Journal of Experimental Psychology:
Learning, Memory, and Cognition

0278-7393/13/$12.00 DOI: 10.1037/a0031670

Category Number Impacts Rule-Based *and* Information-Integration Category Learning: A Reassessment of Evidence for Dissociable Category-Learning Systems

Roger D. Stanton
St. Mary's College of Maryland

Robert M. Nosofsky
Indiana University

Researchers have proposed that an explicit reasoning system is responsible for learning rule-based category structures and that a separate implicit, procedural-learning system is responsible for learning information-integration category structures. As evidence for this multiple-system hypothesis, researchers report a dissociation based on category-number manipulations in which rule-based category learning is worse when the category is composed of 4, rather than 2, response categories; however, information-integration category learning is unaffected by category-number manipulations. We argue that within the reported category-number manipulations, there exists a critical confound: Perceptual clusters used to construct the categories are spread apart in the 4-category condition relative to the 2-category one. The present research shows that when this confound is eliminated, performance on information-integration category learning is worse for 4 categories than for 2 categories, and this finding is demonstrated across 2 different information-integration category structures. Furthermore, model-based analyses indicate that a single-system learning model accounts well for both the original findings and the updated experimental findings reported here.

Keywords: category learning, classification, implicit and explicit learning, dissociation, multiple and single systems

Supplemental materials: http://dx.doi.org/10.1037/a0031670.supp

An important debate in the perceptual classification literature concerns whether category learning is best described in terms of a single representational system or in terms of multiple systems (e.g., Ashby, Alfonso-Reese, Turken, & Waldron, 1998; Erickson & Kruschke, 1998; Love & Gureckis, 2007; Newell, Dunn, & Kalish, 2011; Nosofsky & Johansen, 2000; Nosofsky & Zaki, 1998; Smith, 2008). A variety of models that posit single-system representational formats exist; however, most multiple-system proposals assume that category learning is governed by an explicit-reasoning system and at least one implicit system.

Ashby et al.'s (1998) COVIS (Competition between a Verbal and an Implicit System) model is representative of the class of models that posit multiple systems. According to COVIS, category learning is accomplished by an explicit system that formulates and tests hypotheses based on verbalizable rules and an implicit system that is rooted in procedural learning. Furthermore, the processing assumptions of COVIS are motivated by the functions of underlying neurobiological structures. Whereas learning in the explicit system is mediated largely within the anterior cingulate, prefrontal cortex, and head of the caudate nucleus, learning in the implicit system is dominated by the tail of the caudate nucleus. Learning in the implicit system involves linking clusters of visual cortical cells with a category label and learning to associate this linkage with the correct motor response. Large groups of cells from visual cortex project to single cells in the striatum, which includes the caudate nucleus. Essentially, a map of the perceptual space is represented by the striatal units, and as experience with the task increases, each striatal unit becomes associated with a particular categorization response.

Numerous studies have attempted to verify COVIS, and a major source of supporting evidence has been reported dissociations of classification learning involving two basic categorization tasks—*rule-based* and *information-integration* (for reviews, see Ashby & Maddox, 2005; Maddox & Ashby, 2004;). Rule-based tasks are ones in which the observer can easily verbalize the optimal rule for classifying the stimuli. Rule-based category structures allow the observer to make separate decisions about the percept's value along each of the component dimensions and then combine these decisions at a postdecisional stage to make a classification response. Maddox, Filoteo, Hejl, and Ing's (2004) unidimensional category structure is an example of a rule-based task (see Figure 1a). The stimuli are lines varying in length and angle, and the vertical boundary depicts the optimal boundary for classifying the stimuli. In the Figure 1a example, only the length of the line determines category membership, and the angle of the line is

Roger D. Stanton, Department of Psychology, St. Mary's College of Maryland; Robert M. Nosofsky, Department of Psychological and Brain Sciences, Indiana University.

Correspondence concerning this article should be addressed to Roger D. Stanton, Department of Psychology, St. Mary's College of Maryland, 18952 East Fisher Road, St. Mary's City, MD 20686. E-mail: rdstanton@smcm.edu

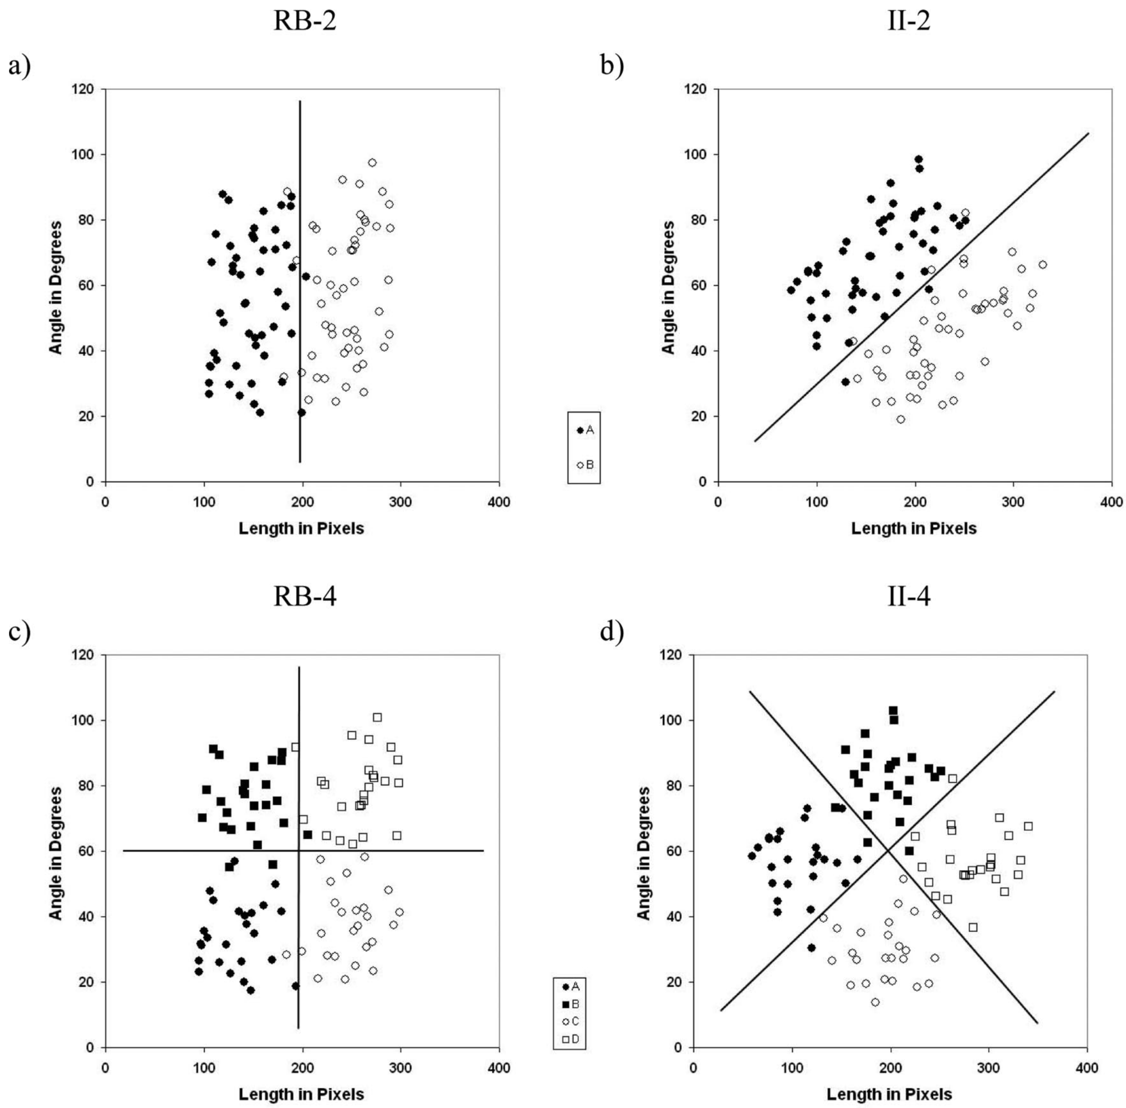

Figure 1. Two rule-based and two information-integration category structures and the corresponding optimal decision boundaries for partitioning the space. RB-2 = two-category rule-based task; RB-4 = four-category rule-based task; II-2 = two-category information-integration task; II-4 = four-category information-integration task.

irrelevant. If the percept falls to the left of the decision bound, it belongs to Category A, whereas if it falls to the right, it belongs to Category B. The hallmark of a rule-based task is that the optimal boundaries are orthogonal to the coordinate axes. For the rule-based task illustrated in Figure 1a, an observer might formulate a verbal rule that corresponds to "Category A lines are short, and category B lines are long."

An example of an information-integration category structure is Maddox, Filoteo, et al.'s (2004) diagonal category structure, depicted in Figure 1b. The diagonal boundary represents the optimal decision bound for this category structure. In an information-integration categorization task, obtaining optimal accuracy requires that the observer combine perceptual information from at least two of the dimensions prior to making his or her response (Ashby & Gott, 1988). Importantly, the optimal boundary for an information-integration categorization problem is difficult or impossible to verbalize. Although it is possible to define a rule such as "Respond A if the angle of the line is greater than its length," this description has little psychological validity because angle and length are not directly comparable.

According to the processing assumptions outlined in COVIS, rule-based tasks are learned predominately by the explicit system, whereas information-integration tasks are learned predominately by the implicit system. Most of the empirical support for the separate explicit and implicit systems proposed in COVIS has involved reported dissociations of rule-based and information-integration categorization performance. In particular, numerous studies have reported experimental manipulations that affect learning of rule-based tasks but not information-integration tasks, or vice versa (for a review, see Maddox & Ashby, 2004).

Although the results of these dissociation studies are consistent with the predictions from COVIS, single-system theorists (e.g., Dunn, Newell, & Kalish, 2012; Lewandowsky, Yang, Newell, & Kalish, 2012; Newell, Lagnado, & Shanks, 2007; Nosofsky, Stanton, & Zaki, 2005; Stanton, 2013; Stanton & Nosofsky, 2007) have argued that many of the dissociations may stem from certain extraneous variables that are not held constant across the rule-based and information-integration tasks. When these extraneous variables are controlled, the reported dissociations are often eliminated and in some cases reversed (we provide a brief review in the General Discussion). The purpose of the present research was to address a fundamentally new type of dissociation between rule-based and information-integration category learning that multiple-system theorists claim provides a major challenge to single-system models.

Maddox, Filoteo, et al. (2004) reported a dissociation in which manipulating the number of response categories impacted rule-based category learning but not information-integration learning. Because COVIS posits that learning in the explicit system requires working memory and executive attention, Maddox, Filoteo, et al. reasoned that increasing the number of response categories in a rule-based task should lower overall performance. In particular, increasing the number of categories would require the observer to form additional decision bounds, which would place increased demand on working memory and executive attention. For example, the rule-based task in Figure 1a is actually composed of four separate bivariate normal distributions that, according to Maddox, Filoteo, et al., constitute four distinct perceptual clusters. Only the length of the line is relevant for the two-category task depicted in Figure 1a, so only a single decision boundary is needed. However, for the four-category rule-based task in Figure 1c, both length and angle are relevant, and adding a second decision bound is needed to partition the perceptual space into the four response regions. According to Maddox, Filoteo, et al., COVIS predicts that this additional decision bound should render the four-category rule-based task (RB-4) more difficult to learn than the two-category rule-based task (RB-2) due to the increased demand on working memory and attention.

By contrast, Maddox, Filoteo, et al. (2004) reasoned that as long as the number of perceptual clusters is maintained, adding decision bounds to an information-integration task will not impact the learning difficulty of the task because the implicit system relies only on linking clusters of cortical cells to category responses. Based on this reasoning, Maddox, Filoteo, et al. hypothesized that the two-category information-integration task (II-2), depicted in Figure 1b, would yield equivalent learning performance to the four-category information-integration task (II-4), depicted in Figure 1d, because both category structures are composed of the same four perceptual clusters.

Across several experiments, Maddox, Filoteo, et al. (2004) demonstrated that increasing the number of response categories for a variety of rule-based tasks lowered overall accuracy. Interestingly, however, accuracy on the II-2 and II-4 tasks was roughly equal across the two levels of category number.[1] Maddox, Filoteo, et al. argued that this striking dissociation was in accord with the a priori predictions from the multiple-system COVIS model. Indeed, as we explain below, under the types of conditions outlined by Maddox, Filoteo, et al., many influential single-system models would tend to predict deleterious effects of category number on information-integration category learning. Thus, on the surface, it appears that Maddox, Filoteo, et al.'s results pose a fundamental challenge to such models.

