Procedural interference in perceptual classification: Implicit learning or cognitive complexity?

ROBERT M. NOSOFSKY and ROGER D. STANTON
Indiana University, Bloomington, Indiana

and

SAFA R. ZAKI
Williams College, Williamstown, Massachusetts

Researchers have argued that an implicit procedural-learning system underlies performance for information integration category structures, whereas a separate explicit system underlies performance for rule-based categories. One source of evidence is a dissociation in which procedural interference harms performance in information integration structures, but not in rule-based ones. The present research provides evidence that some form of overall difficulty or category complexity lies at the root of the dissociation. The authors report studies in which procedural interference is observed for even simple rule-based structures under more sensitive testing conditions. Furthermore, the magnitude of the interference is large when the nature of the rule is made more complex. By contrast, the magnitude of interference is greatly reduced for an information integration structure that is cognitively simple. These results challenge the view that a procedural-learning system mediates performance on information integration categories, but not on rule-based ones.

A modern debate in the perceptual classification literature concerns the issue of whether or not multiple representational systems underlie categorization performance (e.g., Ashby, Alfonso-Reese, Turken, & Waldron, 1998; Erickson & Kruschke, 1998; Nosofsky & Johansen, 2000; Nosofsky & Zaki, 1998). Furthermore, assuming that there are multiple systems, what is their nature, and how do the systems interact?

One highly influential proposal stems from the COVIS (competition between a verbal and an implicit system) model of Ashby et al. (1998). According to the COVIS model, there are at least two separate cognitive systems that underlie perceptual classification. One system learns categories by using explicit reasoning that is closely related to the formation of verbal rules. A second cognitive system relies on a form of implicit learning. Other multiple-system models have also been proposed that posit combinations of explicit and implicit representational systems. However, a distinctive aspect of the COVIS model involves its assumption that the implicit-learning system in categorization is procedurally based. By contrast, in other multiple-system models, the implicit system might

involve the storage of specific exemplars or category prototypes (Erickson & Kruschke, 1998; Pickering, 1997; Reber, Stark, & Squire, 1998).

In recent work, Ashby, Ell, and Waldron (2003) sought to provide distinctive evidence that favors the procedural-learning assumption of the COVIS model. The key idea that motivated their experiments is that procedural learning is closely associated with motor performance and the learning of specific motor responses (Willingham, Nissen, & Bullemer, 1989). By contrast, there is no reason to expect close ties, for example, between exemplar-based learning and motor responses.

To pursue this idea, Ashby et al. (2003) conducted experiments in which forms of response–motor interference were introduced into the learning of different types of categories. The researchers distinguished between two fundamentally different types of category structures. *Rule-based* category-learning tasks are ones in which the category structure can be learned by an explicit reasoning process that generally involves the formation of a verbal rule. From a formal point of view, in rule-based tasks, the observer makes a separate decision along each individual dimension about the category regions where a percept falls. These separate decisions are then combined to determine the overall categorization response (Ashby & Gott, 1988; Nosofsky, Clark, & Shin, 1989). A simple example is the unidimensional category structure that is illustrated in the top panel of Figure 1. In this example, the stimuli vary along two continuous dimensions, but only one dimension is relevant for performing the clas-
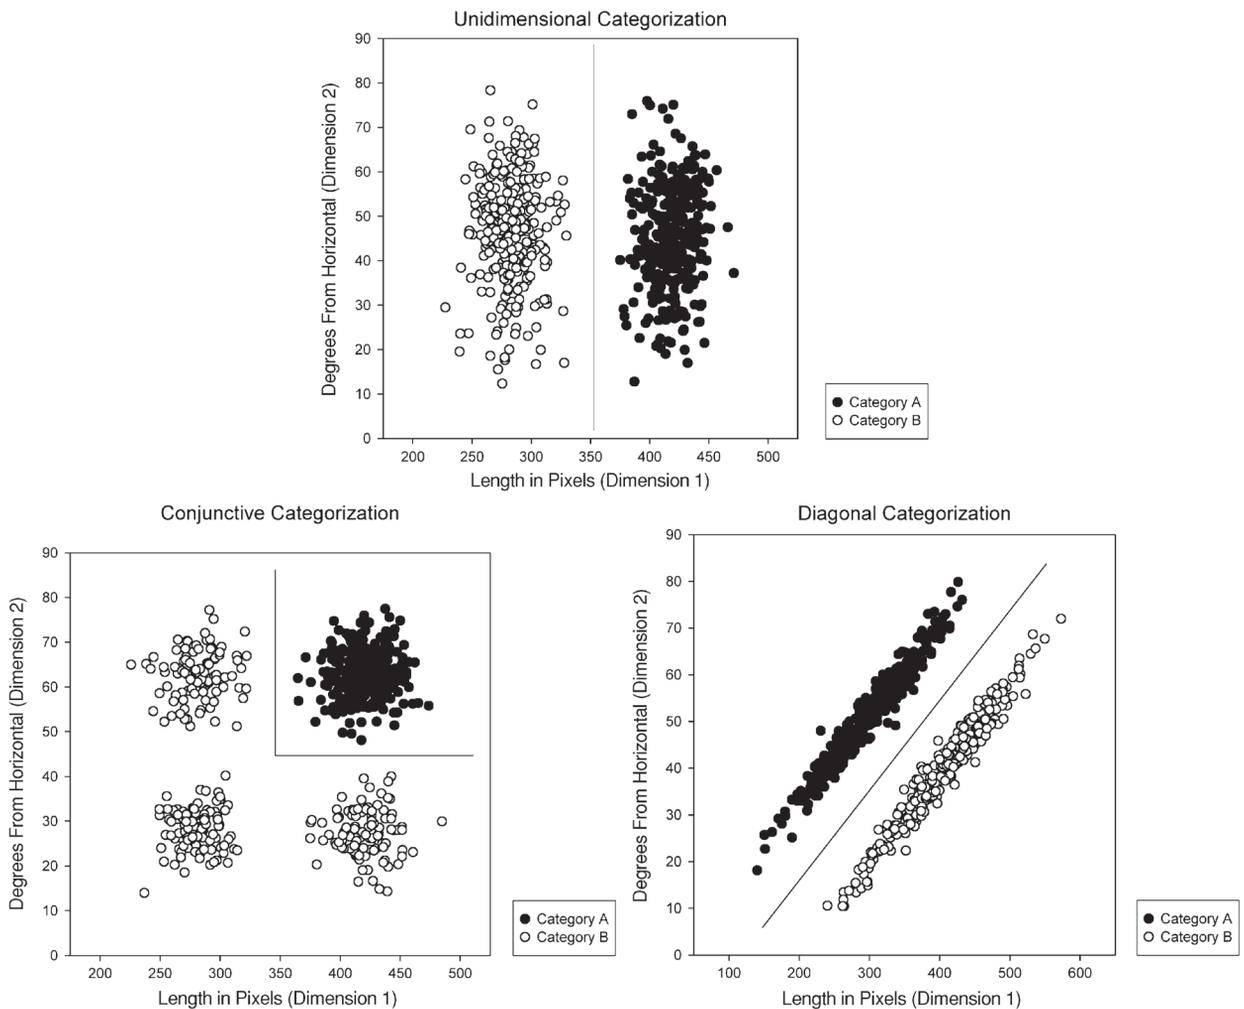
Figure 1. Three different category structures with associated decision boundaries for partitioning the stimuli into response regions.

sification. Stimuli with large values on Dimension 1 are assigned to Category A, whereas stimuli with small values are assigned to Category B. The *decision bound* for dividing the space into response regions is the vertical line illustrated in the figure. Any percept that falls to the right of the boundary is classified into Category A, whereas any percept that falls to the left of the boundary is classified into Category B. The boundary formalizes the simple verbal rule: "Respond A if Dimension 1 has a large value and respond B if Dimension 1 has a small value."

Another example of a rule-based task is the conjunctive rule structure illustrated in the bottom left panel of Figure 1. In this case, if a stimulus has a large value on Dimension 1 *and* has a large value on Dimension 2, it belongs to Category A; otherwise, it belongs to Category B. Note, again, that separate decisions are made along each individual dimension and these individual decisions are then combined to determine the categorization response. In general, the hallmark of a rule-based task involving combinations of independent decisions is that the deci-

sion boundaries that divide the perceptual space into category regions are orthogonal to the coordinate axes of the space.

By contrast, an example of an *information integration* category-learning task is the diagonal categorization illustrated in the bottom right panel of Figure 1. In general, in an information integration task, accuracy can be maximized only if perceptual information from two or more component dimensions is combined prior to making any classification decisions (Ashby & Gott, 1988). In this case, the optimal decision boundaries that divide the perceptual space into category regions are no longer orthogonal to the coordinate axes. In the example shown in the bottom right panel of Figure 1, the optimal decision boundary is the diagonal line illustrated in the figure. Note that whereas combinations of independent-decision boundaries correspond to explicit rules that can be easily verbalized, such is not the case for information integration structures. For example, suppose that Dimensions 1 and 2 correspond to length and angle. Then the "rule"

formalized by the diagonal boundary might be stated as the following: "Members of Category A are those stimuli in which the angle exceeds the length." However, because length and angle are not directly comparable, this statement has no direct psychological implementation.

