

# The effects of category overlap on information-integration and rule-based category learning

SHAWN W. ELL

*University of California, Berkeley, California*

and

F. GREGORY ASHBY

*University of California, Santa Barbara, California*

In three experiments, we investigated whether the amount of category overlap constrains the decision strategies used in category learning, and whether such constraints depend on the type of category structures used. Experiments 1 and 2 used a category-learning task requiring perceptual integration of information from multiple dimensions (an information-integration task) and Experiment 3 used a task requiring the application of an explicit strategy (a rule-based task). In the information-integration task, participants used perceptual-integration strategies at moderate levels of category overlap, but explicit strategies at extreme levels of overlap—even when such strategies were suboptimal. In contrast, in the rule-based task, participants used explicit strategies, regardless of the level of category overlap. These data are consistent with a multiple systems view of category learning, and suggest that categorization strategy depends on the type of task that is used, and on the degree to which each stimulus is probabilistically associated with the contrasting categories.

In the study of category learning, it is often desirable to design tasks in which participants use a particular type of decision strategy. This goal is typically pursued by simply instructing participants to use a specific strategy (see, e.g., Allen & Brooks, 1991) rather than constraining the design of the categorization task. We propose that specifying the amount of overlap between contrasting categories may provide a simple method to constrain decision strategy. Category overlap historically has been manipulated to control task difficulty, and was not thought to affect the qualitative nature of the decision strategy used by participants. This article presents the results of three experiments that challenge this widely held view.

Category learning has been investigated using tasks that vary considerably with regard to stimulus materials, category structures, and procedure. For example, in some tasks, the entire stimulus set comprises just nine exemplars (e.g., Medin & Schaffer, 1978), whereas in other tasks, a single category comprises hundreds of exemplars (e.g., Ashby &

Gott, 1988). Despite this variability, in the majority of tasks, a trial begins with the presentation of a stimulus, followed by a categorization response, and typically, corrective feedback. Thus, at first glance, one might expect any effect of category overlap on decision strategy to be invariant across tasks. Recent research, however, suggests that the choice of task may be critical in determining the particular category-learning system and, consequently, the particular decision strategy that is used to learn the categories (Ashby & Ell, 2001). Thus, an alternative hypothesis is that the effect of category overlap on decision strategy may vary as a function of the task. We investigate this alternative using two category-learning tasks that have received the majority of attention in the multiple systems debate: information-integration and rule-based tasks (see Ashby & Maddox, 2005, and Maddox & Ashby, 2004, for a complete review of the dissociations between these two tasks).

Information-integration tasks are those in which accuracy is maximized by implicit, perceptual-integration strategies, which assume that information from two or more dimensions is integrated at some predecisional stage, outside of conscious awareness (Ashby, Alfonso-Reese, Turken, & Waldron, 1998). The type of perceptual integration required could take any number of forms, from a weighted combination of the two dimensions (Ashby & Gott, 1988; Garner, 1974) to more holistic processing (see, e.g., Kehler Nelson, 1993) to the incremental acquisition of stimulus-response associations (Ashby & Waldron, 1999), but the critical point is that the integration is assumed to occur prior to invoking any decision processes.