In deriving their prediction, Maddox, Filoteo, et al. (2004) emphasized that it was important to control other extraneous factors that might vary with category number, including optimal accuracy, the number of stimulus clusters, and the distributional parameters within each cluster (such as within-cluster scatter). Unfortunately, despite controlling these extraneous factors, the central thesis of our present work is that Maddox, Filoteo, et al.'s category-number dissociation resulted from the presence of a major confound across the II-2 and II-4 tasks. In particular, to equate maximum accuracy across tasks, Maddox, Filoteo, et al. increased the distance between clusters in the four-category condition so that the maximum accuracy that could be achieved in both the two-category and four-category conditions was 95%.[2] That is, they simply spread apart the component clusters in the four-category condition. It is therefore the case that in Maddox, Filoteo, et al.'s design, the variable of category number was confounded with the variable of *between-cluster distance*. In effect, these two variables were pitted against one another, making it difficult to ascertain the effects of each variable considered separately. Furthermore, in our view, because the variable of between-class similarity is generally recognized as one of the most important controlling variables in category learning (e.g., Medin & Schaffer, 1978; Nosofsky, 1986; Rosch, Mervis, Gray, Johnson, & Boyes-Brehm, 1976; Smith & Medin, 1981), this confound is likely to be an extremely serious one.

Importantly, were it not for the confound involving between-cluster distance, then under the types of conditions outlined by Maddox, Filoteo, et al. (2004), many influential single-system models would indeed tend to predict effects of category number on information-integration category learning. To take an example, consider Nosofsky's (1984, 1986) *generalized context model* (GCM). According to the GCM, people represent categories by storing individual exemplars of the categories in memory and classify objects on the basis of their similarity to the stored exemplars. Before proceeding, we should clarify that besides being a single-representation system model, the GCM provides an example of a model that implements an information-integration strategy of classification. That is, perceptual information from multiple dimensions is combined (integrated) in the form of whole exemplars, and classification decisions are based on global similarity comparisons to these exemplars. Indeed, for the types of information-integration tasks illustrated in Figure 1, the model

[1] Because chance performance is lower in the II-4 condition (25%) than in the II-2 condition (50%), Maddox, Filoteo, et al. (2004) predicted and observed that early in training accuracy would be lower in the II-4 condition.

[2] It is unclear whether the 95% optimal performance criterion imposed by Maddox, Filoteo, et al. (2004) was with respect to the theoretical stimulus population or the specific sample on which subjects were trained. Throughout the present research, we interpret *optimal accuracy* as that which can be produced by an ideal observer who uses linear decision boundaries to classify the stimuli in a given sample. A single linear decision boundary partitions the space in the II-2 condition, and a pair of linear decision boundaries partitions the space in the II-4 condition. We consider alternative definitions and boundary conditions on optimal accuracy in our General Discussion.

tends to predict patterns of responding that are similar to those produced by the diagonal boundaries depicted in the figure. We should also clarify that whereas COVIS theorists presume that the information-integration component of their multiple-system model is fully implicit, the GCM does not make strong claims regarding the implicit versus explicit nature of the stored exemplars. We discuss this latter issue in greater detail in our General Discussion.

Consider the GCM's predictions of the effects of category number on learning in the perceptual-cluster paradigm conducted by Maddox, Filoteo, et al. (2004). Suppose there are four perceptual clusters, denoted A1, A2, B1, and B2. Suppose further that in a four-category task, each cluster defines its own category. By contrast, in a two-category task, Clusters A1 and A2 are merged to define Category A, whereas Clusters B1 and B2 are merged to define Category B. Assume that the locations of the perceptual clusters are held fixed across the tasks (unlike in the Maddox, Filoteo, et al., 2004, design) and that test item *i* from Cluster A1 is presented. In a baseline version of the GCM, the probability that item *i* is correctly classified into Category A1 in the four-category task is given by

$$P(\mathrm{A1}|i) = \frac{S_{\mathrm{A1}}}{S_{\mathrm{A1}} + S_{\mathrm{A2}} + S_{\mathrm{B1}} + S_{\mathrm{B2}}}, \tag{1a}$$

where S_{A1} denotes the summed similarity of item *i* to the exemplars of Cluster A1 (and likewise for S_{A2}, S_{B1}, and S_{B2}). By comparison, in the two-category task, the probability that item *i* is correctly classified in Category A is given by

$$P(\mathrm{A}|i) = \frac{S_{\mathrm{A1}} + S_{\mathrm{A2}}}{S_{\mathrm{A1}} + S_{\mathrm{A2}} + S_{\mathrm{B1}} + S_{\mathrm{B2}}}. \tag{1b}$$

As discussed in more detail in the Modeling Analyses section of this article, in the baseline model, the summed similarities S_{A1}, S_{A2}, S_{B1}, and S_{B2} are invariant across the two-category and four-category information-integration tasks. Thus, comparing Equations 1a and 1b, it is easily seen that correct classification in the two-category information-integration task would always be greater than in the four-category one.

Although these predictions stem from an extremely simple, baseline version of the GCM, the core ingredients described above continue to exert a major influence when more elaborate versions of the model are applied. In particular, holding the locations of the perceptual clusters fixed, the model will often predict more accurate classification performance for two-category information-integration tasks than for four-category ones because the summed similarity of test items to their target category will be greater in the two-category case. Of course, if one does *not* hold stimulus conditions constant across tasks and instead spreads farther apart the perceptual clusters in the four-category case (as Maddox, Filoteo, et al., 2004, did in their design), then superior performance for the two-category task may no longer be observed. For example, if the values S_{B1} and S_{B2} are made arbitrarily small by spreading apart the B clusters from the A clusters in the four-category task, then extremely good performance on the four-category task can be achieved. The precise predictions from the model will depend on the detailed structure of the experimental design. Although we have described these predictions for only one representative from the class of single-system models, we believe that similar principles operate for other members from this broad class, including prototype models and clustering models (e.g., Love & Gureckis, 2007; Reed, 1972).

The goal of the present research is to investigate whether the dissociation reported by Maddox, Filoteo, et al. (2004) does indeed challenge single-system views of category learning. We pursue this goal in two ways. First, we examine the role of increasing the number of response categories in information-integration category-learning tasks while avoiding the confound of increased between-cluster distance. Our plan is to test a variety of information-integration structures, using procedures that parallel the ones adopted by Maddox, Filoteo, et al. but in which the locations of the perceptual clusters are held fixed. Under such conditions, we predict that performance will often be worse in the four-category tasks than in the two-category tasks. Such demonstrations would provide a strong challenge to Maddox, Filoteo, et al.'s hypothesis that information-integration category learning is not affected by category number and would remove another source of dissociation evidence that has been taken in support of multiple-system models compared to single-system ones. Indeed, to the extent that Maddox, Filoteo, et al.'s reasoning about the predictions from the COVIS model is correct, such effects of category number on information-integration category learning would appear to constitute a strong challenge to COVIS.

Second, beyond providing these new experimental findings, we also conduct modeling-based investigations of whether the Maddox, Filoteo, et al. (2004) results challenge single-system views of category learning. To anticipate, we demonstrate that a simple version of a single-system learning model (namely, the GCM) accounts well for *both* the original Maddox, Filoteo, et al. information-integration category-learning data and the new data reported herein. Taken together, the empirical evidence and the single-system modeling account of the results place the Maddox, Filoteo, et al. category-number dissociations in a new light and indicate that they do not challenge the general class of single-system category-learning models.

Experiment 1

Our initial experiment sought to replicate the Maddox, Filoteo, et al. (2004) results. Such replication was important in order to ensure correspondence between this study and the Maddox, Filoteo, et al. study. Furthermore, this replication provides a baseline with which to compare the results obtained in the experiments reported later in this article. Because the results pertaining to the rule-based categories tested by Maddox, Filoteo, et al. are not in question nor in opposition to the predictions from single-system models, we sought to replicate the results obtained for only the II-2 and II-4 tasks. In Experiment 1, we followed the exact procedure used by Maddox, Filoteo, et al.

Method

Participants. Forty observers participated in this experiment in partial fulfillment of an introductory psychology course requirement. There were 20 participants in the II-2 condition and 20 in the II-4 condition. The participants were instructed that they could earn a $4 bonus for good performance in the task. If the observer achieved 80% accuracy in the final block, he or she earned the bonus; however, the participants were unaware of this criterion and were instructed that they would earn the bonus if they achieved "good performance."

Stimuli. There were 100 stimuli in both conditions. In the II-2 condition, there were 50 stimuli in Categories A and B, and in the II-4 condition, there were 25 stimuli in each of Categories A, B, C, and D. The stimuli were lines varying in their length and angle, and the procedure used for generating the stimuli was the same as used by Maddox, Filoteo, et al. (2004). The II-2 stimuli were generated by randomly sampling from the bivariate normal distributions defined by the parameters listed in Table 1. The II-4 stimuli were created by shifting each of the 100 II-2 stimuli either 13 or 14 units in the appropriate direction; thus, it was not necessary to resample the distributions when creating the II-4 category structure. The underlying population parameter values for the four-category condition are also listed in Table 1. The population parameters are identical to those used by Maddox, Filoteo, et al. Both the II-2 and II-4 conditions were constrained such that an ideal observer using linear boundaries could achieve 95% correct in both conditions. The population distributions were sampled until both the II-2 and II-4 samples met this criterion. Each point in the distribution was transformed into a line by converting the *x* value to a line length (specified in pixels) and the *y* value to an angle (given in radians). In accord with Maddox, Filoteo, et al., the scaling factor $\pi/500$ was applied to the *y* value so that angle and length were of roughly equal salience. The length values of the stimuli were increased by 50 units so that all stimuli were at least 50 pixels in length. The structures depicted in Figure 1 are based on these length-shifted values. The stimuli were displayed as white lines drawn on a black background and were presented on a 15-in. monitor in a dimly lit room.

Procedure. Category number was a between-subjects variable. The instructions informed participants of the number of categories in their task (either two or four). For both conditions, a fixation appeared for 1 s, and then the stimulus appeared. After the participant provided a categorization response, corrective feedback was provided for 1 s. In accord with the methods used by Maddox, Filoteo, et al. (2004), the correct category label was presented along with the word *wrong* for incorrect responses and the word *right* for correct responses. This was followed by a 500-ms black screen before the next trial began. There were six blocks of trials with all 100 stimuli presented once per block. The order of presentation of the stimuli was randomized anew in each block for each individual participant. Self-paced rest breaks were allowed between each block.

Table 1
Parameter Values Used to Construct the Category Distributions in Experiments 1–3

Distributions	μ_L	μ_A	σ_L^2	σ_A^2	σ_{LA}
II-2, II-2(fixed), II-4(fixed)					
Category A_1	80	150	900	900	0
Category A_2 (B)	150	220	900	900	0
Category B_1 (C)	150	80	900	900	0
Category B_2 (D)	220	150	900	900	0
II-4					
Category A	66	150	900	900	0
Category B	150	233	900	900	0
Category C	150	66	900	900	0
Category D	233	150	900	900	0
Diag-2, Diag-4					
Category A_1	108	167	1,296	1,296	1,275
Category A_2 (B)	136	139	1,296	1,296	1,275
Category B_1 (C)	164	111	1,296	1,296	1,275
Category B_2 (D)	192	83	1,296	1,296	1,275

Note. Condition II-4 of Experiment 1 is sometimes referred to in the text as II-4(spread). μ_L = mean length; μ_A = mean angle; σ_L^2 = variance length; σ_A^2 = variance angle; σ_{LA} = covariance length-angle; II-2 = two-category information-integration task; II-4 = four-category information-integration task; Diag-2 = two-category diagonal information-integration task; Diag-4 = four-category diagonal information-integration task.