According to Ashby et al.'s (1998) COVIS model, performance in simple rule-based tasks, such as the unidimensional category structure in Figure 1, is mediated by the explicit-reasoning system, which is not procedurally based. By contrast, learning and performance in information integration tasks, such as the diagonal categorization, is mediated primarily by the implicit procedural-learning system. Thus, Ashby et al. (2003) reasoned that there should be a close link between learning and motor performance in the diagonal categorization task, but not in the unidimensional categorization task. To test this idea, Ashby et al. (2003) conducted experiments in which subjects first learned the unidimensional or diagonal structures. Following initial learning, Ashby et al. (2003) then conducted transfer tests. During transfer, the category structures remained the same. In a control condition, a response time deadline was imposed, such that subjects needed to make their classification responses within 1,500 msec of the presentation of the stimulus. In the critical experimental condition, in addition to imposing this response time deadline, the subjects were required to switch the response buttons that they used for indicating their categorization assignments. That is, the button that previously had indicated a Category A response now indicated a Category B response, and the button that previously had indicated a Category B response now indicated a Category A response. Thus, in this button switch condition, both the response locations and the motor responses were reversed from what had occurred during training. To the extent that procedural learning underlies performance, this button switch manipulation should induce interference at the time of transfer. The response time deadline was imposed in all conditions to ensure that the subjects could not overcome such interference simply by inhibiting their initial response.

The key result observed by Ashby et al. (2003) was that marked interference was observed in the button switch condition for the diagonal category structure, but no interference was observed in the button switch condition for the unidimensional category structure. Ashby et al. (2003) interpreted these results as support for the hypothesis that rule-based categorizations are learned via the explicit-reasoning system and involve the assignment of abstract category labels to regions of psychological space. By contrast, information integration categorizations are largely mediated by a procedural system that learns specific motor responses.[1]

Although this evidence of selective interference in the diagonal categorization task is intriguing and could, indeed, provide distinctive evidence in favor of the multiple-system view proposed by Ashby et al. (1998), alternative possibilities need to be considered. As was acknowledged by Ashby et al. (2003), there is a sense in which the di-

agonal category structure is more complex and *difficult* than is the unidimensional category structure. Thus, as was stated by Ashby et al. (2003), the selective interference in the diagonal categorization task might "show only that difficult tasks are more susceptible to interference than are simpler tasks" (p. 1123). Ashby et al. (2003) took various steps to try to equate forms of difficulty across the two tasks. For example, they manipulated the overall discriminability of the stimuli across conditions, as well as the amount of training prior to conducting the transfer test. In all cases, they continued to observe button switch interference in the diagonal categorization task, but not in the unidimensional categorization task. Nevertheless, Ashby et al. (2003) concluded: "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" (p. 1123).

Given the importance of the issue of multiple systems in category learning, the central goal of the present research was to pursue further the difficulty explanation of Ashby et al.'s (2003) results. We hypothesize that regardless of the structure of the perceptual classification task (i.e., rule vs. information integration), some procedural-learning component is always involved, in the sense that subjects have formed a habit of using particular motor responses for indicating given category labels. Furthermore, the button switch requirement will introduce some degree of interference with the procedurally learned response. If the task is sufficiently simple, however, and the response time deadline period sufficiently long, the interference can be corrected before an observable decrement in performance occurs. Thus, it is possible that the particular response time deadline period used by Ashby et al. (2003; 1,500 msec) was long enough to mask any observable interference for the simple unidimensional rule-based task but was too short to mask the interference for the more complex diagonal task.

We pursued this hypothesis in the present research by following two main approaches. In the first approach, we continued to test the same unidimensional rule-based task as that used by Ashby et al. (2003). However, instead of using only the single 1,500-msec deadline, we varied parametrically the deadline across conditions. The difficulty explanation suggests that with a sufficiently sensitive paradigm, an interference due to button switching should eventually be observed. In particular, the hypothesis predicts that one should observe an interference in performance in the button switch condition, relative to the control condition, when the response time deadline period is made sufficiently short.

As a second approach to testing the difficulty explanation, we explored a variety of other rule-based and information integration categorization tasks. The idea here was to decouple the class of categorization task (rule-based vs. information integration) from the variable of difficulty. Thus, rather than testing only the simple unidimensional

rule-based task, we tested a rule-based task that was cognitively complex. Conversely, we tested an information integration task that, relatively speaking, seemed cognitively simple. According to Ashby et al.'s (1998) COVIS model and the hypothesis that only information integration tasks are procedurally based, the degree of interference due to button switching should be greater for the information integration tasks than for the rule-based ones, regardless of their difficulty. By contrast, the difficulty explanation suggests that a good deal of interference should be observed for a complex rule-based task, whereas minimal interference should be observed for an easy information integration task.

We should acknowledge at the outset that we do not have a formal theory of *difficulty* or *category complexity* and that, to a large degree, the predictions from the difficulty explanation are based on intuitive judgments. Nevertheless, regardless of this limitation, the experiments still provide strong tests of the hypothesis that procedural interference due to button switching is limited primarily to information integration category structures.

EXPERIMENT 1

The purpose of our initial experiment was simply to replicate the main results from Ashby et al. (2003) involving the unidimensional rule-based and diagonal categorization tasks. It was important to verify that our laboratory procedures yielded the same fundamental results as those in Ashby et al. (2003), before moving on to new experimental conditions. In addition, the results from these initial conditions could then serve as baselines of comparison for the new conditions introduced in this research.

Table 1
Number of Subjects in Each Condition in Each Experiment

Condition	Subjects Included	Subjects Excluded
Experiment 1		
Unidimensional control	21	1
Unidimensional button switch	21	1
Diagonal control	21	3
Diagonal button switch	24	3
Experiment 2		
Unidimensional control (400)	21	1
Unidimensional button switch (400)	19	0
Unidimensional control (750)	20	2
Unidimensional button switch (750)	23	1
Experiment 3		
Biconditional control	21	3
Biconditional button switch	20	4
Experiment 4		
Center control	21	1
Center button switch	23	2

Note—Subjects included = number of experimental subjects who met the learning criterion and were included in the statistical analyses. Subjects excluded = number of experimental subjects who did not meet the learning criterion and were not included in the statistical analyses. Values in parentheses in the Experiment 2 conditions are the response time deadlines (in milliseconds).

Method

Subjects. In this section, we will describe aspects of the subject populations that are common to all four experiments reported in this research. The subjects were undergraduates from Indiana University, who participated as part of a course requirement or who were paid for their participation. An average of 23.08 subjects participated in each condition in each experiment. The exact number of experimental participants in each condition in each experiment is reported in Table 1. Using an analytic procedure similar to one used by Ashby et al. (2003), in all the conditions of all the experiments, we excluded from the statistical analysis any subject who failed to meet a 70% learning criterion in either of the two final blocks of training. An average of 1.83 subjects were excluded in each condition in each experiment, with the exact numbers reported in Table 1.

In all the conditions of all the experiments, to increase motivation, the subjects were informed that they could earn up to a $4 bonus for performing well in the task. They could earn a $2 bonus for achieving a performance criterion in the learning phase, plus an additional $2 for achieving a performance criterion in the test phase. The subjects were not aware of the exact criterion level, only that they needed to perform "sufficiently well" to earn the bonus. (The actual bonus criterion was 90% accuracy in the final two practice blocks and the test blocks for the unidimensional categorizations in Experiments 1 and 2 and 85% accuracy for all other category structures. A response was scored as *correct* only if the correct response button was pressed prior to the response time deadline.)

Stimuli. The stimuli and stimulus generation procedure were the same as those used by Ashby et al. (2003), and we refer the reader to this earlier article for details. In brief, the stimuli were lines varying along the dimensions of length and orientation. The stimuli were sampled randomly from categories defined by bivariate normal distributions with the parameters listed in Table 2. A total of 300 stimuli defined each of Categories A and B for both the rule-based and the diagonal tasks. The resulting category structures are those shown in Figure 1 (top and bottom panels).

Procedure. There were two conditions tested for each category structure: control and button switch. Each task consisted of a training phase and a transfer phase. The training phase was identical for the control and the button switch conditions. On each trial, a stimulus was presented; the subject classified it into either Category A or B, and corrective feedback was then provided. For the unidimensional categorization, training consisted of 8 blocks of 50 trials each; whereas for the diagonal categorization, training consisted of 10 blocks of 50 trials each.[2] The stimuli were randomly assigned to trials for each subject, with the constraint that no single stimulus was presented more than once. Each trial during training had a 5-sec response time deadline; if a response was not provided within that time period, the computer provided a prompt to the subject to respond more quickly.