---

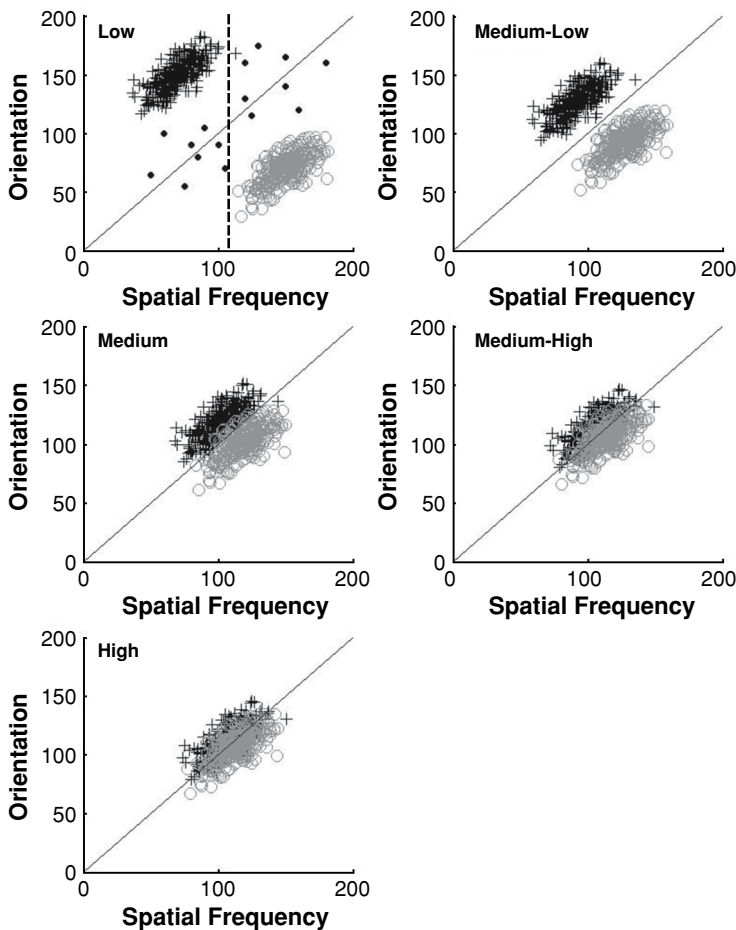
This research was supported in part by Public Health Service Grant MH3760. The authors thank Kelle Karp for her assistance in data collection. We also thank Richard Ivry, Todd Maddox, Shannon McCoy, Rebecca Spencer, Mari Reiss Jones, and three anonymous reviewers for their helpful comments on earlier drafts of the manuscript. Correspondence concerning this article should be addressed to S. W. Ell, Cognition and Action Lab, Helen Wills Neuroscience Institute and Psychology Department, University of California, 3210 Tolman Hall 1650, Berkeley, CA 94720-1650 (e-mail: shawnell@socrates.berkeley.edu).

For example, the stimuli might be Gabor patterns (sine wave gratings in which contrast is modulated by a circular Gaussian filter) that vary across trials in spatial frequency and orientation. In a typical information-integration task, the optimal decision bound might be set at the  $y = x$  line, requiring participants to attend to both spatial frequency and orientation in order to maximize accuracy. In this example, a verbal description of the optimal strategy would be to “Respond A if the difference between orientation and spatial frequency is positive and Respond B if the difference is negative.” However, this is not a strategy that participants readily verbalize, since it involves comparing dimensions measured in different units (Ashby et al., 1998). Nonetheless, given enough practice, neurologically healthy individuals are able to learn information-integration tasks,

even though they are rarely able to accurately describe their decision strategy (see, e.g., Ashby & Maddox, 1992).

Early in training, however, participants often use sub-optimal, explicit strategies (Ashby et al., 1998; Ashby, Queller, & Berretty, 1999). The simplest explicit strategy would assume that participants attend selectively to spatial frequency or orientation while ignoring the other dimension. In this example, the decision rule might be described as: “Respond A if the line is low in spatial frequency; otherwise Respond B.” Unlike perceptual-integration strategies, explicit strategies are assumed to be accessible to conscious awareness and, consequently, easily verbalizable (Ashby et al., 1998).

One question for the present research is whether varying the overlap between contrasting categories will fun-



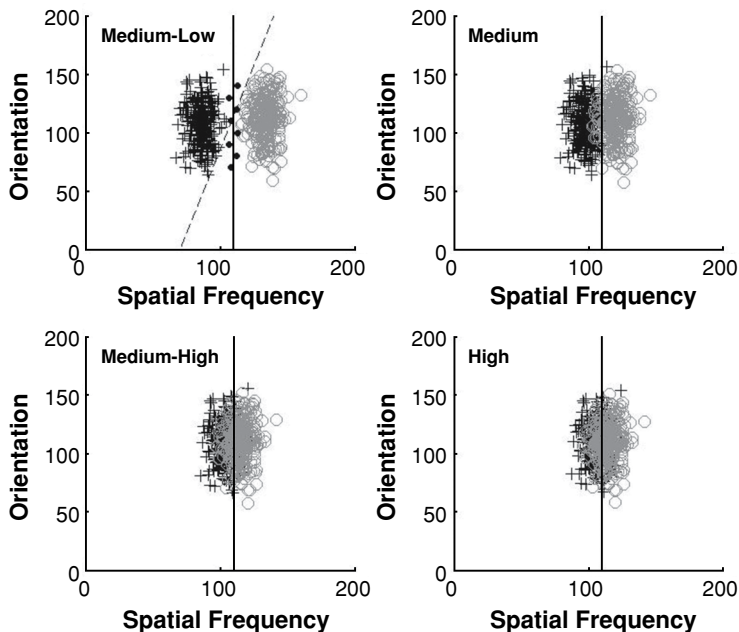
**Figure 1.** The information-integration task of Experiment 1 defined at five different levels of category overlap. Each symbol represents a Gabor stimulus in spatial frequency–orientation space (with spatial frequency in arbitrary units and orientation in degrees from horizontal). The black plus signs and gray circles denote Category A and B exemplars, respectively. The solid lines are the optimal perceptual-integration decision strategies. In the low condition, the dotted line is an example of an explicit strategy that would also maximize accuracy. Note that an explicit strategy assuming selective attention to orientation (i.e., a horizontal criterion) would also maximize accuracy. The filled circles in the low condition denote probe stimuli that were included to aid in distinguishing between explicit and perceptual-integration strategies.