Results and Discussion

The mean proportion of correct responses is displayed in Figure 2 as a function of category number (two or four) and block. The pattern of results mirrors that obtained by Maddox, Filoteo, et al. (2004). As in the Maddox, Filoteo, et al. study, performance in the II-4 task was initially worse due to lower chance performance, but performance was roughly equal by the third learning block. A 2 (category number) × 6 (learning block) analysis of variance (ANOVA) was performed on the accuracy data. The main effect of blocks was significant, $F(5, 190) = 94.561$, $MS_E = .003$, $p < .001$, reflecting that performance improved across the blocks. There was no significant difference in II-2 and II-4 learning performance, $F(1, 38) = 0.828$, $MS_E = .054$, $p = .369$. However, the Block × Category Number interaction was significant, $F(5, 190) = 16.119$, $MS_E = .003$, $p < .001$, likely reflecting the performance advantage in early blocks for the II-2 condition. To verify that a performance advantage in the early blocks was the source of the interaction, we tested the simple main effect of category number at each level of block. The simple main effects analysis revealed that accuracy on II-2 was significantly higher than II-4 for the first block, $F(1, 38) = 15.851$, $p < .001$, but that there was not a significant difference in II-2 accuracy and II-4 accuracy at any other level of block (all $ps > .20$).

The results of Experiment 1 replicate the main finding of Maddox, Filoteo, et al. (2004) that learning performance in the later blocks does not differ between the II-2 and II-4 conditions. This pattern of data is consistent with the Maddox, Filoteo, et al. hypothesis that as long as the number of perceptual clusters is maintained, adding response categories to an information-integration task does not affect the learning difficulty of the task. As stated previously, however, it is unclear what aspect of the task produced this pattern of results because category number is confounded with between-cluster distance in this design. Therefore, we now test designs in which between-cluster distance is held fixed across tasks, while category number is varied. Under such conditions, we hypothesize that performance in the four-category tasks will be worse than in the two-category tasks.

Experiment 2

In Experiment 2, we test two information-integration conditions based on a resampling of the same population parameters that were used to generate the II-2 condition from Experiment 1. Crucially, in the current design, the locations of the four bivariate normal distributions (i.e., the perceptual clusters) are held fixed across the

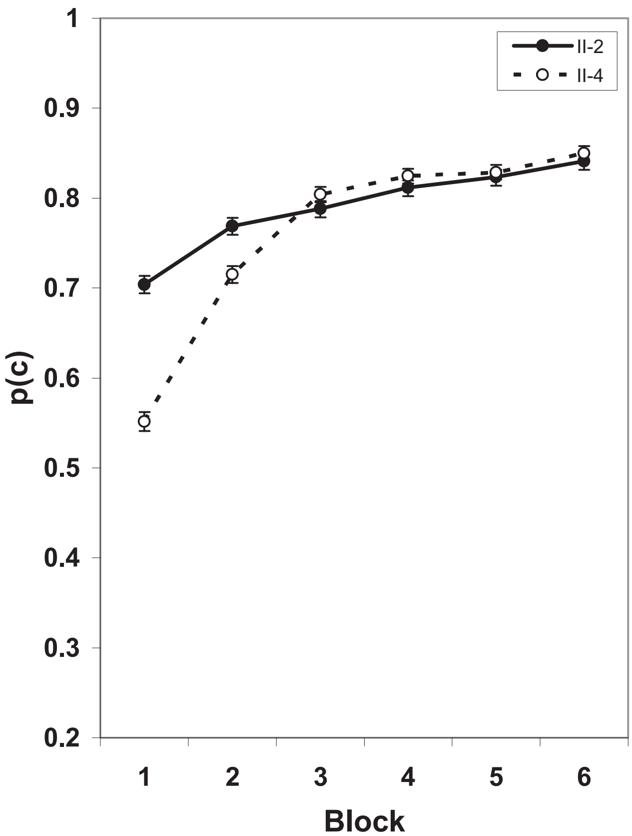

Figure 2. Experiment 1: Mean probability of correct classifications as a function of category number and learning block. Error bars are within-block standard errors pooled across subjects. II-2 = two-category information-integration task; II-4 = four-category information-integration task.

two-category and four-category tasks. To avoid the need to increase the between-cluster distance across tasks, we constrained the sampled stimulus configuration such that an ideal observer using linear decision boundaries could achieve 100% accuracy in both the two-category and four-category conditions. For ease of discussion, we refer to these modified category structures in which the clusters are held fixed across conditions as the II-2(fixed) and II-4(fixed) conditions. By contrast, we refer to the Experiment 1 four-category condition, in which the perceptual clusters were spread apart, as II-4(spread).

The II-2(fixed) and II-4(fixed) categories are illustrated together in Figure 3. For the II-2(fixed) condition, the solid symbols represent the Category A stimuli, the open symbols represent the Category B stimuli, and the solid diagonal boundary is the optimal linear boundary for classifying the stimuli. In the II-4(fixed) condition, each of the four categories is represented by a unique shading/shape combination (e.g., Category A stimuli are solid circles), and the solid and dashed boundaries are the optimal linear boundaries.

Importantly, note that the present design was generated in a manner that is closely parallel to the one used in the original Maddox, Filoteo, et al. (2004) study. In particular, there are four perceptual clusters defined by bivariate normal distributions. Each cluster forms its own category in the four-category task, whereas two adjacent pairs of clusters are each merged to form categories in the two-category task. The boundary conditions for the category-number manipulation that were deemed important by Maddox, Filoteo, et al. (2004, p. 229) are all satisfied in the present design. In particular, both the two-category and four-category structures are composed of four perceptual clusters, within-cluster scatter is held fixed across conditions, the distributional parameters within each cluster are held fixed, and maximum accuracy for an observer using linear decision boundaries is equated across conditions. Thus, according to Maddox, Filoteo, et al.'s stated reasoning, the multiple-system model predicts that performance in the later learning blocks should be the same across the two-category and four-category tasks.

Method

Participants. Forty-eight observers participated in this experiment in partial fulfillment of an introductory psychology course requirement. There were 24 participants in the II-2(fixed) condition and 24 in the II-4(fixed) condition. All other aspects pertaining to the participants were the same as in Experiment 1.

Stimuli. The stimuli were again lines varying in length and angle. The category structures were generated by sampling from four bivariate normal distributions. The parameter values for these distributions are listed in Table 1, and the resulting category structures are depicted in Figure 3. The sampled distributions were constrained such that an observer using linear decision boundaries could achieve 100% accuracy in both conditions. The same stimulus values were used in both the II-2(fixed) and II-4(fixed) conditions. There were 25 stimuli sampled from each of the four bivariate normal distributions. In the II-2(fixed) condition, there

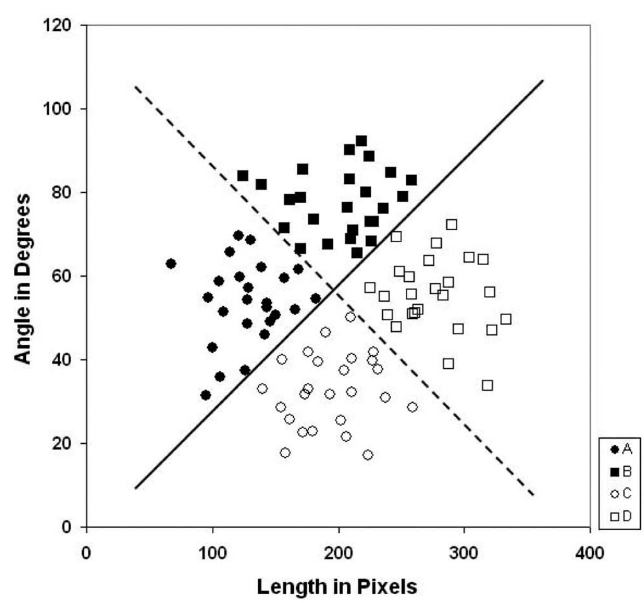

Figure 3. Experiment 2: The II-2(fixed) and II-4(fixed) category structures and their corresponding decisions boundaries. II-2 = two-category information-integration task; II-4 = four-category information-integration task.

were 50 stimuli in each of Categories A and B, and in the II-4(fixed) condition, there were 25 stimuli in each of Categories A, B, C, and D.

Procedure. The procedure was identical to the one used in Experiment 1.

Results and Discussion

The mean proportion of correct responses is displayed in Figure 4 as a function of category number (two or four) and block. In contrast to the results of Experiment 1, it is clear that performance in the two-category condition was better than performance in the four-category condition. To verify this observation, a 2 (category number) × 6 (learning block) ANOVA was performed on the accuracy data. The main effect of block was significant, $F(5, 230) = 51.347$, $MS_E = .004$, $p < .001$, indicating that accuracy was increasing with additional learning blocks. Performance in the II-4(fixed) condition was significantly worse than performance in the II-2(fixed) condition, $F(1, 46) = 32.639$, $MS_E = .063$, $p < .001$. The Category Number × Block interaction was significant, $F(5, 230) = 8.173$, $MS_E = .004$, $p < .001$, reflecting greater improvement across the early learning blocks in the II-4(fixed) condition compared to the II-2(fixed) condition. It is clear from inspection, however, that performance in the II-4(fixed) condition remained worse than performance in the II-2(fixed) condition during the later learning blocks.

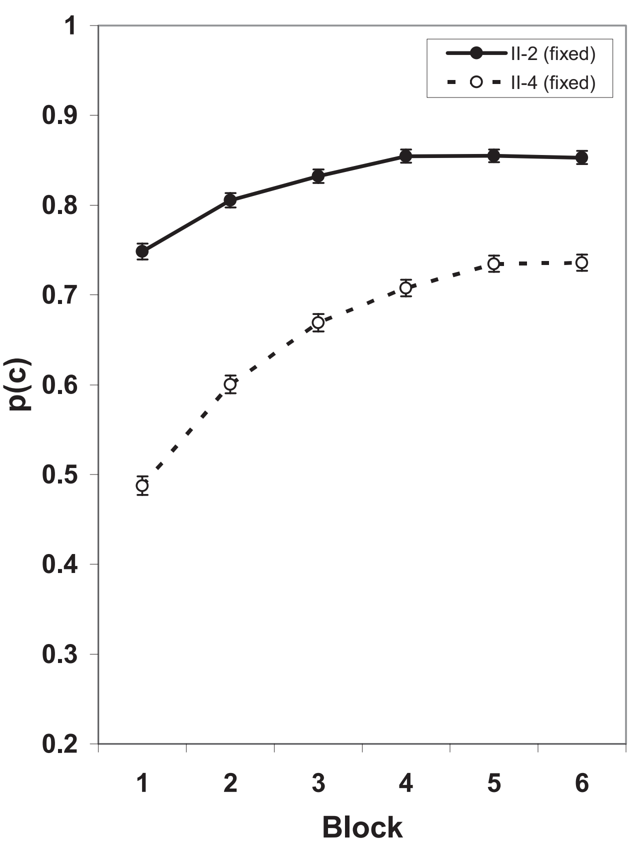

Figure 4. Experiment 2: Mean probability of correct classifications as a function of category number and learning block. Error bars are within-block standard errors pooled across subjects. II-2 = two-category information-integration task; II-4 = four-category information-integration task.

To confirm that the pattern of results from Experiment 2 differs significantly from the one observed in Experiment 1, we conducted a three-way ANOVA using as factors experiment, category number, and learning block. The most important result was that there was a significant two-way interaction between experiment and category number, $F(1, 84) = 11.17$, $MS_E = .059$, $p < .001$. This result reflects the much larger effect of category number that we observed in Experiment 2 compared to Experiment 1.

In sum, in experimental conditions that directly parallel the ones generated by Maddox, Filoteo, et al. (2004), we found that the variable of category number significantly impacted the learning of information-integration category structures. Furthermore (based on the stated reasoning from Maddox, Filoteo, et al., 2004, p. 229), we observed these effects of category number under conditions in which the implicit, procedural-learning system of COVIS predicts they should have been absent.

The results suggest strongly that Maddox, Filoteo, et al. (2004) failed to find an advantage for their II-2 task compared to their II-4(spread) task because they increased the distance between stimulus clusters in the II-4(spread) condition.

Experiment 3

The purpose of Experiment 3 was to provide another test of whether category number impacts information-integration category learning by using another category structure. To maintain consistency with the Maddox, Filoteo, et al. (2004) procedures, we sought to test a structure that is closely related to the sort of tasks used by these researchers.