Immediately following the training phase, instructions for the transfer phase were provided on the computer screen. For the control condition, the only change in instructions was that the deadline for making responses was now 1,500 msec, instead of 5,000 msec. For the button switch condition, in addition to the shorter response time deadline, the subjects were instructed about the switch in response buttons used for indicating their categorization decisions. Specifically, they were instructed that "the response buttons for each category are now reversed. That is, lines that were correctly classified with an 'A' response in the previous blocks are now correctly classified with a 'B' response. Likewise, lines that were correctly classified with a 'B' response in previous blocks are now correctly classified with an 'A' response." Otherwise, the procedure was the same as that for the training phase. There were two 50-trial blocks of transfer.

Results

Here and in all of the other experiments, we used two modes of analysis of the data. Following Ashby et al.

Table 2
Population Parameter Values Used to Generate the
Category Structures Tested in Experiments 1–4

	μ_L	μ_A	σ_L^2	σ_A^2	σ_{LA}
Experiments 1 and 2, Unidimensional					
Category A	420	150	330	1,500	0
Category B	280	150	330	1,500	0
Experiment 1, Diagonal					
Category A	300	171.4	3,500	1,500	1,462
Category B	400	128.5	3,500	1,500	1,462
Experiment 3, Biconditional					
Category A					
Distribution A1	280	208.28	330	275	0
Distribution A2	420	91.67	330	275	0
Category B					
Distribution B1	280	91.67	330	275	0
Distribution B2	420	208.28	330	275	0
Experiment 4, Center					
Category A					
Distribution A1	350	150	200	100	0
Category B					
Distribution B1	350	250	200	100	0
Distribution B2	490	150	200	100	0
Distribution B3	350	50	200	100	0
Distribution B4	210	150	200	100	0

Note—μ_L = mean length, μ_A = mean angle, σ_L^2 = variance length, σ_A^2 = variance angle, σ_{LA} = length-angle covariance. Length is measured in pixels. The angle values, measured in radians, were multiplied by the scale factor $\pi/600$.

(2003), in the first mode, we deleted from analysis any trial in which the subject failed to meet the response time deadline. In the second mode, we did not delete any trials from the analysis and defined a response to be correct only if the correct category label was provided and the subject met the response time deadline. In all cases, both modes of analysis led to the same pattern of results and conclusions. For simplicity, therefore, we will report the results from only the first mode of analysis.

The mean probability of correct classifications is shown as a function of category task (unidimensional vs. diagonal), condition (control vs. button switch), and block of training/transfer in Figure 2. The results replicated closely the pattern of results observed by Ashby et al. (2003) in their previous study. First, accuracy was consistently higher for the unidimensional rule-based task than for the diagonal task. Second, during the critical transfer blocks, there was no decrement in performance (relative to the final training blocks) in either condition of the unidimensional categorization task. Third, there was an overall decrement in performance at the time of transfer in the diagonal task. Most important, this decrement in performance was greater in the button switch condition than in the control condition. The latter result is the key interference effect of interest.

To confirm these observations, we conducted a series of ANOVAs. First, for each categorization task taken separately, we conducted focused mixed-model two-way ANOVAs with condition (control vs. button switch) and block

(final two blocks of training vs. the two blocks of transfer) as factors. (In this analysis, condition is a between-subjects factor, and block is a within-subjects factor.) In the unidimensional task, the analysis yielded no main effects of condition or block and no interaction. By contrast, in the diagonal task, the analysis yielded the following: a main effect of condition [$F(1,43) = 7.06$, $MS_e = 0.014$, $p = .011$], reflecting the fact that overall performance was better in the control condition than in the button switch condition; a main effect of block [$F(1,43) = 9.61$, $MS_e = 0.003$, $p = .003$], reflecting the fact that overall performance was better at the end of training than during transfer; and a condition × block interaction [$F(1,43) = 7.53$, $p = .009$], reflecting the fact that the decrement in transfer performance was greater in the button switch condition than in the control condition.

A concern that might be raised is that, for the diagonal structure, the subjects in the button switch condition appear to have performed worse than the subjects in the control condition, even during the training phase (in which the two conditions were logically identical). We computed the mean probability of correct classifications across Blocks 1–10 for each subject in each condition. The difference in performance across conditions did not approach statistical significance [$t(18) = 1.001$, $p = .33$].

Finally, a three-way ANOVA with category task (unidimensional vs. diagonal), condition, and block as factors also revealed a three-way interaction [$F(1,83) = 6.06$, $MS_e = 0.013$, $p = .016$], reflecting the fact that any button switch decrement in transfer performance was greater in the diagonal task than in the unidimensional task.

Discussion

The results from Experiment 1 replicated closely the pattern of results observed previously by Ashby et al. (2003). To review, according to Ashby et al.'s (2003) interpretation, the results are consistent with the hypothesis that performance in the diagonal task has a procedural-learning component, which is disrupted by the button switching requirement. By contrast, the lack of interference in the unidimensional task is consistent with the idea that rule-based tasks rely on a separate explicit-learning system that is not procedurally based. Having replicated this main pattern of results, we now turned to the investigation of the alternative idea that all category-learning tasks involve a procedural-learning component. Instead, the present pattern of results may be better explained in terms of some form of overall task difficulty, rather than in terms of a fundamental distinction between rule-based and information integration category learning.

EXPERIMENT 2

In this experiment, we again tested subjects with the unidimensional categorization task in both the control and the button switch conditions. However, at time of transfer, instead of using only the single 1,500-msec deadline, we placed greater speed stress on the subjects by also testing
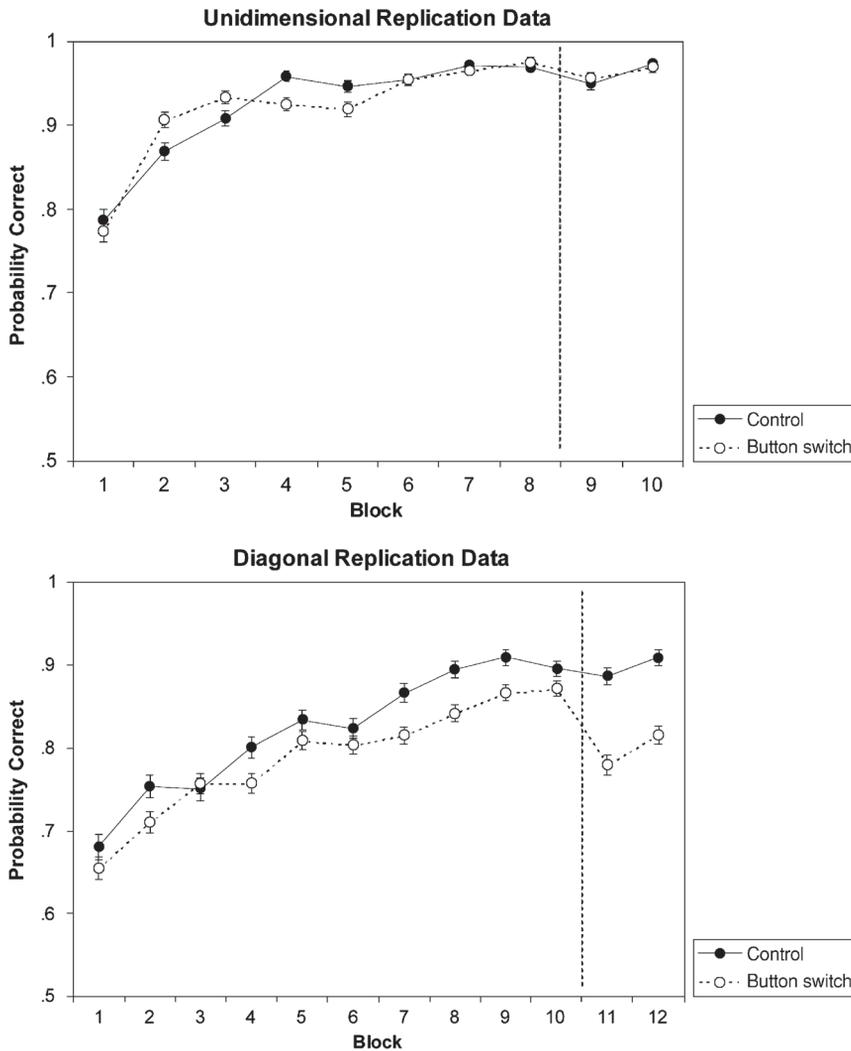
Figure 2. Mean probability of correct responses as a function of condition (control vs. button switch) and block in Experiment 1. The dashed vertical line in each panel indicates the break between training and transfer. Top panel, unidimensional categorization; bottom panel, diagonal categorization.

a 750-msec and a 400-msec deadline. Possibly, placing greater speed stress on the subjects might yield a more sensitive indicator of a procedural-learning component even in the simple rule-based task. Specifically, the overall difficulty explanation predicts that button switch interference should be observed even in the rule-based task if there is a sufficiently stringent response time deadline. Because interference was already demonstrated in the diagonal categorization task even with the longer response time deadline (1,500 msec), we did not test diagonal categorization with the more stringent deadlines.

