARACENI, T., & MARSDEN, C. D. (1997). Learning manual pursuit tracking skills in patients with Parkinson's disease. *Brain*, **120**, 1325-1337.
- SQUIRE, L. R., KNOWLTON, B. J., & MUSEN, G. (1993). The structure and organization of memory. *Annual Review of Psychology*, **44**, 453-495.
- THOMAS-OLLIVIER, V., REYMAN, J. M., LE MOAL, S., SCHÜCK, S., LIEURY, A., & ALLAIN, H. (1999). Procedural memory in recent-onset Parkinson's disease. *Dementia & Geriatric Cognitive Disorders*, **10**, 172-180.
- WALDRON, E. M., & ASHBY, F. G. (2001). The effects of concurrent task interference on category learning: Evidence for multiple category learning systems. *Psychonomic Bulletin & Review*, **8**, 168-176.
- 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 Psychology: Learning, Memory, & Cognition*, **15**, 1047-1060.
- WILLINGHAM, D. B., WELLS, L. A., FARRELL, J. M., & STEMWEDDEL, M. E. (2000). Implicit motor sequence learning is represented in response locations. *Memory & Cognition*, **28**, 366-375.

## NOTES

1. A rule-plus-exception model that allows generalization from exceptions (Nosofsky & Palmeri, 1998) fares no better. First, it predicts that

because amnesiacs cannot memorize exceptions, they should always use unidimensional strategies, which is incorrect (Filoteo, Maddox, & Davis, 2001). Second, it predicts that no factor should interfere with rule-based learning more strongly than with information-integration learning. Several studies have disconfirmed this prediction (Ashby, Noble, Filoteo, Waldron, & Ell, 2003; Waldron & Ashby, 2001).

2. We have used similar category structures in a number of other experiments (Ashby, Maddox, & Bohil, 2002; Ashby, Queller, & Berretty, 1999), in addition to a second set of unidimensional and diagonal structures that were obtained by rotating the categories shown in Figures 1A and 1B each by 90°. Thus, in these other experiments, two unidimensional categories (horizontal and vertical tending) and two diagonal categories (positive and negative slope) were used. In each of these experiments, there was no qualitative difference between the observers' performance with either unidimensional category structure, nor was there a difference between performance with either diagonal category structure. For this reason, in the present study, we included only a single example of each category structure type.

3. The length and orientation dimensions in Figure 1 are expressed in degrees of visual angle and degrees from horizontal, respectively, for

presentation purposes only. To generate the stimuli, length was defined in pixels, and orientation was defined in arbitrary units that resulted in a range of 0 to 700 units on both the length and the orientation dimensions.

4. All pairwise comparisons were made using Sidak-corrected 95% confidence intervals.

5. The expertise hypothesis predicts that increased practice should make categorization less susceptible to interference when response instructions are suddenly changed. No interference was observed in the hand-switch condition with the diagonal categories in Experiment 1. As a result, the expertise hypothesis predicts that extra practice should have no effect in this condition. For this reason, the hand-switch condition was dropped for the diagonal categories in Experiment 2.

6. Parkinson's disease (PD) patients are impaired in procedural learning (e.g., Soliveri, Brown, Jahanshahi, Caraceni, & Marsden, 1997; Thomas-Ollivier et al., 1999), so we predict PD deficits in information-integration category learning. In fact, such deficits do exist (Knowlton, Mangels, & Squire, 1996; Maddox & Filoteo, 2001). However, Ashby et al. (1998) also predicted that the extensive damage that PD causes to the head of the caudate nucleus would lead to an even greater PD deficit in rule-based category learning.

## APPENDIX

Here, we briefly describe the decision bound models. For more details, see Ashby (1992a) or Maddox and Ashby (1993).

### Rule-Based Models

Two models assume that the observers use an explicit rule-based strategy.

**The unidimensional model.** This model assumes that the observer sets a criterion on a single perceptual dimension and then makes an explicit decision about the level of the stimulus on that dimension (Ashby & Gott, 1988; Shaw, 1982). It has two free parameters: a decision criterion on the relevant perceptual dimension and the variance of internal (perceptual and criterial) noise (i.e.,  $\sigma^2$ ). In the unidimensional conditions, a special case of the unidimensional model assumes that observers use the unidimensional decision bound that maximizes accuracy (i.e., the vertical bound shown in Figure 1A). This special case has only one free parameter (i.e., noise variance).

**The conjunction model.** Another possible rule-based strategy is that the observers use a conjunction rule in which they make separate decisions about the levels on the two dimensions and then select a response on the basis of the outcome of these two decisions. Conjunction models have three parameters (a criterion on each dimension and  $\sigma^2$ ).

### Information Integration Models

**The general linear classifier (GLC).** This model 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 length and orientation. The GLC has three parameters (slope and intercept of the linear bound and  $\sigma^2$ ). In the diagonal conditions, a special case of the GLC assumes that observers use the linear bound that maximizes accuracy (i.e., the diagonal bound shown in Figure 1B). This model has only one free parameter (noise variance).

**The general quadratic classifier (GQC).** A natural extension of the GLC is to assume that the observer uses a quadratic, rather than a linear, decision bound. This model also produces an information-integration strategy, but the integration of perceived length and orientation is nonlinear. The GQC has six free parameters (five describing the form of the decision bound and  $\sigma^2$ ).

### Model Fitting

Each of these models were fit separately to the data for every observer from the last pretransfer block (i.e., Block 10) and the first post-transfer block (i.e., Block 11). The model parameters were estimated using maximum likelihood (Ashby, 1992b; Wickens, 1982), and the goodness-of-fit statistic was

$$\text{BIC} = r \ln N - 2 \ln L,$$

where  $N$  is the sample size,  $r$  is the number of free parameters, and  $L$  is the likelihood of the model given the data (Schwarz, 1978). The BIC statistic penalizes a model for bad fit and for extra free parameters. To find the best model among a set of competitors, one simply computes a BIC value for each model and then chooses the model with the smallest BIC.