Recall that in our introduction, we reviewed two rule-based category-learning tasks tested by Maddox, Filoteo, et al. (2004)—see Figures 1a and 1c. The RB-4 task required that the observer attend to both stimulus dimensions; however, there was only a single relevant dimension in the RB-2 task. As a result, it was possible that increased attentional demands, rather than a greater number of response categories, caused the lower performance in the RB-4 task. To assess the effects of category number on rule-based categorization while holding the number of relevant dimensions fixed, Maddox, Filoteo, et al. (Experiment 3) tested two-category and four-category unidimensional rule-based tasks (Uni-2 and Uni-4). A slightly modified version of these category structures is depicted in Figure 5a. For the Uni-2 condition, Category A stimuli are represented by solid symbols, Category B stimuli are represented by open symbols, and the solid vertical boundary is the optimal boundary for classifying the stimuli. Each of the four Uni-4 categories is represented by a unique shading/shape combination (e.g., Category A stimuli are solid circles), and the solid and dashed boundaries are the optimal classification boundaries. Maddox, Filoteo, et al. found that performance for Uni-2 was superior to that for Uni-4, and thus, they concluded that category number impacts rule-based category learning regardless of the number of relevant dimensions.

These unidimensional category structures provide a starting point for further testing the Maddox, Filoteo, et al. (2004) hypothesis concerning the effect of category number on information-integration tasks. In the same manner that Maddox, Filoteo, et al.

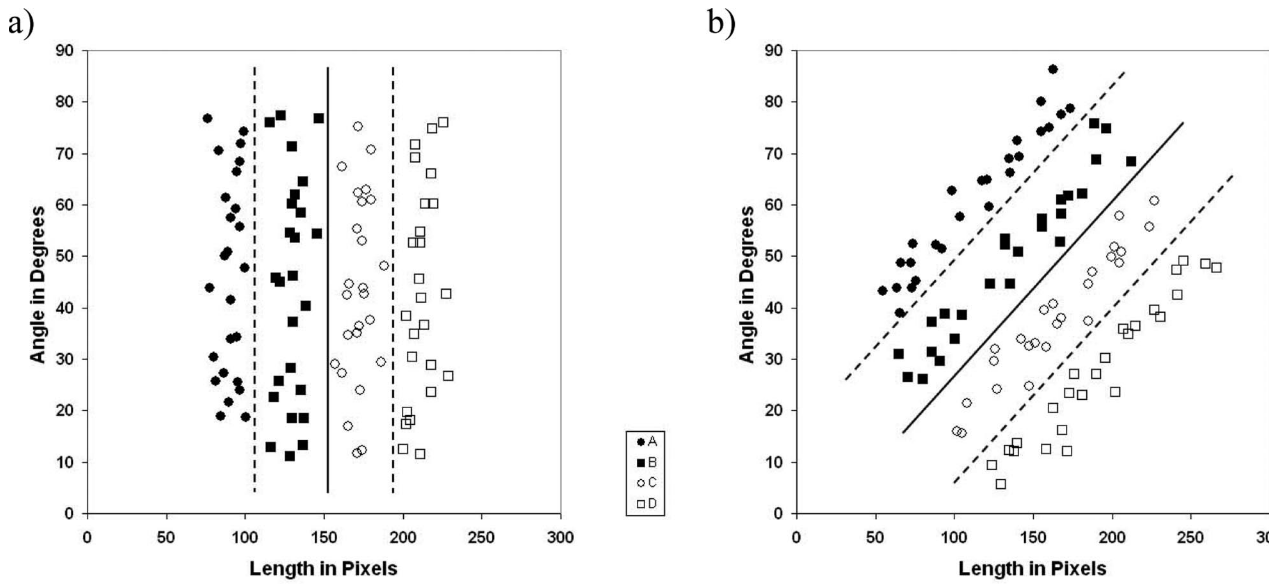

Figure 5. Experiment 3: The Uni-2 and Uni-4 category structures and the Diag-2 and Diag-4 category structures, as well as the optimal decision boundaries. Diag-2 = two-category diagonal information-integration task; Diag-4 = four-category diagonal information-integration task; Uni-2 = two-category unidimensional rule-based task; II-4 = four-category unidimensional rule-based task.

produced their II-2 and II-4 conditions by basically rotating the RB-2 and RB-4 structures by 45° (see Figure 1), diagonal information-integration category structures can be produced from the Uni-2 and Uni-4 conditions. In particular, rotating the unidimensional category structures 45° clockwise produces two- and four-category diagonal information-integration category structures (see Figure 5b). Diagonal categorization tasks have been used extensively in many of the dissociation studies reported in support of COVIS (e.g., Ashby, Ell, & Waldron, 2003; Ashby, Maddox, & Bohil, 2002; Maddox, Ashby, & Bohil, 2003). Importantly, the present stimulus configuration is composed of four separate perceptual clusters, each of which can be produced by sampling from a bivariate normal distribution. Furthermore, the configuration can be tested as either a two-category task (Diag-2), where the solid and open symbols represent the two categories, or as a four-category task (Diag-4), where each unique shading/shape combination represents a category. Finally, as was the case in Experiment 2, the distributions are sampled such that an ideal observer using linear boundaries can achieve 100% accuracy in both tasks.

Based on the reasoning outlined by Maddox, Filoteo, et al. (2004), each of the stimulus clusters qualifies as a distinct perceptual cluster because they all were sampled from separate bivariate normal distributions. Also, as was the case in our Experiment 2, the number of clusters is held fixed across the Diag-2 and Diag-4 conditions, and the other extraneous variables listed by Maddox, Filoteo, et al. are also controlled. Thus, the Maddox, Filoteo, et al. hypothesis predicts that learning performance will be roughly equal during the latter training blocks of these two information-integration conditions. By contrast, if our hypothesis is correct and category number does impact information-integration learning (when the number of categories is not confounded with changes to the physical structure of the stimulus set), then the Diag-4 condition should yield worse learning performance than the Diag-2 condition.

Method

Participants. Forty-five observers participated in Experiment 3 in partial fulfillment of an introductory psychology course requirement. Twenty-three observers participated in the Diag-2 condition, and 22 participated in the Diag-4 condition. All other aspects of the participants were identical to those of Experiments 1 and 2.

Stimuli. As in Experiments 1 and 2, the stimuli were lines varying in length and angle. The category structures were generated by sampling from four bivariate normal distributions. The parameter values for these distributions are listed in Table 1, and the category structure is depicted in Figure 5. The same stimulus values were used in both the Diag-2 and Diag-4 conditions. Across both conditions, the sampled distributions were constrained such that an ideal linear classifier (using a single boundary in the Diag-2 task and three boundaries in the Diag-4 task) could achieve 100% accuracy on the tasks. Twenty-five stimuli were sampled from each of the four bivariate normal distributions. For the Diag-2 condition, there were 50 stimuli in each of Categories A and B, and there 25 stimuli in each of the four categories tested in the Diag-4 condition.

Procedure. The procedure was identical to the one used in Experiments 1 and 2.

Results and Discussion

The mean proportion of correct responses is displayed in Figure 6 as a function of category number and learning block. It is immediately apparent that performance in the Diag-2 condition was much better than performance in the Diag-4 condition. To

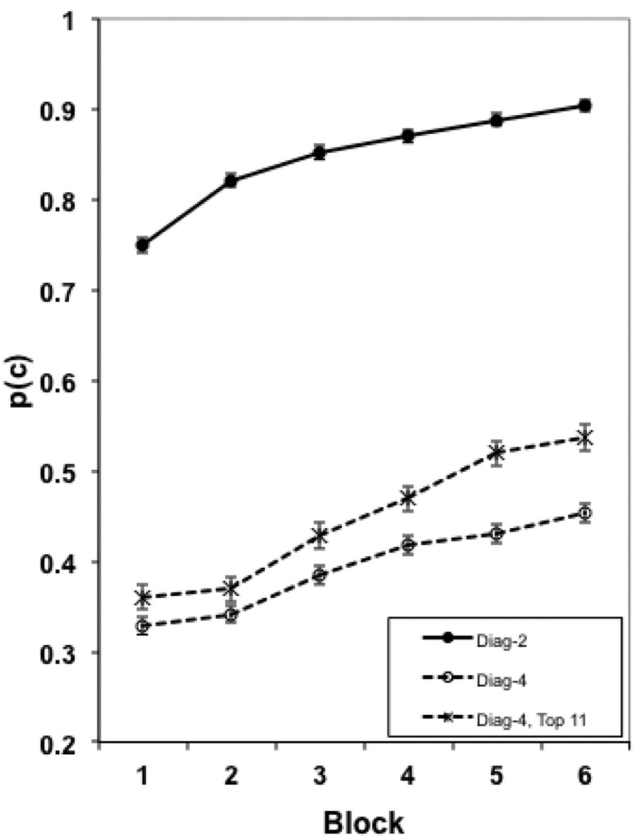

Figure 6. Experiment 3: Mean probability of correct classifications as a function of category number and learning block. Error bars are within-block standard errors pooled across subjects. Diag-2 = two-category diagonal information-integration task; Diag-4 = four-category diagonal information-integration task.

verify this observation, we performed a 2 (category number) × 6 (learning block) ANOVA on the learning data obtained in Experiment 3. The main effect of learning block was significant, $F(5, 215) = 24.784$, $MS_E = .003$, $p < .001$, indicating that accuracy was improving with additional training. Performance in the Diag-2 condition was significantly better than performance in the Diag-4 condition, $F(1, 43) = 451.789$, $MS_E = .023$, $p < .001$. The Learning Block × Category Number interaction was significant, $F(5, 215) = 13.158$, $MS_E = .003$, $p < .001$.

Performance in the Diag-4 condition was very poor, and perhaps some observers found the task too demanding and stopped trying to learn the task. If some subset of observers did stop trying and randomly guessed, this would lower average performance in this task. Possibly a subset of the observers performed quite well, but this effect was masked by the other subset that simply gave up. To further assess the ability of the observers to learn the Diag-4 task, we analyzed the accuracy data for the top 50% of performers in this condition, that is, the 11 participants with the highest average accuracy in the final three blocks. The mean proportion of correct responses for this subset of participants is displayed in Figure 6 as asterisks with a dotted connecting line. It is clear from inspection that, even restricting the analysis to this subset of participants, performance was quite poor in the Diag-4 condition.

A potential objection to our Experiment 3 design is that, despite being generated by single bivariate normal distributions, the individual categories in Figure 5b do not constitute perceptual clusters. To reiterate, in constructing the designs tested here, we followed the logic and methods set forth in the original Maddox, Filoteo, et al. (2004) article. Specifically, these researchers defined perceptual clusters in terms of bivariate normal distributions, which were sometimes circular in form and other times elongated along individual dimensions (see Figures 1 and 5a). They produced information-integration structures by rotating the space by 45°, sometimes producing categories that were elongated in diagonal directions (see, e.g., Figure 1b). In the present design, we produced the two-category and four-category tasks by using parallel methods. Future research into these issues might benefit, however, by formalizing more precisely what counts as a perceptual cluster.

In sum, the results from Experiment 3 provide another example in which a four-category information-integration task yielded worse learning performance than a two-category information-integration task formed from identical perceptual clusters. This result again challenges the Maddox, Filoteo, et al. (2004) hypothesis that, for designs with fixed perceptual clusters, the number of response categories does not impact performance on information-integration tasks. Furthermore, the results of Experiments 1–3 support the claim that Maddox, Filoteo, et al. failed to obtain an advantage for their two-category information-integration task because the variable of category number was confounded with the between-cluster distances in their design.

Model-Based Analyses

Our argument is that in Maddox, Filoteo, et al.'s (2004) original study, II-4(spread) performance was roughly equal to II-2 performance because the stimulus space was altered across the two conditions. Thus, we argue that those results do not provide evidence against the predictions from single-system models. To make this argument more convincing, it is important to demonstrate that a single-system model can in fact account for the data. Thus, in this section, we report analyses from the exemplar-based GCM (Nosofsky, 1984, 1986), which is a well-known representative from the single-system class. In our introduction, we briefly described predictions from a simple, baseline version of the model. In this section, we apply a fuller version designed to account for more of the details of the learning performance across the different conditions.