Method

Subjects. The subjects were from the same population as that in Experiment 1. The number of subjects who participated in each condition is reported in Table 1.

Stimuli. The stimuli and category structure were the same as those for the unidimensional rule-based task in Experiment 1.

Procedure. The procedure was the same as that for the rule-based task in Experiment 1, except that we tested two new response deadlines at time of transfer: one set at 750 msec and the other at 400 msec.

Results

The mean probability of correct classifications is shown as a function of deadline (750 vs. 400 msec), condition (control vs. button switch), and block in Figure 3. A comparison with the top panel of Figure 2 reveals a fundamentally different pattern of results when a more stringent response deadline was used in the rule-based task. Recall that at the longer deadline (1,500 msec), there was no decrement in transfer performance in either the control or the button switch condition (see Figure 2, top panel). By

contrast, at these shorter deadlines, there was a decrement in performance in both the control and the button switch conditions. Most critically, the decrement in performance was greater in the button switch condition than in the control condition, which is the key interference effect that is indicative of a disruption of procedural learning.

We conducted a mixed-model three-way ANOVA of all of the unidimensional task data (including the unidimensional task data from Experiment 1). The between-subjects factors were response deadline (1,500, 750, or 400 msec) and condition (control vs. button switch), whereas the within-subjects factor was block (final two blocks of training vs. the two blocks of transfer). The analysis revealed the following: a main effect of response deadline [$F(1,119) = 57.3, MS_e = 0.295, p < .001$], reflecting worse overall performance in the more stringent deadline conditions; a main effect of condition [$F(1,119) = 9.01, MS_e = 0.046, p = .003$], reflecting worse overall performance in the button switch condition than in the control condition; a main effect of block [$F(1,119) = 211.4, MS_e = 0.505, p < .001$], reflecting worse overall performance during transfer than during training; and a two-way interaction between condition and block [$F(1,119) = 6.65, MS_e = 0.016, p = .011$], reflecting the fact that the overall decrement in transfer performance was greater in the button switch condition than in the control condition. Again, this two-way interaction is the key interference effect. Finally, the three-way interaction among response deadline, condition, and block was significant [$F(1,119) = 2.61, MS_e = 0.006, p = .039$, one-tailed],[3] reflecting the fact that the interfer-
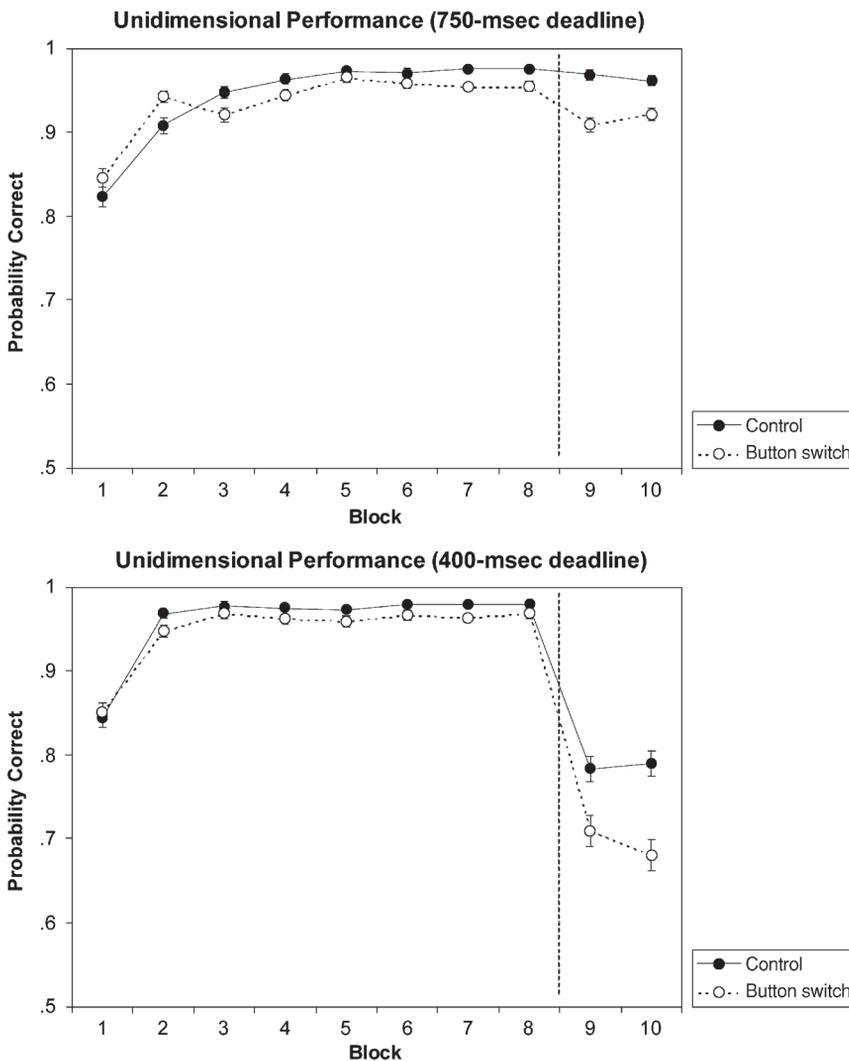

Figure 3. Mean probability of correct responses as a function of condition (control vs. button switch) and block for the unidimensional categorization task in Experiment 2. The dashed vertical line in each panel indicates the break between training and transfer. Top panel, 750-msec response deadline; bottom panel, 400-msec response deadline.

ence effect was observed only as the response deadline was made more stringent.

Discussion

When the response time deadline was made more stringent, an interference effect was observed in the button switch condition even for the simple unidimensional rule-based task. This result is consistent with the idea that even the unidimensional rule task involves a procedural-learning component. A reasonable interpretation is that, given the simplicity of the unidimensional task, the subjects had time to "undo" any interference caused by button switching when the 1,500-msec deadline was imposed. However, under the more stringent deadlines, the underlying interference was revealed.

It is likely that had we also tested diagonal categorization with the more stringent response time deadlines, the magnitude of interference due to button switching would have exceeded the magnitude of interference in the unidimensional categorization task. Such a result is irrelevant, however, to the key issue at hand. Both the multiple-system hypothesis and the overall difficulty hypothesis predict that such a pattern of results would be observed. By contrast, the strong prediction from the multiple-system COVIS model is that no procedural interference should have been observed for the unidimensional categorization task, regardless of the response time deadline.

EXPERIMENT 3

In Experiment 3, we attempted to provide further evidence in support of a difficulty explanation of Ashby et al.'s (2003) results. Specifically, instead of testing subjects on the simple unidimensional rule-based task, we tested them on a complex rule-based task—namely, the biconditional structure illustrated in Figure 4. The explicit rule is that a stimulus is a member of Category A if it has a small value on Dimension 1 and has a large value on Dimension 2 *or* if it has a large value on Dimension 1 and a small value on Dimension 2. Otherwise, the stimulus is a member of Category B. Note that, according to Ashby and Gott's (1988) decision boundary framework, the biconditional structure is an example of a rule-based task, because the decision boundaries for dividing the space into response regions are orthogonal to the coordinate axes. In other words, the observer makes independent decisions about the values of a stimulus on each of its dimensions and then combines these separate decisions to determine the appropriate categorization response. To avoid any ambiguity about the use of such an explicit reasoning strategy, the subjects were given explicit instructions about the form of the rule prior to the start of the experiment. They needed only to learn the location of the various decision boundaries in order to implement the rule. In all other respects, the procedure was the same as that in Experiment 1, with both a control and a button switch condition tested at a 1,500-msec response deadline. Given the complexity of the biconditional rule, the difficulty explanation predicts that button switch interference will be observed (even at this longer 1,500 msec deadline), just as it was for diagonal categorization. By contrast, if procedural learning is not involved in explicit rule-based tasks, any decrement in transfer performance should be no greater in the button switch condition than in the control condition.