damentally affect decision strategy in an information-integration task. In Experiment 1, we varied category overlap from low to high by decreasing the distance between category means (see Figure 1). In the low condition, the overlap is low enough that explicit (the dashed line) and perceptual-integration (the solid line) strategies predict similar accuracy. Thus, we would predict that participants would continue to use the seemingly default explicit strategies and never transition to perceptual-integration strategies. As overlap is increased, the accuracy of explicit strategies becomes far less than the accuracy of perceptual-integration strategies. Therefore, consistent with previous data, under normal conditions, participants should come to rely on perceptual-integration strategies and learn the information-integration task (see, e.g., Ashby & Maddox, 2005).

This line of reasoning suggests that perceptual-integration strategies will be used as long as they are more accurate than any competing explicit strategy. It is not clear, however, how large this accuracy advantage must be. In Figure 1, the accuracy advantage for the optimal perceptual-integration strategy is at its peak in the medium condition (see the Methods section for details of this computation), but steadily declines in the medium-high and high condi-

tions. Thus, an alternative possibility is that the accuracy advantage at greater levels of overlap may be insufficient to trigger the transition to perceptual-integration strategies. If so, then suboptimal, explicit strategies would dominate in the high (and perhaps medium-high) condition instead of perceptual-integration strategies.

In contrast, we would predict that category overlap has little effect on the decision strategy in tasks in which accuracy is maximized by explicit strategies (i.e., a rule-based task; Ashby et al., 1998). In Experiment 3, we tested this hypothesis by varying category overlap in a rule-based task. In the rule-based task of Figure 2, the optimal explicit strategy always requires participants to attend selectively to spatial frequency while ignoring orientation at all levels of category overlap.<sup>1</sup> In all cases, this strategy could be described as: “Respond A if the line is low in spatial frequency; otherwise Respond B.” At relatively low levels of overlap (i.e., the medium-low condition), we would expect that explicit strategies would dominate, given that they seem to be the default decision strategy (Ashby et al., 1999; Medin, Wattenmaker, & Hampson, 1987). As category overlap increases, the use of explicit strategies will consistently result in higher accuracy than



**Figure 2.** The rule-based task used in Experiment 3 defined at four different levels of category overlap. Each symbol represents a Gabor stimulus in spatial frequency–orientation space (with spatial frequency in arbitrary units and orientation in degrees from horizontal). The black plus signs and gray circles denote Category A and B exemplars, respectively. The solid lines are the optimal explicit decision strategies. In the medium-low condition, the dotted line is an example of a perceptual-integration strategy that would also maximize accuracy. The filled circles in the medium-low condition denote probe stimuli that were included to aid in distinguishing between explicit and perceptual-integration strategies. Note that the low condition was omitted because of its similarity to the medium-low condition.

with perceptual-integration strategies. Thus, we would expect explicit strategies to continue to dominate, regardless of the amount of overlap.

Across three experiments, we tested the predictions detailed above regarding the differential impact of category overlap on the decision strategy used in information-integration (Experiments 1 and 2) and rule-based (Experiment 3) category-learning tasks. The extant data suggest that category overlap can be used to constrain decision strategy, but only in information-integration tasks. Specifically, we predict that explicit strategies will be used when overlap is low and that perceptual-integration strategies will be used as overlap is increased. Whether participants will be constrained to use explicit or perceptual-integration strategies in an information-integration task at high levels of overlap is unclear.