According to the GCM, people represent categories in terms of exemplars stored in memory and make classification judgments on the basis of the similarity between a test item and the stored exemplars. Each exemplar is represented as a point in psychological space, and the similarity between two items is a decreasing function of their distance in this space. Let x_{im} denote the value of exemplar i on dimension m. In the present case, the stimuli are composed of separable dimensions (Garner, 1974; Shepard, 1964); therefore, the distance between exemplars i and j is computed by using a weighted city-block distance metric

$$d(i, j) = \sum w_m |x_{im} - x_{jm}|, \tag{2}$$

where x_{im} and x_{jm} are the psychological values of items i and j on dimension m, and w_m ($0 \leq w_m \leq 1$, $\Sigma w_m = 1$) is the attention weight given to dimension m. For simplicity, in the present analyses, the psychological coordinate values of the exemplars were assumed to be identical to those in the logically defined category space (see Table 1). In addition, in these initial analyses, the attention weights were set equal to one another (i.e., $w_1 = w_2 = .5$) because both dimensions were roughly equally relevant to classification in the present information-integration tasks. The similarity of exemplar i to exemplar j is an exponential decay function of their psychological distance (Shepard, 1987) given by

$$s_{ij} = \exp(-c \cdot d_{ij}), \tag{3}$$

where c is an overall scaling parameter.

The probability of classifying item i as a member of Category A is determined by the summed similarity of the item to the Category A exemplars relative to the item's summed similarity to all exemplars in all categories,

$$P(\mathrm{A}|i) = \frac{\left(\sum_{a \in \mathrm{A}} s(i,a) + \frac{B}{N}\right)^{\gamma}}{\sum_{K}\left(\sum_{k \in K} s(i,k) + \frac{B}{N}\right)^{\gamma}} \tag{4}$$

where $\Sigma s(i, a)$ is the summed similarity of item i to all previously experienced exemplars of Category A and γ is a response scaling parameter (Ashby & Maddox, 1993; McKinley & Nosofsky, 1995). In Equation 4, B represents a fixed level of background noise (Nosofsky, Kruschke, & McKinley, 1992), and N is the number of response categories in the task. A process interpretation is that background elements exist in memory prior to the start of training. If a background element is retrieved from memory, the observer chooses randomly among the available category responses. Note that as learning proceeds, more and more category exemplars will be experienced and stored in memory. Thus, the summed similarity of presented items to category exemplars increases with additional experience. Early in the learning sequence, the level of background noise is high relative to the summed similarities, so there is a good deal of random responding. However, with increased experience, the magnitude of the summed similarities begins to outweigh the background noise, producing higher levels of accuracy as learning continues.[3]

Qualitative Predictions of Averaged Learning Data

In our initial analyses, we sought to determine the ability of the GCM to account for the main qualitative pattern of results in our Experiments 1 and 2. To do so, we fitted the model to the averaged learning data for each of the conditions. In particular, we conducted a computer search for the values of the free parameters c, γ, and B that minimized the sum of squared deviations between predicted and observed values. We constrained the model by forcing it to fit all four conditions with a single set of parameter values (best fitting parameters: $c = .052$, $\gamma = 1.30$, $B = 3.549$).

The predicted learning curves are shown alongside the observed data in Figure 7. The model does a good job of capturing the main qualitative pattern of results in these averaged data. It predicts nearly equal performance during the later blocks of training across the II-2 and II-4(spread) conditions, as was observed by Maddox, Filoteo, et al. (2004) and in the present replication (see the top panels of Figure 7). At the same time, it predicts clearly better performance for the II-2(fixed) than for the II-4(fixed) condition throughout all blocks of training (see the bottom panels of Figure 7). The quantitative fit of the model is reasonably good as well. However, because these are averaged data, the focus here is on the general qualitative pattern of predictions rather than on fine-grained quantitative fit.

In our view, if even a narrow range of parameter values from the model yields an excellent account of the qualitative pattern of results, it already begins to cast doubt that such data strongly challenge single-system approaches.[4] In the present case, however, the argument is even stronger because it turns out that these general qualitative predictions hold over a wide range of parameter values from the model. To document this point, in Table 2, we report learning-curve predictions for an assortment of parameter-value combinations. For purposes of generality, we chose different combinations of c and γ that produced asymptotic predictions of performance in the II-2 condition that ranged from approximately .70 to .85. (Because c and γ tend to trade off with one another, different combinations of c and γ can yield similar predictions of overall percentage correct scores.) Holding fixed each parameter combination, we also display learning curves predicted for conditions II-4(spread), II-2(fixed), and II-4(fixed). (In each case, a fixed value of B was chosen to produce learning, but the asymptotic predictions do not depend on B.) As can be seen from inspection of the table, the deterioration in performance for II-4(fixed) compared to II-2(fixed) is always far greater in magnitude than is any deterioration in performance for II-4 (spread) compared to II-2. Furthermore, it is often the case that predicted asymptotic performance for the II-2 and II-4(spread) conditions is nearly equal. In a nutshell, the complete pattern of information-

[3] This background-noise approach to modeling learning in the GCM was first applied by Nosofsky et al. (1992). In addition, Nosofsky and Palmeri (1997) developed a mechanistic account of categorization decision making based on a random-walk process that yields the Equation 4 decision rule as a special case. It is important to understand that the asymptotic predictions from the model (i.e., the performance levels where the learning curves flatten out) are not influenced by the background-noise construct. The background-noise version of the model is applied here simply to produce the learning curves. The present version of the learning model (involving the term B/N) assumes that there is a fixed *total* amount of background noise (divided evenly among the categories) that is independent of the number of categories to be learned. Alternative versions might assume that the total amount of background noise is proportional to the total number of categories. (Such versions would substitute the term B for B/N.) These alternative versions would produce slightly different predictions of the rate at which learning approaches asymptote across the two-category and four-category conditions (i.e., the steepness of the learning curves during the early learning blocks). Because the patterns of performance during the early learning blocks are not central to the present investigation, we made no attempt to distinguish among these more fine-grained alternatives.

[4] We acknowledge that if one model predicts a result across a broad range of parameter settings whereas another predicts the result only within a narrow range of its parameters, then the former model provides the favored explanation. In practice, however, the demonstration of interesting empirical results such as dramatic dissociations may sometimes involve a highly specific and crafted design. In that case, it may not be surprising that the second model requires highly specific parameter settings to account for the effect. For that reason, dissociations become far more convincing when they are demonstrated in very general fashion rather than under highly specified conditions. If the dissociation disappears with minor variations in procedure, then the multiple-system explanation becomes suspect.

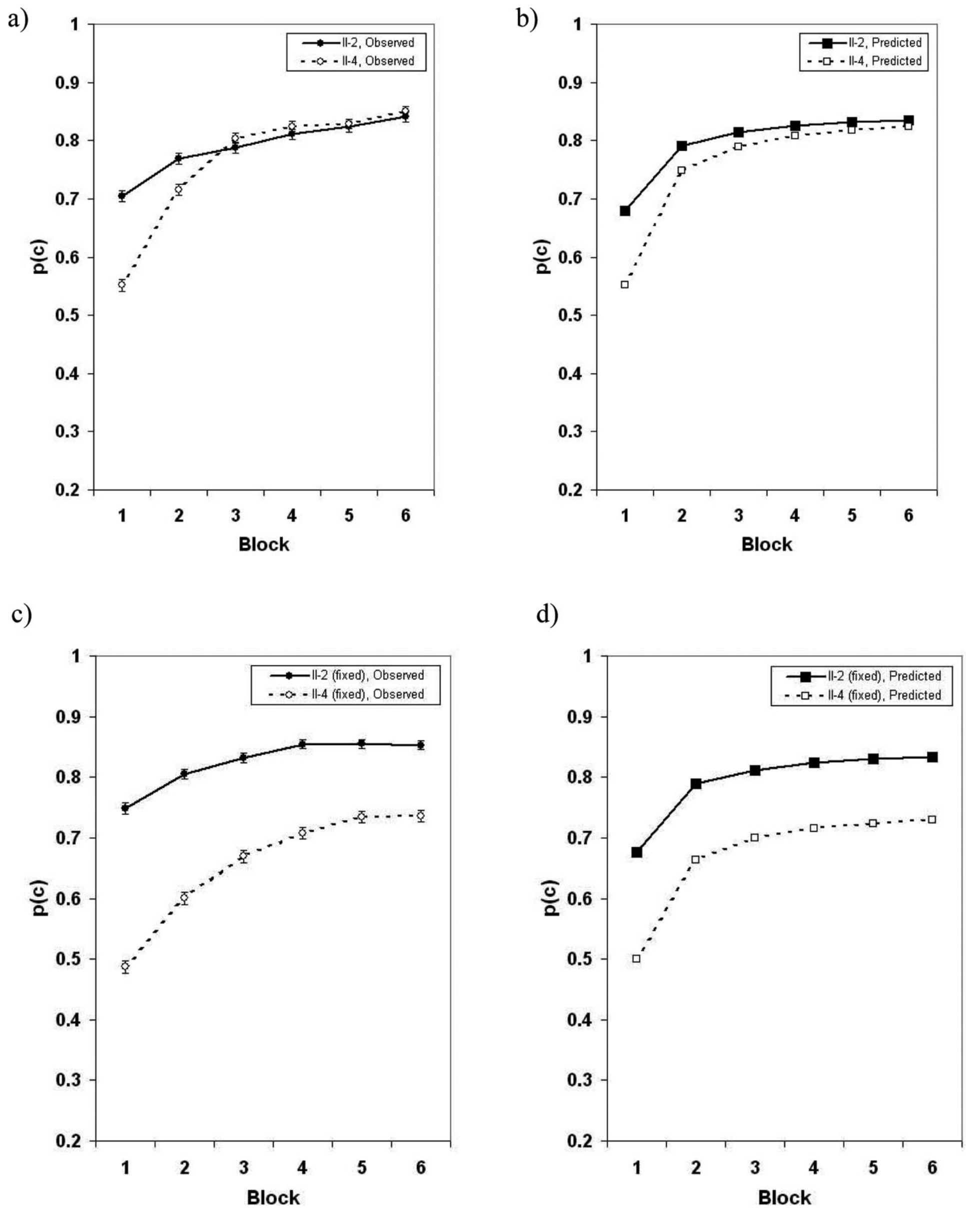

Figure 7. Observed and predicted accuracies for Experiments 1 and 2 shown as a function of category number and learning block. Error bars are within-block standard errors pooled across subjects. II-2 = two-category information-integration task; II-4 = four-category information-integration task.

integration category-learning data reported by Maddox, Filoteo, et al. (2004) and in our Experiments 1 and 2 lies solidly within the model's scope.

As a further test, we examined the model's predictions for the diagonal categorization tasks of Experiment 3. Holding fixed the best fitting parameters from the previous analysis, the predicted learning curves are shown alongside the observed learning curves in Figure 8. Although the model underpredicts the magnitude of the effect, it clearly predicts that performance in the Diag-4 condition is much worse than in the Diag-2 condition, as was observed in the data reported here. Again, this general qualitative pattern of predictions held across a wide range of parameter values, so it may be considered a true a priori prediction from the model.