Method

Subjects. The subjects were from the same population as that in Experiments 1 and 2. The number of subjects who participated in each condition is reported in Table 1.
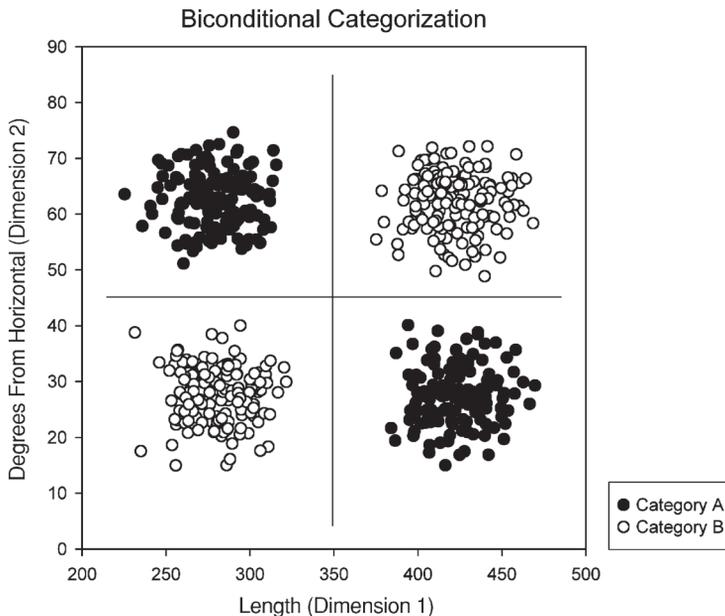

Figure 4. Structure of the biconditional categorization tested in Experiment 3.

Stimuli. The stimuli were again lines varying in their length and orientation. The biconditional category structure was generated by using mixtures of bivariate normal distributions. The parameters of these distributions are reported in Table 2, with the resulting category structure illustrated in Figure 4. The structure has the property that the standardized distance between the means of each adjacent bivariate distribution is approximately equal to the standardized distance between the means of the two categories from the unidimensional structure on its relevant dimension.

Procedure. The procedure was identical to the one used for the diagonal task in Experiment 1, except that the subjects were provided with explicit instructions concerning the form of the rule at the outset of the experiment. This portion of the instructions read as follows. "There is a simple rule you can use to correctly classify all of the lines. If the line is short and at a steep angle, or if the line is long and at a shallow angle it belongs to one category. Alternatively, if the line is short and at a shallow angle, or if the line is long and at a steep angle it belongs to the other category."

Results

The mean probability of correct classifications in the biconditional task is shown as a function of conditions and blocks in Figure 5. It is immediately apparent that the button switch manipulation led to dramatic interference, with transfer performance in the button switch condition being substantially worse than that in the control condition. Indeed, the pattern of results for this biconditional structure is extremely similar to that observed previously for the diagonal structure (compare Figures 2 and 5), despite the fact that the biconditional structure is a rule-based task, whereas the diagonal structure is an information integration task.

To confirm these observations, we conducted a two-way ANOVA on the biconditional data, using condition (control vs. button switch) and block (last two blocks of training vs. the two blocks of transfer) as factors. The analysis yielded the following: a main effect of condition [$F(1,39) = 8.59$, $MS_e = 0.091$, $p = .006$], reflecting the worse overall performance in the button switch condition; a main effect of block [$F(1,39) = 46.6$, $MS_e = 0.128$, $p < .001$], reflecting the worse overall performance at time of transfer; and a condition $\times$ block interaction [$F(1,39) = 18.7$, $p < .001$], reflecting the fact that the decrement in transfer performance was significantly greater in the button switch condition than in the control condition.

Discussion

The results from the biconditional structure support the difficulty explanation of the pattern of interference effects that result from button switching. Although it is a rule-based task and the subjects were provided with explicit instructions concerning the form of the rule, the pattern of interference is extremely similar to that observed previously for the information integration (diagonal) structure. Thus, the pattern of interference does not seem to be linked to whether the task involves explicit verbal rules or implicit information integration but, rather, to factors involving the overall difficulty of the task.

EXPERIMENT 4

Our final experiment was exploratory in nature. Our goal was to provide a final converging source of evidence for the difficulty explanation of the pattern of interference effects. Specifically, whereas in Experiment 3 we tested a difficult and cognitively complex rule-based task, the idea in Experiment 4 was to test a simple and relatively easy information integration task. We were unable, however, to think of an information integration structure that seemed intuitively as simple as the unidimensional rule tested in Experiment 1.

The structure that we did test is illustrated in Figure 6. For ease of description, we will refer to the structure as the *center* categorization task. As can be seen, Category A
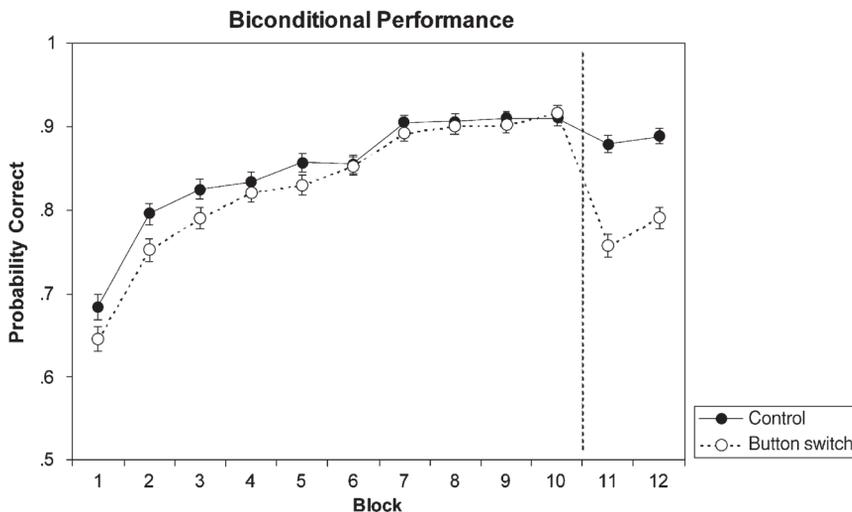

Figure 5. Mean probability of correct responses as a function of condition (control vs. button switch) and block for the biconditional categorization task in Experiment 3. The dashed vertical line indicates the break between training and transfer.

consists of a single bivariate normal distribution in the center of the two-dimensional space, whereas Category B consists of a collection of other bivariate normal distributions located along the periphery. We reasoned that such a task might be very easy to perform. Observers could perform the task, for example, by forming a single prototype to represent Category A. Any percept that was sufficiently similar to the prototype would be classified in Category A; else, the stimulus would be classified in Category B. Such a strategy is represented schematically in terms of the circular boundary of equal similarity to the prototype that is illustrated in the top panel of Figure 6.

We should acknowledge that we cannot prove that all subjects use such an information integration strategy to perform the task. For example, a possible rule-based strategy is to form four decision boundaries that are orthogonal to the coordinate axes, as is illustrated in the bottom panel of the figure. In this case, the subject classifies a percept as a member of Category A if it has a value greater than $X1$ on Dimension 1 and less than $X2$ on Dimension 1 *and* has a value greater than $Y1$ on Dimension 2 and less than $Y2$ on Dimension 2; otherwise, the subject classifies the stimulus into Category B. However, the number of explicit decision boundaries that is used for partitioning the space in this manner is even greater than what was involved for the biconditional structure. Thus, to the extent that performance is quite good in the center categorization task, it would seem to place strain on such a multiple independent-decision boundaries explanation.

Method

Subjects. The subjects were from the same population as those in the previous experiments, with the number of subjects who participated in each condition reported in Table 1.

Stimuli. The stimuli were the same as those used in the previous experiments. The population parameter values used to generate the category structure are reported in Table 2. To maintain equal category base rates, exemplars were sampled four times as often from Distribution A1 as from Distributions B1–B4. The resulting category distribution is illustrated in Figure 6.

Procedure. The procedure was the same as that in the previous experiments. The subjects were not provided with any explicit instructions concerning the form of the category structure but, rather, learned it online via induction over the training exemplars.

Results and Discussion

The mean probability of correct classifications in the center categorization task is shown as a function of condition and block in Figure 7. Although it took some time for the subjects to learn the task, by the end of training (Blocks 9 and 10), the subjects were performing at very high levels, essentially the same as those observed for the unidimensional task. An inspection of Figure 7 reveals interference due to button switching (i.e., transfer performance in the button switch condition lies below transfer performance in the control condition). This observation was confirmed by statistical test, with a two-way ANOVA revealing a significant interaction between condition and block [$F(1,42) = 10.16$, $MS_e = 0.003$, $p = .003$]. However, the magnitude of the interference effect was fairly

small, with mean performance dropping only from 96.1% to 93.1%. By contrast, for the biconditional structure tested in Experiment 3, mean performance dropped from 91.0% to 77.4%. Some caution is needed in comparing the magnitude of these effects, because they are almost certainly scale dependent. Nevertheless, the results can be characterized as providing preliminary evidence of a case in which there was a relatively small degree of interference due to button switching in an information integration task—at the same response deadline that simultaneously produced a good deal of interference in a rule-based task. Further research, however, is needed to confirm such results, using alternative information integration structures and response time deadlines.