## EXPERIMENT 1

### Method

**Participants and Design.** Twenty-five participants, ranging in age from approximately 18 to 26, were recruited from the University of California, Santa Barbara student community, and paid \$15 per experimental session (10 blocks of 60 trials) for their participation. There were a total of five experimental conditions, varying in category overlap. Five participants participated in each condition, and the number of training sessions varied across conditions as follows: low, 1; medium-low, 3 (1 participant completed 2 days of training); medium, 3; medium-high, 3; high, 4. No one participated in more than one experimental condition. All of the participants reported 20/20 vision or vision corrected to 20/20. Each session was approximately 45 min in duration and consecutive sessions were separated by 24 h, on average.

**Stimuli and stimulus generation.** Experiment 1 used an information-integration task at five levels of category overlap (Figure 1). Category overlap was varied from low to high by decreasing the distance between category means. The labels low, high, and so on are purely descriptive and are intended to represent ordinal relations among increasing levels of category overlap.

The experiment used the randomization technique introduced by Ashby and Gott (1988), in which each category was defined as a bivariate normal distribution. Each category distribution was specified by a mean and a variance on each dimension, and by a covariance between dimensions. The exact parameter values are displayed in Table 1. On each trial, a random sample  $(x, y)$  was drawn from the Category A or B distribution, and these values were used to construct a sine wave grating of spatial frequency  $x' = 0.2x - 1$  cycles/100 pixels and orientation  $y' = \pi/180y$  radians. The resulting range on

the spatial frequency dimension was .5 to 6 cycles/degree of visual angle. In each condition, all stimuli were generated offline and a linear transformation was applied to the sample stimuli to ensure that the sample statistics matched the population parameters. For all conditions, the two categories had identical spatial frequency and orientation variances and an identical spatial frequency–orientation covariance. Under these conditions, the optimal strategies in all conditions are linear (i.e., the solid decision boundaries plotted in Figure 1).<sup>2</sup> Participants who consistently responded using the optimal strategy would achieve the following accuracy rates: low = 100%, medium-low = 100%, medium = 96%, medium-high = 78%, high = 70%. In contrast, the participants who consistently responded using the most accurate explicit strategy would achieve accuracy rates of 100%, 92%, 75%, 62%, and 58% correct, respectively.<sup>3</sup>

In the low condition, it was predicted that the participants would use the explicit strategy depicted in Figure 1 (the dashed line). However, it would be impossible to distinguish between this explicit strategy and alternative perceptual-integration strategies (e.g., the solid line in the low condition of Figure 1) on the basis of accuracy alone. Probe stimuli (the filled black circles in Figure 1) were included to test this critical prediction. The coordinates of the probe stimuli were chosen so that the explicit and perceptual-integration strategies depicted in Figure 1 predicted different categorization responses. For example, those probe stimuli less than approximately 110 spatial frequency units would be classified as Category A exemplars, according to the explicit strategy, but only half of these probe stimuli would be classified as Category A exemplars according to the perceptual-integration strategy. A total of 16 probe stimuli were generated and were included in the final block of training. The coordinates of the probe stimuli are presented in Appendix A.

The stimuli were computer generated and displayed on a ViewSonic 15-in. CRT with  $832 \times 624$  pixel resolution in a dimly lit room. Each stimulus was presented on a gray background and subtended a visual angle of approximately  $3^\circ$ . The stimuli were generated and presented using the Psychophysics Toolbox (Brainard, 1997; Pelli, 1997) in the MATLAB software environment.

**Procedure.** The participants were told that there were two equally likely categories, and were informed of the optimal accuracy (e.g., in the high condition, the participants were instructed that, across the entire experiment, 70% correct was the highest accuracy they could achieve). All of the participants within a condition were presented with a different random ordering of the same 600 stimuli in each experimental session. On a trial, a single stimulus was presented and the participant was instructed to make a category assignment by depressing one of two response keys (labeled *A* or *B*) with his or her index fingers, and trial-by-trial feedback was provided for stimuli from the Category A and B distributions. In the low condition, the 16 probe stimuli were randomly interleaved with 34 Category A and B stimuli during the final block of trials. The probe trials differed, in that trial-by-trial feedback was omitted in an effort to minimize the likelihood that the probe stimuli would contribute to the forma-

**Table 1**  
Parameter Values Used to Generate the Stimuli From the Information-Integration Task at Each Level of Category Overlap

Condition	Means				Variances		
	Spatial Frequency		Orientation		Spatial Frequency	Orientation	Covariance
	Category A	Category B	Category A	Category B			
Low	70	150	150	70	162.5	162.5	112
Medium-low	92.5	127.5	127.5	92.5	162.5	162.5	112
Medium	101.5	118.5	118.5	101.5	162.5	162.5	112
Medium-high	106.1	113.9	113.9	106.1	162.5	162.5	112
High	107.4	112.6	112.6	107.4	162.5	162.5	112

Note—The spatial frequency values are in arbitrary units, and the orientation values are in degrees rotated clockwise from horizontal. Variances and covariance are for both categories.

tion of a participant's decision strategy. Specifically, the participants were instructed that trial-by-trial feedback would be provided, but that there would be some trials near the end of the experiment in which feedback would be omitted. The probe trials were used to aid in the identifiability of a participant's decision strategy in the model-based analyses (as described above), but were not considered when computing the participant's accuracy.