The focus of the present study was on the role of category number in influencing category learning in information-integration tasks. Nevertheless, for completeness, it is of interest to also consider the GCM's predictions of learning performance for the RB-2 and RB-4 tasks that were tested by Maddox, Filoteo, et al. (2004)—see the left panels of our Figure 1 for a review of those category structures. Recall that Maddox, Filoteo, et al. observed significantly better learning for the RB-2 structure than for the RB-4 structure, despite the fact that the perceptual clusters for the RB-4 task were spread farther apart than for the RB-2 task. As is apparent from inspection of Figure 1, however, a crucial difference between the RB-2 structure compared to the RB-4, II-2, and II-4 structures is that only a single dimension is relevant for the RB-2

Table 2
Learning-Curve Predictions for Different Parameter Combinations From the Generalized Context Model

Condition	Block 1	2	3	4	5	6	Parameter value c	γ	B
				Asymptote ≈ .85					
II-2	0.682	0.802	0.827	0.839	0.847	0.851	0.070	1.083	2.137
II-4	0.543	0.749	0.796	0.818	0.829	0.837			
II-2(fixed)	0.677	0.802	0.830	0.843	0.850	0.855			
II-4(fixed)	0.502	0.678	0.718	0.735	0.746	0.753			
II-2	0.754	0.829	0.837	0.840	0.843	0.843	0.027	2.562	2.911
II-4	0.716	0.840	0.849	0.856	0.857	0.859			
II-2(fixed)	0.749	0.822	0.831	0.836	0.837	0.838			
II-4(fixed)	0.614	0.736	0.749	0.754	0.758	0.760			
				Asymptote ≈ .80					
II-2	0.662	0.771	0.795	0.807	0.815	0.819	0.057	1.091	3.190
II-4	0.514	0.704	0.750	0.773	0.784	0.792			
II-2(fixed)	0.656	0.766	0.793	0.806	0.812	0.817			
II-4(fixed)	0.471	0.626	0.663	0.680	0.690	0.696			
II-2	0.711	0.785	0.795	0.799	0.803	0.803	0.028	1.960	3.209
II-4	0.640	0.770	0.785	0.794	0.795	0.799			
II-2(fixed)	0.706	0.777	0.787	0.792	0.794	0.795			
II-4(fixed)	0.549	0.662	0.675	0.681	0.686	0.688			
				Asymptote ≈ .75					
II-2	0.682	0.743	0.750	0.753	0.756	0.756	0.019	2.324	3.520
II-4	0.599	0.710	0.721	0.729	0.730	0.733			
II-2(fixed)	0.678	0.733	0.740	0.744	0.745	0.746			
II-4(fixed)	0.507	0.600	0.610	0.615	0.618	0.621			
II-2	0.675	0.737	0.742	0.745	0.748	0.747	0.011	4.068	3.172
II-4	0.592	0.715	0.725	0.732	0.732	0.736			
II-2(fixed)	0.672	0.727	0.731	0.736	0.736	0.737			
II-4(fixed)	0.494	0.602	0.611	0.617	0.620	0.622			
				Asymptote ≈ .70					
II-2	0.631	0.701	0.706	0.709	0.713	0.712	0.006	6.233	3.034
II-4	0.504	0.651	0.666	0.675	0.675	0.680			
II-2(fixed)	0.628	0.691	0.695	0.701	0.701	0.702			
II-4(fixed)	0.426	0.546	0.559	0.565	0.568	0.571			
II-2	0.648	0.707	0.713	0.716	0.719	0.719	0.011	3.431	5.013
II-4	0.541	0.655	0.668	0.676	0.677	0.681			
II-2(fixed)	0.645	0.697	0.703	0.707	0.708	0.709			
II-4(fixed)	0.458	0.551	0.562	0.567	0.570	0.573			

Note. II-2 = two-category information-integration task; II-4 = four-category information-integration task.

task. Since its inception, a fundamental assumption of the GCM is that people learn to attend selectively to relevant dimensions and to ignore irrelevant ones (Nosofsky, 1984, 1986). Following previous approaches, in applying the GCM to the present rule-based tasks, we assumed that subjects learned to optimize their attention weights by attending to just the single relevant dimension of the RB-2 task and giving equal attention to both relevant dimensions of the RB-4 task (Nosofsky, 1984). That is, whereas we continued to assume $w_1 = w_2 = .5$ in Equation 2 for the RB-4 task, we set $w_1 = 1$ and $w_2 = 0$ for the RB-2 task.

In Figure 9, we display the predicted learning curves from the model for all four Figure 1 structures simultaneously. The particular values of c, γ, and B that we used to generate these learning curves were the best fitting parameters from our previous fits to the four information-integration conditions of our Experiments 1 and 2 (see Figure 7). However, the same qualitative pattern of predictions as is displayed in Figure 9 holds over a very broad range of parameter settings from the model. To begin, note that at the same time that the model predicts nearly identical asymptotic performance for the II-2 and II-4(spread) conditions, it predicts clearly better performance for the RB-2 task compared to the RB-4 task. Thus, the model naturally predicts the fundamental dissociation of the effect of category number on rule-based and information-integration category learning that was the central theme of Maddox, Filoteo, et al.'s (2004) original report. The results, however, are far more impressive than simply capturing this dissociation result. In particular, in summarizing their empirical data, Maddox, Filoteo, et al. (2004, p. 231) also emphasized that learning performance in the

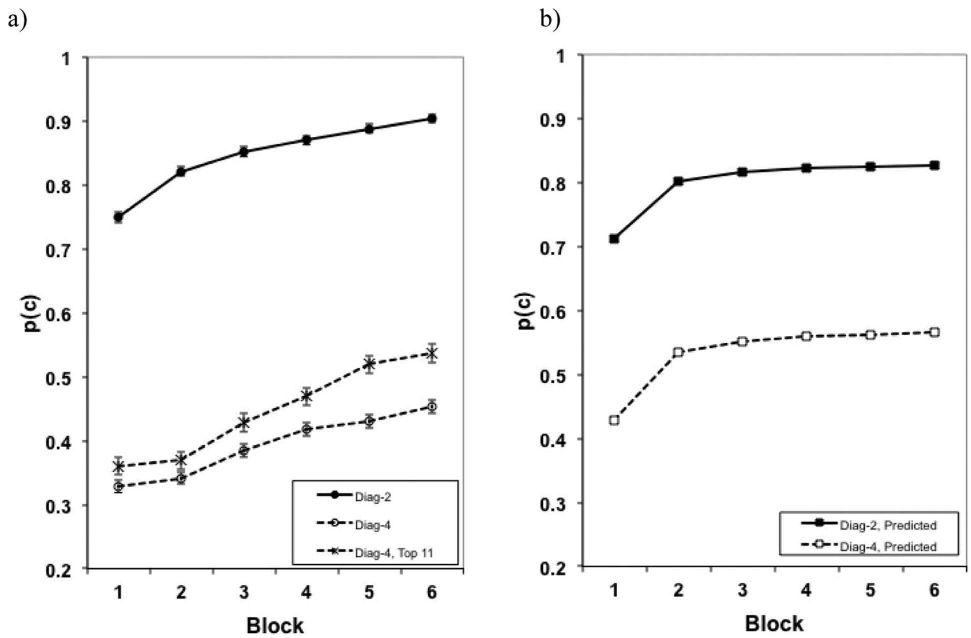

Figure 8. Observed and predicted accuracies for Experiment 3 shown as a function of category number and learning block. Error bars are within-block standard errors pooled across subjects. Diag-2 = two-category diagonal information-integration task; Diag-4 = four-category diagonal information-integration task.

RB-2 task was better than in either information-integration task, whereas learning performance in the RB-4 task was *worse* than in either information-integration task (see Maddox, Filoteo, et al., 2004, Figure 2). As can be seen in our Figure 9, the GCM provides a natural account of this more intricate pattern as well.

Although Maddox, Filoteo, et al. (2004, p. 231) made special note of the ordering of performance across their RB-2, II-2, II-4, and RB-4 conditions, they did not provide a theoretical explanation for the ordering. Their reasoning led them to predict that RB-2 performance would be better than RB-4 performance and that II-2 and II-4(spread) performance would be nearly equal, but nothing in their theoretical account allowed them to make direct comparisons *across* the rule-based and information-integration conditions. Because they provided no formalization of how learning actually took place in the explicit rule-based system and the implicit information-integration system, such theoretical comparisons were not possible. The success of the GCM in generating these qualitative predictions across the four conditions should therefore be viewed as a major achievement of that single-system model.

It is straightforward to understand the reason why the GCM predicts better performance in the RB-2 condition compared to the other conditions. As described in previous articles, the geometric interpretation of attention-weighting is that the psychological space is stretched along attended relevant dimensions and shrunk along unattended irrelevant ones (Nosofsky, 1984, 1986). When the observer attends selectively to the single relevant dimension of the RB-2 structure, it has the consequence of maximizing within-category similarities among exemplars and minimizing between-category similarities, thereby leading to greatly enhanced learning performance. The reason why the GCM predicts worse performance in the RB-4 condition compared to the II-4 condition is more subtle. Although the II-4 condition is a 45° rotation of the RB-4 condition and observers are presumed to attend equally to both dimensions in both tasks, it does not follow that similarities among exemplars are invariant across these conditions. In particular, for the highly separable-dimension stimuli used in these experiments, distances between exemplars are computed using a city-block distance metric (Equation 2; Garner, 1974; Shepard, 1964). Although distance is rotation invariant according to an (unweighted) Euclidean distance metric, such is not the case when distances are computed using a city-block distance metric. The 45° rotation of the RB-4 coordinate space in fact tends to produce slightly *greater* city-block distances between exemplars of contrasting categories than when the space is not rotated, thereby leading to improved category learning in the II-4 condition.

In sum, a fairly simple model based on a single representational system of stored exemplars accounts for the main qualitative pattern of results of the effects of category number on learning difficulty. Although our main focus was on the role of category number in influencing learning in information-integration tasks, an added bonus is that the model predicts the effects of category number on rule-based category learning as well, capturing the complete pattern of dissociation effects reported by Maddox, Filoteo, et al. (2004).

Quantitative Predictions of Individual-Subject Data

The focus of the previous section was on the general qualitative predictions from the GCM. However, a limitation of those analyses is that they pertained to patterns of averaged subject data. As is well known, averaged data are often not representative of the behavior patterns of individual subjects. Indeed, in evaluating the present work, an anonymous reviewer noted that although the manipulation of category structures across our Experiments 1 and 2 seemed relatively minor, a large performance difference was observed. This result led

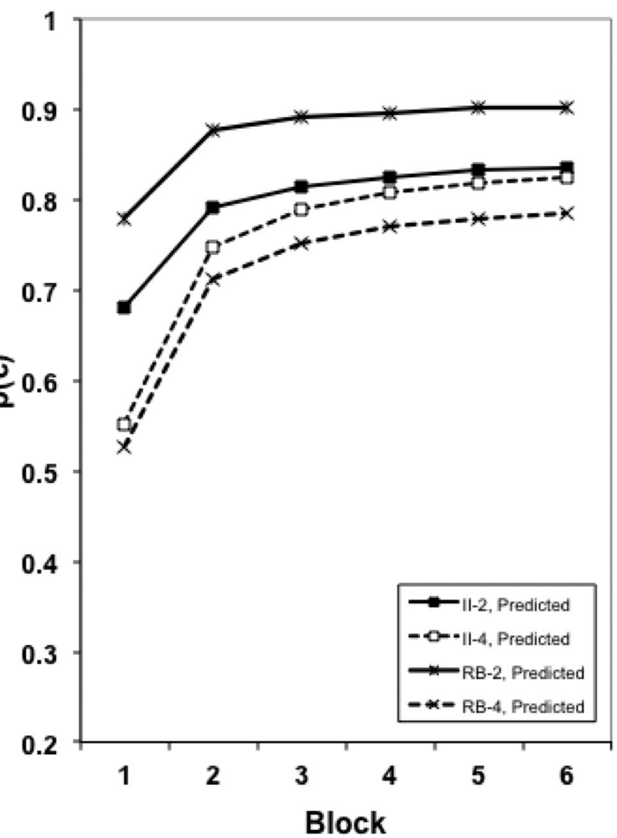

Figure 9. Predicted learning curves for all four structures depicted in Figure 1. II-2 = two-category information-integration task; II-4 = four-category information-integration task; RB-2 = two-category rule-based task; RB-4 = four-category rule-based integration task.

the reviewer to hypothesize that there may have been a qualitative change in strategies across the experiments (despite our finding that the GCM naturally predicts the large performance change—see Figure 7 and Table 2). In particular, the reviewer hypothesized that individual observers may have reverted to the use of suboptimal verbal rules in the II-4(fixed) condition of our Experiment 2, which would explain why performance in that condition was so much poorer than in the II-2(fixed) condition.

To evaluate this possibility, we conducted formal modeling analyses at the individual- subject level in which we compared the quantitative fits of information-integration and rule-based models. Although the key issue pertained to the II-4(fixed) condition, we fitted the models to the individual-subject data of all four information-integration conditions across Experiments 1 and 2. Following similar analyses conducted by Maddox, Filoteo, et al. (2004), our focus was on the performance data from the final three blocks, where most subjects appeared to have reached asymptote. As a representative from the class of single-system information-integration models, we again fitted the GCM to the data. Because the background-noise parameter B contributes negligibly to the predictions of asymptotic performance, we set $B = 0$ in these fits. However, because our focus was now on quantitative fitting of the data, we allowed the attention-weight parameter w_1 to vary freely for each individual subject. Except for extreme values, the magnitude of the attention weight has little influence on the general qualitative predictions of overall learning performance for the information-integration category structures. However, its value will have an important influence on the response-probability predictions associated with individual stimuli in the space.