GENERAL DISCUSSION

Dissociations and Multiple-Categorization Systems

A important debate in the modern perceptual classification literature concerns the issue of whether or not multiple cognitive systems underlie the learning and representation of perceptual categories (Ashby & Casale, 2003; Knowlton & Squire, 1993; Nosofsky & Johansen, 2000; Palmeri & Flanery, 2002; Shanks & St. John, 1994; Smith, Patalano, & Jonides, 1998). Much of the evidence for distinct representational systems is based on the demonstration of a variety of experimental dissociations involving classification performance. The general interpretation is that damage or interference involving one representational system will harm performance in some types of cognitive tasks, although leaving intact performance in other tasks that do not rely on the system. Well-known examples include the studies of Squire, Knowlton, and associates (Knowlton & Squire, 1993; Reed, Squire, Patalano, Smith, & Jonides, 1999; Squire & Knowlton, 1995). These researchers have demonstrated that amnesic subjects often display normal perceptual classification performance, despite showing dramatic deficits in various explicit memory tasks, such as recognition and recall. The interpretation is that recognition and recall rely on an explicit memory system, which is damaged in amnesics, whereas perceptual classification relies on an implicit system, which remains intact.

However, other researchers have argued that such dissociations are also naturally predicted by single-representation-system models (Kinder & Shanks, 2001; Nosofsky & Zaki, 1998; Palmeri & Flanery, 1999). For example, Nosofsky and Zaki demonstrated that a single-system exemplar model predicted the essential features of Knowlton and Squire's recognition–categorization dissociations, and Zaki and Nosofsky (2001) demonstrated that the single-system model predicted dissociations between cued recall and classification in amnesics. The key idea in these demonstrations is that a single representational system based on stored exemplars underlies both memory and classification performance. Naturally, it is assumed that the exemplar-based memory system of the amnesics is
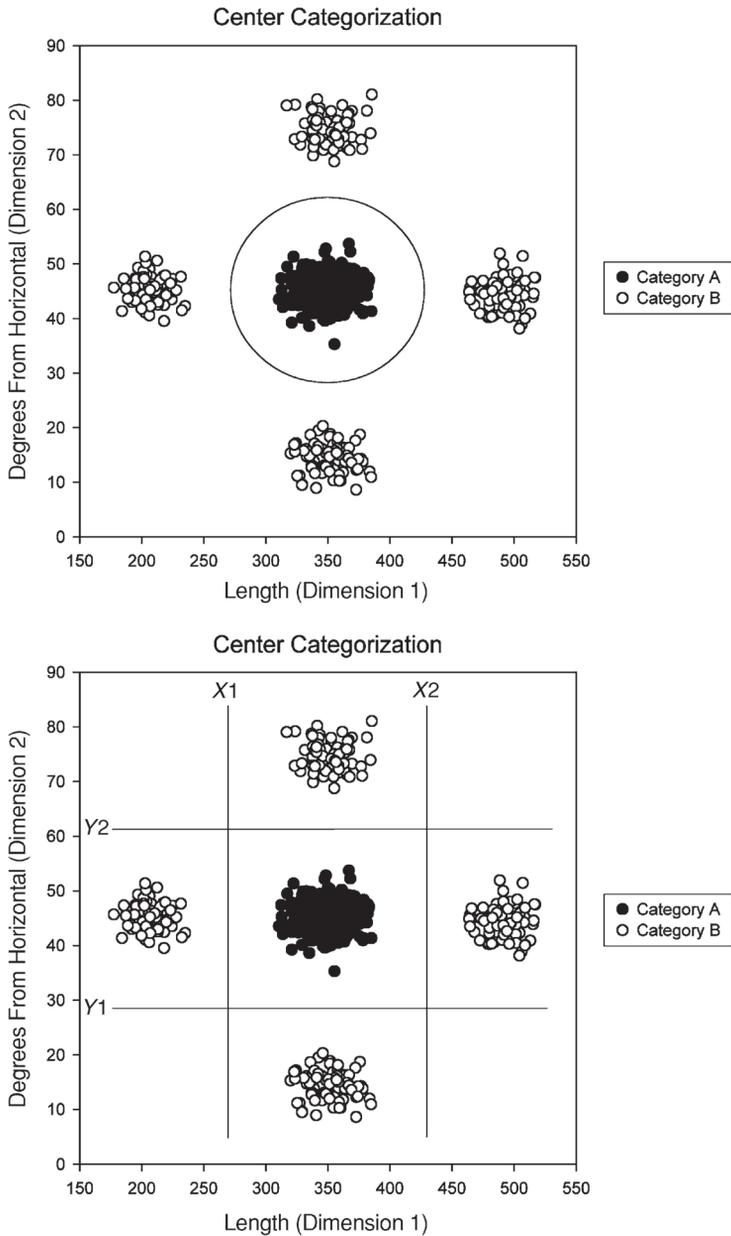
Figure 6. Structure of the center categorization task tested in Experiment 4. Top panel, schematic similarity-to-prototype boundary; bottom panel, schematic multiple independent-decision boundaries.

impaired. However, in the tasks conducted by Knowlton, Squire, and colleagues, it turns out that a fine-grained exemplar-based memory is not needed for achieving near-normal classification performance. Intuitively, in these studies, gross-level discriminations between distinct exemplars are sufficient to support categorization. By contrast, an exemplar-based memory impairment leads to marked reductions in recognition and cued recall, which require an ability to make fine-grained discriminations between old and new items.

Importantly, in further support of the single-system interpretation, researchers have also demonstrated that these dissociations disappear when more powerful statistical approaches and/or diagnostic experimental procedures are applied. For example, in a meta-analysis of categorization memory experiments involving amnesic subjects, Zaki (2004) demonstrated that although the effects in any given study may be small and not statistically significant, amnesics do, in fact, appear to display consistent deficits in classification performance, as compared with normals,
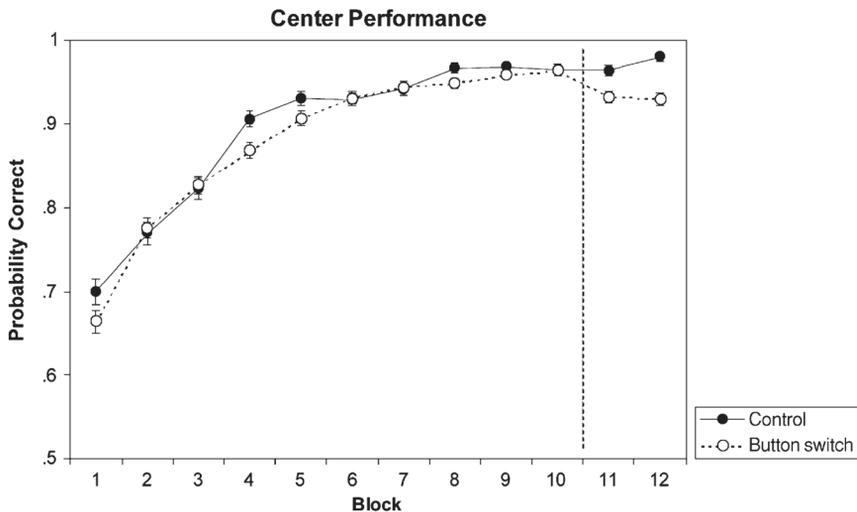
Figure 7. Mean probability of correct responses as a function of condition (control vs. button switch) and block for the center categorization in Experiment 4. The dashed vertical line indicates the break between training and transfer.

when the results are considered collectively. Likewise, when more diagnostic classification tasks are tested that make larger demands on exemplar-based memory, amnesic subjects show large deficits in classification performance (Zaki, Nosofsky, Jessup, & Unverzagt, 2003).

The theme of the present research was similar in spirit to those of the previous research summarized above. Specifically, the goal of the present research was to begin to address a new line of classification dissociations. These dissociations are consistent with the hypothesis that distinct cognitive systems underlie the representation of rule-based versus information integration categories, as hypothesized in the multiple-system COVIS model (Ashby et al., 1998). In an important representative study along these lines, Ashby et al. (2003) demonstrated a dissociation in which a procedural interference manipulation based on button switching disrupted performance in an information integration categorization task, but not in a rule-based task. This result was consistent with the COVIS hypothesis that rule-based tasks are learned via an explicit verbal-reasoning system, whereas information integration tasks are learned via an implicit procedural-learning system. However, as was acknowledged by Ashby et al. (2003), an alternative possibility was that some form of overall difficulty or category complexity was at the root of the effects.

In the present research, we hypothesized that all perceptual classification learning tasks generally involve some form of procedural learning and decided to pursue the alternative, *difficulty* explanation that was offered by Ashby et al. (2003). We followed two main approaches. First, we demonstrated that even in the simple unidimensional rule-based tasks tested by Ashby et al. (2003), the button switch manipulation does harm performance when the response time deadline is made more stringent. Our interpretation is that when a sufficiently sensitive experi-

mental test is applied, the underlying procedural interference is revealed. Second, we demonstrated that even with long response time deadline periods, the button switch manipulation leads to dramatic procedural interference for a more complex rule-based categorization—namely, the biconditional structure. Finally, we provided preliminary evidence that it may be possible to find easy types of information integration tasks that lead to fairly small magnitudes of interference due to button switching. In a nutshell, therefore, we believe that the reported dissociation involving effects of procedural interference on rule-based versus information integration tasks seems to have disappeared as well.

The Ashby et al. (2003) study is a single representative from a larger body of experiments that have demonstrated similar dissociations between rule-based and information integration tasks (for a review, see Ashby & Casale, 2003). Importantly, we believe that the dissociations in various of these other studies may have had a similar locus in terms of the role of overall category complexity. For example, Maddox, Bohil, and Ing (2004) manipulated whether or not there was a consistent response location for a given category. In their design, subjects learned either an information integration or a unidimensional rule-based category structure. The key manipulation was the specific question asked during training. In one condition, the subjects were shown a particular stimulus and were asked "Is this an A or a B?" Note that in this A–B condition, there was a consistent mapping of category label to computer key. In another condition, the subjects were randomly asked one of two questions on each trial: "Is this item from Category A?" or "Is this item from Category B?" In this *yes–no* condition, no particular key was associated with a particular category. Maddox, Bohil, and Ing reasoned that because procedural learning is closely

associated with motor responses, a learning deficit should be observed in the *yes–no* condition, relative to the A–B condition, but only for the information integration category structure. In fact, this pattern of results was observed. However, as Maddox, Bohil, and Ing acknowledged, the *yes–no* condition may simply be more difficult than the A–B condition, and this difficulty interacts with the complexity of the information integration category structure.[4] We hypothesize, for example, that strong interference would be observed in the *yes–no* condition if a biconditional structure were tested, instead of a simple unidimensional rule structure.

We do not claim that all of the dissociations demonstrated in this emerging literature on multiple classification systems have their locus in overall difficulty. Interpretation of the basis for each dissociation requires a careful task analysis, as well as a consideration of alternative single-representation accounts of category learning. For example, Waldron and Ashby (2001) observed that performing a concurrent task caused greater interference in the learning of a unidimensional rule-based categorization than in the learning of a multiple-dimension information integration categorization. Because the unidimensional task was easier than the information integration task, Waldron and Ashby argued that the results challenged single-representation accounts of category learning. However, Nosofsky and Kruschke (2002) demonstrated that a single-system exemplar memory model (Kruschke's [1992] ALCOVE model) predicted naturally the results by assuming that its selective-attention learning process was disrupted by the concurrent task.

In another example from this literature, Maddox, Ashby, Ing, and Pickering (2004) conducted an experiment in which subjects learned either a rule-based or an information integration categorization. Following the presentation of category feedback, the subjects performed a visual-memory–scanning task. Maddox, Ashby, et al. found that if feedback-processing time was short, there was a good deal of interference for the rule-based categorization, but no interference for the information integration categorization, a result that was consistent with predictions from the multiple-system COVIS model. However, to equate *difficulty* across the rule-based and the information integration categorizations, Maddox, Ashby, et al. reduced the physical spacing between the contrasting categories in the rule-based task such that the stimuli were very hard to discriminate. Conceivably, the introduction of visual memory scanning may have disrupted the subjects' perceptual memories of the stimuli from the categorization task. Such an effect would be particularly devastating in a condition in which discriminability is so poor to begin with. Thus, although the results from Maddox, Ashby, et al. are quite intriguing, our view is that further research is needed to confirm the basis for the observed dissociation.

Toward a Single-System Modeling Account of the Procedural Interference

The research reported in this article challenges the multiple-system hypothesis of a procedural-learning system linked solely to information integration category structures. A natural question, however, is whether or not a single-representation system model may be formulated that could provide an account of the complete set of procedural interference effects. Developing and providing rigorous tests of such a model entails a highly ambitious project that goes beyond the scope of the present research. Nevertheless, it seems worthwhile to outline the form that such a model might take on.

One promising approach involves an extension of current exemplar-based models of speeded classification performance (Cohen & Nosofsky, 2003; Lamberts, 1998, 2000; Nosofsky & Palmeri, 1997b). According to these models, people represent categories by storing individual exemplars in memory and classify objects on the basis of their similarity to the stored exemplars. For example, according to the exemplar-based random-walk (EBRW) model (Nosofsky & Palmeri, 1997b), presentation of a test item causes individual exemplars to be retrieved, and the retrieved exemplars drive a random-walk process that leads to classification decisions. Classification response times are determined by the speed with which the random walk is completed.

The most recent formulation of the EBRW model (Cohen & Nosofsky, 2003) borrows a key idea from Lamberts's (2000) exemplar-based model of speeded classification. Specifically, there is a stochastic process that determines which of the perceptual dimensions of a test object have been encoded at any given point in time. The rate at which the dimensions are sampled and encoded is determined jointly by their intrinsic psychological salience and by the degree of selective attention that an observer devotes to each dimension. The similarity between a test object and the stored exemplars changes systematically depending on the progress of this stochastic perceptual encoding process. Thus, when implemented within the framework of the EBRW, there are dynamic changes in which exemplars tend to be retrieved as the random walk proceeds.

To apply the model to the present experimental results, further assumptions would be needed that link the classification-decision–making process with subsequent processes involving response selection and motor execution. The general processing framework that we envision is as illustrated in Figure 8. The dimensional encoding, exemplar retrieval, and classification-decision–making stages, which are themselves highly interactive, are as already formalized in the EBRW. Once a classification decision has been made, a response must be selected, and the response must then be executed by an appropriate motor program. Presumably, during the course of training, associations are developed between given classification decisions and the responses that signify those decisions. Likewise, associations are developed for the motor programs that execute the responses.

The EBRW model has already been applied successfully to account for performance in speeded classification tasks such as those tested in the control conditions in the present experiments (Nosofsky & Palmeri, 1997a, 1997b; Nosofsky & Stanton, 2005). It is well documented, for example, that the model accounts in quantitative detail for such phenomena as *distance from boundary* effects, in which

objects that are far from the category decision boundary are classified more accurately and more rapidly than are objects that lie close to the boundary. Furthermore, the model also accounts well for speed–accuracy trade-offs. For example, if placed under time pressure, subjects may need to reduce the magnitude of the decision criteria that assess the information accumulated in the random walk. Such a process leads to faster classifications at the expense of lowered accuracy. In situations involving highly separable-dimension stimuli, increased time pressure may also lead to a failure to encode all of the dimensions of a test object, a process that results in other systematic changes in patterns of classification (Lamberts & Freeman, 1999).

Within the framework of this EBRW modeling approach, the button switch manipulation likely has consequences at several different stages of the information-processing sequence. In general, because associations between classification decisions and selected responses need to be "undone" at the time of transfer, the total time to execute a correct response should tend to increase. If the execution time still falls below the imposed deadline, the process can operate as in the control conditions, and an interference effect will not be observed. The hypothesis that motivated our Experiment 2 is that such a situation arises for the easy unidimensional categorization in the long deadline period condition. However, to meet the deadline under the more stringent conditions, the observer may need to complete the encoding, exemplar retrieval, and decision-making stages more quickly by reducing the magnitude of the random-walk decision criteria. In this case, the observer would thereby sacrifice accuracy for speed. Furthermore, on some proportion of trials, the random-walk decision process may not have time to come to completion. In such cases, the system might emit a random guess. We expect that noise is probably introduced at the response selection stage as well. For example, even if the classification decision process is completed, an error may arise because there is inadequate remaining time to undo the previously learned response association. Each of these events would lead to lowered overall accuracy, thereby contributing to an account of the button switch interference effects.

The modeling approach outlined above also has the potential to explain more microlevel aspects of the effects of button switch interference on speeded classification performance. For example, Ashby et al. (2003) conducted analyses in which simple linear decision boundary models were fitted to the individual-trial data from the speeded classification tasks. The two main decision bound models

that were contrasted were a *rule-based* model, in which the decision boundary was assumed to be a vertical line; and an *information integration* model, in which the slope of the decision boundary was allowed to be a free parameter. (See Ashby et al., 2003, p. 1125, for a fuller discussion of the nature of these decision boundary models.) Not surprisingly, Ashby et al. (2003) found that, for the unidimensional classification data, the rule-based model yielded more parsimonious fits than did the information integration model, regardless of whether the data were collected during training or transfer or in the control or the button switch condition. Likewise, they found that, for the vast majority of subjects, the information integration model yielded more parsimonious fits to the diagonal classification data.