Following Maddox, Filoteo, et al. (2004), as a plausible representative from the class of rule-based models, we fitted an extreme-values model to the data from the II-4(spread) and II-4(fixed) conditions. Using the layout illustrated in Figure 3, according to this model, the observer responds “A” if the line is short, responds “D” if the line is long, responds “B” if the line is medium in length and high in angle, and responds “C” if the line is medium in length and low in angle. Fitting this model requires estimation of the locations of two decision boundaries along the length dimension (to partition short lines from medium lines and to partition medium lines from long lines), a decision boundary along the angle dimension (to partition low angles from high angles), and an overall noise parameter to estimate the proportion of each stimulus distribution that falls into each of the response regions (see Maddox, Filoteo, et al., 2004, pp. 232–233, for details).

We fitted the models to each individual subject’s data by conducting computer searches for the values of the free parameters that minimized the Bayesian information criterion (BIC), a well-known statistic for comparing fits of models with differing numbers of free parameters. (A very similar pattern of results was obtained with use of the alternative Akaike information criterion statistic.) The BIC is given by

$$\mathrm{BIC} = -2\ln L + p\ln(M), \tag{5}$$

where $\ln L$ is the log-likelihood of the data given the best fitting parameters from the model, p is the number of free parameters in the model, and M is the total number of data points. The term $p\ln(M)$ in Equation 5 is a penalty term that penalizes a model for its number of free parameters. The model that yields the smaller value of BIC is considered to provide the more parsimonious account of the data.

For the present II-4 structures, the results were overwhelming: The GCM provided a better BIC fit than did the rule-based model for 18 of the 20 participants from Experiment 1 and for 21 of the 24 participants from Experiment 2. These analyses provide no evidence that participants switched to use of a rule-based strategy in the critical II-4(fixed) condition.

On the other hand, we should acknowledge that evidence for use of information-integration strategies was less clear-cut in the II-2 conditions. (Maddox, Filoteo, et al., 2004, observed a similar pattern of model-fitting results in their study.) For the two-category conditions, in addition to fitting the extreme-values rule model described above,[5] we also fitted plausible unidimensional rule-based models to the data. In particular, we fitted a model that assumed that subjects established a single decision boundary along the length dimension and also fitted a model that assumed that subjects established a single decision boundary along the orientation dimension. As it turned out, the GCM provided a better BIC

[5] In application to the II-2 conditions, both models use the same free parameters as in the II-4 conditions. However, the decision rule in the rule-based model is modified as follows: Respond “A” if the line is short or if the line is medium and has a high angle; respond “B” if the line is long or if the line is medium and has a low angle.

fit than did any of the rule-based models for 14 of the 20 participants in Experiment 1 and for 13 of the 24 participants in Experiment 2.

Although this evidence for use of information-integration strategies in the two-category conditions is equivocal, we do not believe that this modeling result compromises the main conclusions of our research. First, both II-2 conditions are basically replications of the original Maddox, Filoteo, et al. (2004) design. Thus, there is nothing about these conditions that might promote greater use of rule-based strategies. Second, the hypothesis that, for some reason, observers may have adopted rule-based strategies to a greater extent in the II-2 task of Experiment 2 than of Experiment 1 goes *against* the proposed explanation of our findings. In particular, to the extent that observers used suboptimal rule-based strategies in the II-2(fixed) condition, one would expect to see *lower* overall performance in that condition. Instead, we observed a greater difference between two-category and four-category performance in Experiment 2 compared to Experiment 1.

Most important, in light of these individual-subject modeling results, we reanalyzed the overall learning performance from Experiments 1 and 2 by including only those subjects for which the information-integration model (i.e., the GCM) provided a better quantitative fit than did the rule-based models. The pattern of learning curves for these subsets of subjects was identical to that reported for the full sets of subjects. Furthermore, statistical analysis of these subjects' data in each experiment yielded the same results as already reported for the full sets of subjects. Thus, restricting the analysis to only those subjects for which modeling suggested use of information-integration strategies, there is still clear evidence that learning performance in the II-4(fixed) condition is worse than in the II-2(fixed) condition.

Our inferences in this section were based on comparing one particular information-integration model (namely, the Equations 2–4 version of the GCM) to some particular plausible rule-based models. We acknowledge that alternative inferences may be derived through comparisons of alternative representatives of these model classes. Because it is impossible to compare the entire class of information-integration models to the entire class of rule-based ones, any such undertaking seems ultimately limited in principle. Nevertheless, in an online supplement, we report an investigation that considers a broader set of models from each class (and that also considers the results from the diagonal conditions of Experiment 3). The results from this investigation led to the same conclusions we have already reported above. There was no evidence that the poor performance observed in our II-4(fixed) and Diag-4 conditions resulted from subjects using suboptimal rule-based strategies in these information-integration tasks.

General Discussion

Summary

An extensive body of research, based on the demonstration of a variety of experimental dissociations, has sought to support the COVIS hypothesis that an explicit system learns rule-based categories and a separate implicit system learns information-integration categories (for reviews, see Ashby & Maddox, 2005; Ashby & O'Brien, 2005; Maddox & Ashby, 2004). In the present research, we addressed a single study that is representative of the larger set of studies that claim to have dissociated rule-based and information-integration category learning. Maddox, Filoteo, et al. (2004) derived a prediction from COVIS that if the number of perceptual clusters is constant across conditions, then the number of response categories will not impact performance on information-integration tasks. They argued, however, that rule-based tasks with a greater number of response categories would yield worse performance than rule-based tasks with fewer response categories. Consistent with this hypothesis, Maddox, Filoteo, et al. reported several experiments in which increasing the number of categories did lower performance on rule-based tasks, yet increasing the number of categories did not adversely affect performance on an information-integration task. They interpreted this dissociation as being consistent with the predictions from COVIS and as challenging single-system models of category learning.

We have argued, however, that within the Maddox, Filoteo, et al. (2004) studies, there exists a critical confound. To equate maximum accuracy across the two-category and four-category conditions, Maddox, Filoteo, et al. increased the distance between the perceptual clusters in the four-category task. Thus, the variables of category number and between-cluster distance were basically pitted against one another, making it difficult to ascertain the effect of category number per se.

The goal of the present studies was to test whether Maddox, Filoteo, et al.'s (2004) dissociation results do indeed challenge single-system models. We pursued this goal through both new empirical investigations and through model-based analysis.

Empirical results. First, we conducted information-integration category-learning studies that paralleled closely those tested by Maddox, Filoteo, et al. (2004), but where the between-cluster-distance confound was removed. We generated two-category and four-category structures in the same manner as did Maddox, Filoteo, et al., by creating perceptual clusters from bivariate normal distributions and merging pairs of adjacent clusters in the two-category tasks. In addition, we controlled the same extraneous variables that Maddox, Filoteo, et al. deemed to be important in demonstrating their dissociation, including the number of perceptual clusters across tasks, within-cluster scatter, distributional parameters of each cluster, and optimal accuracy. We controlled the latter variable not by spreading apart the perceptual clusters in the four-category task but instead by resampling values from the underlying population distributions.

In dramatic contrast to the results reported by Maddox, Filoteo, et al. (2004), we observed significantly worse learning in the four-category tasks compared to the two-category tasks. This significantly worse learning was observed for two different information-integration category structures. These results strongly challenge the generality of the category-number dissociation reported by Maddox, Filoteo, et al. They suggest strongly that the main basis for the null effect of category number in Maddox, Filoteo, et al.'s original design is that it confounded changes in between-cluster distance with the variable of interest. Furthermore, to the extent that Maddox, Filoteo, et al.'s reasoning about the predictions from the multiple-system COVIS model is correct, the result appears to challenge the predictions from COVIS.

Boundary conditions. It is crucial to acknowledge that it is impossible to control all extraneous factors while manipulating only category number in tasks of information-integration category learning. Although our design controlled all of the stated extraneous factors deemed important by Maddox, Filoteo, et al. (2004), one can always identify new factors that are no longer controlled. For example, using our procedure, an ideal observer's accuracy for classifying objects into their underlying populations would be higher for the II-2 task than for the II-4 task because there is overlap in the population distributions. Although such overlapping stimuli were never presented in our design, a multiple-system theorist could still posit that the factor forms an important boundary condition on whether there will be effects of category number on information-integration category learning.

In our view, the uncertainty over what are the crucial boundary conditions arises because the original category-number hypothesis was not strongly motivated. The stated reasoning from Maddox, Filoteo, et al. (2004, pp. 229–230) was as follows:

> Category learning in the II2 and II4 conditions is predicted to be the same . . . . The reasoning is as follows: Learning in the II2 and II4 conditions relies on an implicit, procedural-learning based system . . . . This system links clusters of "percepts" to categorization responses . . . . Because the complexity of the categories was held constant across conditions (i.e., the number of stimulus clusters was held fixed at four), and because each observer was informed of the number of categories prior to each condition, there is no a priori reason to expect a performance difference between information-integration category learning with two versus four categories.

Although these ideas are intriguing, they do not seem to provide a direct line of reasoning that explains *why* a system that "links clusters of 'percepts' to categorization responses" predicts a null effect of category number on learning performance. Therefore, it is difficult to know what are the crucial extraneous factors to control when testing that hypothesis.

Modeling results. In addition to the empirical results, our theoretical investigations involving single-system models provide an equally important demonstration that the Maddox, Filoteo, et al. (2004) result does not challenge the single-system viewpoint. The bottom line is that we showed that a simple representative from the class of single-system models (namely, the GCM) provided a good account of the complete set of results. That is, investigation of the qualitative predictions from the model showed that, across a fairly wide range of parameter settings, it naturally predicted the near-absent effect of category number on information-integration category learning in Maddox, Filoteo, et al.'s spread-cluster design, yet the big effect of category number in our own fixed-cluster designs. (As noted earlier, the model also accounts in natural fashion for the effects of category number in Maddox, Filoteo, et al.'s, 2004, rule-based category-learning tasks.) Thus, the category-number dissociation reported by Maddox, Filoteo, et al. does not appear to discriminate between the predictions from the single-system versus multiple-system approaches.

On the Status of Dissociations of Rule-Based and Information-Integration Category Learning

As noted in the introduction, Maddox, Ashby, and their colleagues have reported a wide variety of dissociations between rule-based and information-integration category learning that they have taken as supporting evidence for the COVIS model. Although we cannot address each and every individual study, in past work, researchers have tried to address studies that are representative of several main themes identified and reviewed by Ashby and Maddox (2005). The present research has continued that approach by addressing a new theme of dissociations involving the effect of category number on classification learning. To place the present work in perspective, however, we briefly review the manner in which some of the other major dissociations have been addressed. This brief review is intended to make clear that the present demonstrations are not an isolated response to the multiple-system-theorist claims but rather part of a more cohesive picture.

One of the themes identified by Ashby and Maddox (2005) uses procedural manipulations to interfere with the procedural learning component of the implicit system (e.g., Ashby et al., 2003; Maddox, Bohil, & Ing, 2004). Seeking to interfere with the implicit system's performance on information-integration categories, Ashby et al. (2003) implemented a procedural manipulation based on switching the mapping of category membership to response buttons. The manipulation impaired performance on an information-integration category structure but not a rule-based one. Ashby et al. interpreted this result as support for the view that rule-based categories are learned by an explicit system and information-integration categories are learned by a separate implicit system. The button-switch manipulation interfered with performance, according to these researchers, because the implicit system learns to associate a procedure with categorization decisions. However, Nosofsky et al. (2005) argued that all category learning involves a procedural learning component. They hypothesized that the rule-based task was unaffected by the procedural manipulation because the task Ashby et al. used was too simple and the observers had sufficient time to overcome any interference caused by the button-switch. Nosofsky et al. demonstrated that the button-switch does interfere with processing of the simple rule-based task when a more stringent time constraint is imposed. Furthermore, they were able to produce procedural interference for a more complex rule-based task even when using the longer response-time deadline.

A second theme of dissociation manipulations seeks to interfere with the explicit system's reliance on working memory and executive attention (e.g., Maddox, Filoteo, et al., 2004; Waldron & Ashby, 2001). Maddox, Filoteo, et al. (2004) reported experiments in which performing a concurrent memory-scanning task impaired learning of a rule-based category structure, leaving performance on an information-integration category structure intact. However, an important confound existed in their study. To equate learning difficulty across the rule-based and information-integration tasks, Maddox, Filoteo, et al. used a rule-based task with lower perceptual discriminability than the corresponding information-integration task. Stanton and Nosofsky (2007) demonstrated that when rule-based and information-integration category structures with highly discriminable stimuli are used, the memory-scanning task does not interfere with learning of either category structure. Furthermore, they demonstrated an example of interference from the memory-scanning task for an information-integration category structure with low perceptual discriminability. Interestingly, across these conditions, Stanton and Nosofsky demonstrated the reverse dissociation from that reported by Maddox, Filoteo, et al. Recent work reported by Newell, Moore, Wills, and Milton (in press) has also strongly questioned previous research purporting to show that concurrent memory load does not interfere with learning of information-integration categories.

In the third theme of dissociation manipulations, several studies have altered the delivery of corrective feedback to produce poor information-integration performance while rule-based performance is unaffected (e.g., Ashby et al., 2002; Maddox et al., 2003; Maddox & Ing, 2005). Ashby et al. (2002) demonstrated that learning performance on an information-integration category structure was worse for observational training than for feedback training; however, learning of a rule-based task was unaffected by the type of training. However, Stanton (2013) argued that the rule-based task Ashby et al. used involved learning of a simple category structure based on a single relevant dimension, whereas the information-integration category structure was cognitively complex. He argued that learning of all cognitively complex category structures may be impaired under observational training compared to feedback training. In accord with this hypothesis, Stanton provided examples of two cognitively complex rule-based category structures that yielded worse performance when the method of training was observational rather than feedback based. Related studies have strongly questioned previous research purporting to demonstrate dissociations between information-integration and rule-based category learning based on the timing of corrective feedback (e.g., Dunn et al., 2012; Stanton, 2013).

In sum, although numerous dissociations between rule-based and information-integration categorization have been reported, when critical extraneous variables are controlled, the dissociations often disappear or can even be reversed.[6] Moreover, many of the findings contradict the predictions derived from the multiple-systems COVIS model.

An interesting question that arises is how the debate between single-system versus multiple-system approaches in the domain of category learning relates more generally to the distinction between explicit and implicit memory. In our view, there may be no simple one-to-one mapping between alternative hypothesized categorization processes and these posited memory systems. Consider, for example, the exemplar-based representative from the class of single-system models. Conceivably, some types of exemplar retrieval may rise to the level of consciousness and have explicit properties, whereas others may operate under the surface and be implicit. Regarding neural underpinnings, although an interesting goal is to build bridges between the cognitive level of explanation provided by the GCM and specific neural systems, such bridges may be exceedingly complex. For example, Nosofsky, Denton, Zaki, Murphy-Knudsen, and Unverzagt (2012) suggested that the summed similarity signal that serves as the main decision variable in the GCM most likely receives contributions from a wide variety of neural components and structures, some of which have been implicated as part of explicit memory systems and others in implicit memory systems. Although COVIS theorists have hypothesized very direct mappings between rule-based processing and explicit neural systems and between information-integration processing and implicit/procedural neural systems, we wonder if the connecting bridges here may be complex as well. These questions are made even more complicated because, although the distinction between explicit and implicit memory is a major one in the fields of cognitive psychology and cognitive neuroscience, debate continues on the extent to which those phenomena reflect completely separate systems (e.g., Berry, Shanks, Speekenbrink, & Henson, 2012).

In any case, the present work converges with the studies reviewed above in questioning the dissociation logic that has been used to argue that rule-based category learning and information-integration category learning rely on separate, isolable systems. It is probably impossible to simultaneously control for all extraneous variables when comparing performance on rule-based and information-integration category structures. In addition, as illustrated in this article, single-system models often *predict* that dissociations between rule-based and information-integration category learning will be observed. As such, alternative approaches that do not rely solely on demonstrating dissociations may be useful in making progress in this field (e.g., see Newell & Dunn, 2008). Another fruitful avenue might be to develop more fully specified versions of both single-system and multiple-system models and evaluate their ability to provide comprehensive accounts of performance across a wide array of categorization tasks (e.g., Wills & Pothos, 2012).

[6] We should also note that in a closely related recent literature, new fMRI studies also raise questions about the interpretation of past brain-imaging dissociations that have been used as evidence of separate implicit and explicit category-learning systems (e.g., Gureckis, James, & Nosofsky, 2011; Milton & Pothos, 2011; Nosofsky, Little, & James, 2012). Some of these studies address another type of hypothesized implicit system involving prototype extraction (for a review, see Smith, 2008).

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. doi:10.1037/0033-295X.105.3.442

Ashby, F. G., Ell, S. W., & Waldron, E. M. (2003). Procedural learning in perceptual categorization. *Memory & Cognition, 31,* 1114–1125. doi:10.3758/BF03196132

Ashby, F. G., & Gott, R. E. (1988). Decision rules in the perception and categorization of multidimensional stimuli. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 14,* 33–53. doi:10.1037/0278-7393.14.1.33

Ashby, F. G., & Maddox, W. T. (1993). Relations between prototype, exemplar, and decision bound models of categorization. *Journal of Mathematical Psychology, 37,* 372–400. doi:10.1006/jmps.1993.1023

Ashby, F. G., & Maddox, W. T. (2005). Human category learning. *Annual Review of Psychology, 56,* 149–178. doi:10.1146/annurev.psych.56.091103.070217

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. doi:10.3758/BF03196423

Ashby, F. G., & O'Brien, J. B. (2005). Category learning and multiple memory systems. *Trends in Cognitive Sciences, 9,* 83–89. doi:10.1016/j.tics.2004.12.003

Berry, C. J., Shanks, D. R., Speekenbrink, M., & Henson, R. N. (2012). Models of recognition, repetition priming, and fluency: Exploring a new framework. *Psychological Review, 119,* 40–79. doi:10.1037/a0025464

Dunn, J. C., Newell, B. R., & Kalish, M. L. (2012). The effect of feedback delay and feedback type on perceptual category learning: The limits of multiple systems. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 38,* 840–859. doi:10.1037/a0027867

Erickson, M. A., & Kruschke, J. K. (1998). Rules and exemplars in category learning. *Journal of Experimental Psychology: General, 127,* 107–140. doi:10.1037/0096-3445.127.2.107

Garner, W. R. (1974). *The processing of information and structure*. New York, NY: Wiley.

Gureckis, T. M., James, T. W., & Nosofsky, R. M. (2011). Reevaluating dissociations between implicit and explicit category learning: An event-related fMRI study. *Journal of Cognitive Neuroscience, 23,* 1697–1709. doi:10.1162/jocn.2010.21538

Lewandowsky, S., Yang, L.-X., Newell, B. R., & Kalish, M. L. (2012). Working memory does not dissociate between different perceptual categorization tasks. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 38,* 881–904. doi:10.1037/a0027298

Love, B. C., & Gureckis, T. M. (2007). Models in search of the brain. *Cognitive, Affective, & Behavioral Neuroscience, 7,* 90–108. doi:10.3758/CABN.7.2.90

Maddox, W. T., & Ashby, F. G. (2004). Dissociating explicit and procedural-learning based systems of perceptual category learning. *Behavioural Processes, 66,* 309–332. doi:10.1016/j.beproc.2004.03.011

Maddox, W. T., Ashby, F. G., & Bohil, C. J. (2003). Delayed feedback effects on rule-based and information-integration category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 29,* 650–662. doi:10.1037/0278-7393.29.4.650

Maddox, W. T., Bohil, C. J., & Ing, D. A. (2004). Evidence for a procedural-learning-based system in perceptual category learning. *Psychonomic Bulletin & Review, 11,* 945–952. doi:10.3758/BF03196726

Maddox, W. T., Filoteo, J. V., Hejl, K. D., & Ing, A. D. (2004). Category number impacts rule-based but not information-integration category learning: Further evidence of dissociable category-learning systems. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 30,* 227–245. doi:10.1037/0278-7393.30.1.227

Maddox, W. T., & Ing, A. D. (2005). Delayed feedback disrupts the procedural-learning system but not the hypothesis-testing system in perceptual category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 31,* 100–107. doi:10.1037/0278-7393.31.1.100

McKinley, S. C., & Nosofsky, R. M. (1995). Investigations of exemplar and decision-bound models in large-size, ill-defined category structures. *Journal of Experimental Psychology: Human Perception and Performance, 21,* 128–148. doi:10.1037/0096-1523.21.1.128

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

Milton, F., & Pothos, E. M. (2011). Category structure and the two learning systems of COVIS. *European Journal of Neuroscience, 34,* 1326–1336. doi:10.1111/j.1460-9568.2011.07847.x

Newell, B. R., & Dunn, J. C. (2008). Dimensions in data: Testing psychological models using state-trace analysis. *Trends in Cognitive Sciences, 12,* 285–290. doi:10.1016/j.tics.2008.04.009

Newell, B. R., Dunn, J. C., & Kalish, M. (2011). Systems of category learning: Fact or fantasy? In B. H. Ross (Ed), *The psychology of learning & motivation* (Vol. 54, pp. 167–215). Burlington, VT: Academic Press.

Newell, B. R., Lagnado, D. A., & Shanks, D. R. (2007). Challenging the role of implicit processes in probabilistic category learning. *Psychonomic Bulletin & Review, 14,* 505–511. doi:10.3758/BF03194098

Newell, B. R., Moore, C. P., Wills, A. J., & Milton, F. (in press). Replacing the frontal lobes? Having more time to think improves "implicit" categorization: A comment on Filoteo, Lauritzen, & Maddox, 2010. *Psychological Science*.

Nosofsky, R. M. (1984). Choice, similarity, and the context theory of classification. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 10,* 104–114. doi:10.1037/0278-7393.10.1.104

Nosofsky, R. M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General, 115,* 39–57. doi:10.1037/0096-3445.115.1.39

Nosofsky, R. M., Denton, S. E., Zaki, S. R., Murphy-Knudsen, A. F., & Unverzagt, F. W. (2012). Studies of implicit prototype extraction in patients with mild cognitive impairment and early Alzheimer's disease. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 38,* 860–880. doi:10.1037/a0028064

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., & McKinley, S. C. (1992). Combining exemplar-based category representations and connectionist learning rules. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 18,* 211–233. doi:10.1037/0278-7393.18.2.211

Nosofsky, R. M., Little, D. R., & James, T. W. (2012). Activation in the neural network responsible for categorization and recognition reflects parameter changes. *PNAS: Proceedings of the National Academy of Sciences, USA, 109,* 333–338. doi:10.1073/pnas.1111304109

Nosofsky, R. M., & Palmeri, T. J. (1997). An exemplar-based random-walk model of speeded classification. *Psychological Review, 104,* 266–300. doi:10.1037/0033-295X.104.2.266

Nosofsky, R. M., Stanton, R. D., & Zaki, S. R. (2005). Procedural interference in perceptual classification: Implicit learning or cognitive complexity? *Memory & Cognition, 33,* 1256–1271. doi:10.3758/BF03193227

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. doi:10.1111/1467-9280.00051

Reed, S. K. (1972). Pattern recognition and categorization. *Cognitive Psychology, 3,* 382–407. doi:10.1016/0010-0285(72)90014-X

Rosch, E., Mervis, C. B., Gray, W., Johnson, D., & Boyes-Brehm, P. (1976). Basic objects in natural categories. *Cognitive Psychology, 8,* 382–439. doi:10.1016/0010-0285(76)90013-X

Shepard, R. N. (1964). Attention and the metric structure of the stimulus space. *Journal of Mathematical Psychology, 1,* 54–87. doi:10.1016/0022-2496(64)90017-3

Shepard, R. N. (1987, September 11). Toward a universal law of generalization for psychological science. *Science, 237,* 1317–1323. doi:10.1126/science.3629243

Smith, E. E. (2008). The case for implicit category learning. *Cognitive, Affective, & Behavioral Neuroscience, 8,* 3–16. doi:10.3758/CABN.8.1.3

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

Stanton, R. D. (2013). *Corrective feedback manipulations do not produce dissociations between explicit and implicit categorization*. Manuscript in preparation.

Stanton, R. D., & Nosofsky, R. M. (2007). Feedback interference and dissociations of classification: Evidence against the multiple learning-systems hypothesis. *Memory & Cognition, 35,* 1747–1758. doi:10.3758/BF03193507

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. doi:10.3758/BF03196154

Wills, A. J., & Pothos, E. M. (2012). On the adequacy of current empirical evaluations of formal models of categorization. *Psychological Bulletin, 138,* 102–125. doi:10.1037/a0025715

Received April 10, 2012
Revision received November 20, 2012
Accepted November 21, 2012 ■