However, a potentially interesting result was their finding that, during transfer in the button switch condition for the diagonal task, the rule-based model provided more parsimonious fits than did the information integration model for 34.3% of the subjects who were tested. We conducted the same decision bound model analyses on our Experiment 1 data and observed the same basic pattern of results as that reported by Ashby et al. (2003). (We report the detailed results of these model-based analyses in the Appendix.) This pattern of results has a number of possible explanations in terms of the EBRW model outlined above. For example, because of the increased time pressure induced by button switching, the subjects may not have had time to encode both dimensions of the test stimuli on some proportion of the trials. If an observer tended to encode only one of the dimensions, the pattern of data predicted by the EBRW model would be well characterized in terms of a *rule-based* decision boundary—that is, a boundary that is orthogonal to the coordinate axes of the space. Lamberts and Freeman (1999) documented the manner in which patterns of classification performance change dynamically with changing response deadlines. And both Lamberts and Freeman (1999) and Cohen and Nosofsky (2003) have demonstrated that exemplar models of speeded classification that incorporate assumptions involving stochastic, time-based encoding of dimensional information account well for these dynamic effects.

In sum, extensions of current exemplar retrieval models appear to offer a promising avenue for accounting for the influence of procedural interference on speeded classification performance. Future research is needed to develop fully specified versions of such models and to conduct rigorous tests that compare such models with multiple-system alternatives.
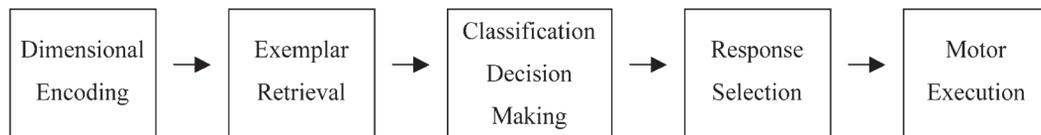

Figure 8. Schematic information-processing sequence for an extended version of the exemplar-based random-walk model.

REFERENCES

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., & CASALE, M. B. (2003). The cognitive neuroscience of implicit category learning. In L. Jimenez (Ed.), *Attention and implicit learning* (Advances in Consciousness Research, Vol. 48, pp. 109-142). Amsterdam: Benjamins.

ASHBY, F. G., ELL, S. W., & WALDRON, E. M. (2003). Procedural learning in perceptual categorization. *Memory & Cognition*, **31**, 1114-1125.

ASHBY, F. G., & GOTT, R. E. (1988). Decision rules in the perception and categorization of multidimensional stimuli. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **14**, 33-53.

COHEN, A. L., & NOSOFSKY, R. M. (2003). An extension of the exemplar-based random-walk model to separable-dimension stimuli. *Journal of Mathematical Psychology*, **47**, 150-165.

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

HOMMEL, B. (1993). Inverting the Simon effect by intention. *Psychological Research*, **55**, 270-279.

KINDER, A., & SHANKS, D. R. (2001). Amnesia and the declarative/non-declarative distinction: A recurrent network model of classification, recognition, and repetition priming. *Journal of Cognitive Neuroscience*, **13**, 648-669.

KNOWLTON, B. J., & SQUIRE, L. R. (1993). 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**, 22-44.

LAMBERTS, K. (1998). The time course of categorization. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **24**, 695-711.

LAMBERTS, K. (2000). Information-accumulation theory of speeded categorization. *Psychological Review*, **107**, 227-260.

LAMBERTS, K., & FREEMAN, R. P. J. (1999). Building object representations from parts: Tests of a stochastic sampling model. *Journal of Experimental Psychology: Human Perception & Performance*, **25**, 904-926.

MADDOX, W. T., ASHBY, F. G., ING, A. D., & PICKERING, A. D. (2004). Disrupting feedback processing interferes with rule-based but not information-integration category learning. *Memory & Cognition*, **32**, 582-591.

MADDOX, W. T., BOHIL, C. J., & ING, A. D. (2004). Evidence for a procedural-learning–based system in category learning. *Psychonomic Bulletin & Review*, **11**, 945-952.

NOSOFSKY, R. M., CLARK, S. E., & SHIN, H. J. (1989). Rules and exemplars in categorization, identification, and recognition. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **15**, 282-304.

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. (1997a). Comparing exemplar-retrieval and decision-bound models of speeded perceptual classification. *Perception & Psychophysics*, **59**, 1027-1048.

NOSOFSKY, R. M., & PALMERI, T. J. (1997b). An exemplar-based random walk model of speeded classification. *Psychological Review*, **104**, 266-300.

NOSOFSKY, R. M., & STANTON, R. D. (2005). Speeded classification in a probabilistic category structure: Contrasting exemplar-retrieval, decision-boundary, and prototype models. *Journal of Experimental Psychology: Human Perception & Performance*, **31**, 608-629.

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.

PALMERI, T. J., & FLANERY, M. A. (1999). Learning about categories in the absence of training: Profound amnesia and the relationship between perceptual categorization and recognition memory. *Psychological Science*, **10**, 526-530.

PALMERI, T. J., & FLANERY, M. A. (2002). Memory systems and perceptual categorization. In B. H. Ross (Ed.), *The psychology of learning and motivation* (Vol. 41, pp. 141-189). San Diego: Academic Press.

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., STARK, C. E. L., & SQUIRE, L. R. (1998). Contrasting cortical activity associated with category memory and recognition memory. *Learning & Memory*, **5**, 420-428.

REED, J. M., SQUIRE, L. R., PATALANO, A. L., SMITH, E. E., & JONIDES, J. (1999). Learning about categories that are defined by object-like stimuli despite impaired declarative memory. *Behavioral Neuroscience*, **113**, 411-419.

SCHWARZ, G. (1978). Estimating the dimension of a model. *Annals of Statistics*, **6**, 461-464.

SHANKS, D. R., & ST. JOHN, M. F. (1994). Characteristics of dissociable human learning systems. *Behavioral & Brain Sciences*, **17**, 367-447.

SMITH, E. E., PATALANO, A. L., & JONIDES, J. (1998). Alternative strategies of categorization. *Cognition*, **65**, 167-196.

SQUIRE, L. R., & KNOWLTON, B. J. (1995). Learning about categories in the absence of memory. *Proceedings of the National Academy of Sciences*, **92**, 12470-12474.

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.

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.

ZAKI, S. R. (2004). Is categorization performance really intact in amnesia? A meta-analysis. *Psychonomic Bulletin & Review*, **11**, 1047-1060.

ZAKI, S. R., & NOSOFSKY, R. M. (2001). A single-system interpretation of dissociations between recognition and categorization in a task involving object-like stimuli. *Cognitive, Affective, & Behavioral Neuroscience*, **1**, 344-359.

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

NOTES

1. We should note that Ashby et al. (2003) also tested a hand switch condition, in which response locations remained the same at time of transfer but subjects reversed the hands that they used for executing the motor responses. Little if any interference was observed in this condition for either category structure. Thus, it appears that processes related to response selection, rather than to motor execution, were the key ones that mediated the interference effects. (For similar results in another paradigm, see, e.g., Hommel, 1993.) Because Ashby et al. (2003) found that only the button switch condition led to interference effects, we focused only on the button switch condition in the present research.

2. On the basis of pilot work, we decided to train using slightly fewer blocks in the unidimensional condition than in the diagonal condition. The unidimensional condition is extremely easy, and we feared that the subjects would become bored and unmotivated if they trained too long. On the other hand, we wished to train for a sufficiently long time that a procedurally learned motor habit might develop for indicating the category labels with given buttonpresses. Our decision to train with eight blocks of trials represents a compromise between these competing factors.

3. Use of the one-tailed test is justified because the prior prediction was that the magnitude of the interference effect would increase as the response deadline was made more stringent.

4. Although Maddox, Bohil, and Ing (2004) tried to address this issue by removing the memory load requirements of the various questions in another version of this task, other forms of difficulty may have been at work. For example, feedback may be more difficult to interpret in the *yes–no* condition than in the A–B condition.

APPENDIX
Application of Simple Decision Boundary Models to the Conditions in Experiment 1

Following the procedures described by Ashby et al. (2003), we fitted the rule-based and information integration decision boundary models to the data from our Experiment 1. We used maximum likelihood as the criterion of fit and compared the fits of the models by using the Bayesian information criterion (BIC) statistic (Schwarz, 1978), which penalizes a model for its number of free parameters. The rule-based model uses two free parameters: a decision criterion on the relevant dimension and the variance of the perceptual-criterial noise. The information integration model uses three free parameters: the slope and intercept of the linear decision boundary and the variance of the perceptual-criterial noise. In Table A1 we report the percentage of subjects whose data were better fit by the rule-based model than by the information integration model in each of the conditions. As has been described in the text, the rule-based model provided better BIC fits in all conditions involving the unidimensional categorization, whereas the information integration model yielded better BIC fits in all conditions involving the diagonal categorization. However, the rule-based model yielded better BIC fits for 20.8% of the subjects in the transfer phase of the button switch condition of the diagonal categorization.

Table A1
Percentages of Subjects Whose Data Were Best Fit by the Rule-Based Model During the Final Two Blocks of Training and During the Two Blocks of Transfer in Experiment 1

Category Structure	Response Instruction	Training	Transfer
Unidimensional	Control	95.2	85.7
	Button switch	85.7	90.5
Diagonal	Control	4.8	4.8
	Button switch	8.3	20.8