The trials were self-paced, with an upper time limit of 5 sec. If a response was not given in that time period, the participant was prompted to speed up his or her response, and that trial was discarded. A brief (1-sec) high-pitched tone (500 Hz) was presented if the response was correct, and a low-pitched tone (200 Hz) was presented if the response was incorrect. In addition, feedback was given at the end of each block of 60 trials regarding the participant's accuracy during that block. The response-stimulus interval was 1 sec.

**Results**


**Accuracy-based analyses.** The learning curves for each of the five overlap conditions are plotted in Figure 3. Visual inspection of the learning curves indicates that average accuracy improved with training, with the exception of the high condition, in which average accuracy hovered around chance levels and never approached optimal (i.e., 70%). As expected, asymptotic accuracy decreased with increasing category overlap. A one-way ANOVA conducted on the average accuracy during the final session (low,  $M = 96.64, SD = 2.50$ ; medium-low,  $M = 98.03, SD = .43$ ; medium,  $M = 87.45, SD = 5.06$ ; medium-high,  $M = 60.79, SD = 5.69$ ; high,  $M = 51.21, SD = 3.35$ ) confirmed this observation [ $F(4,20) = 152.84, MS_e = 15.12, p < .001$ ] with post hoc tests (Tukey) revealing that all conditions were significantly different from each other ( $p < .05$ ), with the exception of the high and medium-high conditions.

To investigate whether or not accuracy improved with training in the high condition, a within-subjects ANOVA was conducted on the average accuracy from each of the four sessions (Session 1,  $M = 53.11, SD = 2.92$ ; Session 2,

$M = 53.00, SD = 3.55$ ; Session 3,  $M = 52.08, SD = 2.35$ ; Session 4,  $M = 51.21, SD = 3.35$ ). This analysis revealed that there was no significant difference in accuracy across sessions [ $F(3,12) = .69, MS_e = 5.71, p = .58$ ]. The average accuracy was marginally significantly greater than chance during the first session [ $t(4) = 2.39, p = .07$ ], but not significantly greater than chance in the remaining sessions [Session 2,  $t(4) = 1.89, p = .13$ ; Session 3,  $t(4) = 1.98, p = .12$ ; Session 4,  $t(4) = .81, p = .47$ ].

**Model-based analyses.** As expected, the analysis of the accuracy data revealed that performance generally decreased with increasing overlap. The primary question, however, concerns the effect of overlap on the decision strategy. For example, participants might use the optimal perceptual-integration strategy (see Figure 1) at all levels of overlap. Alternatively, participants might rely on suboptimal, explicit strategies. The following analyses present a quantitative approach to evaluating these hypotheses. Specifically, we fit a number of different decision bound models (Ashby, 1992a; Maddox & Ashby, 1993) to each participant's responses. Decision bound models assume each participant partitions the perceptual space into response regions by constructing a decision bound. On each trial, the participant determines which region the percept is in, and then produces the associated response. To be clear, in the present application, decision bound models are used as an analytic tool to provide a description of each participant's data.

Two different types of models were fit to each participant's responses (see Appendix B for more details). One type was compatible with the assumption that participants used an explicit strategy and one type assumed a perceptual-integration strategy. These models make no detailed process assumptions, in the sense that a number of different process accounts are compatible with each of the models (e.g., Ashby, 1992a; Ashby & Waldron, 1999). For ex-



**Figure 3.** Average accuracy across levels of category overlap in the information-integration task of Experiment 1.

ample, if a perceptual-integration model fits significantly better than an explicit model, we can be confident that participants did not use an explicit strategy, but we cannot specify which perceptual-integration strategy was used (e.g., a weighted combination of the two dimensions vs. more holistic processing). Thus, the modeling described in this section provides a powerful vehicle to test hypotheses about the decision strategies used by participants, even though it has little to say about psychological process.

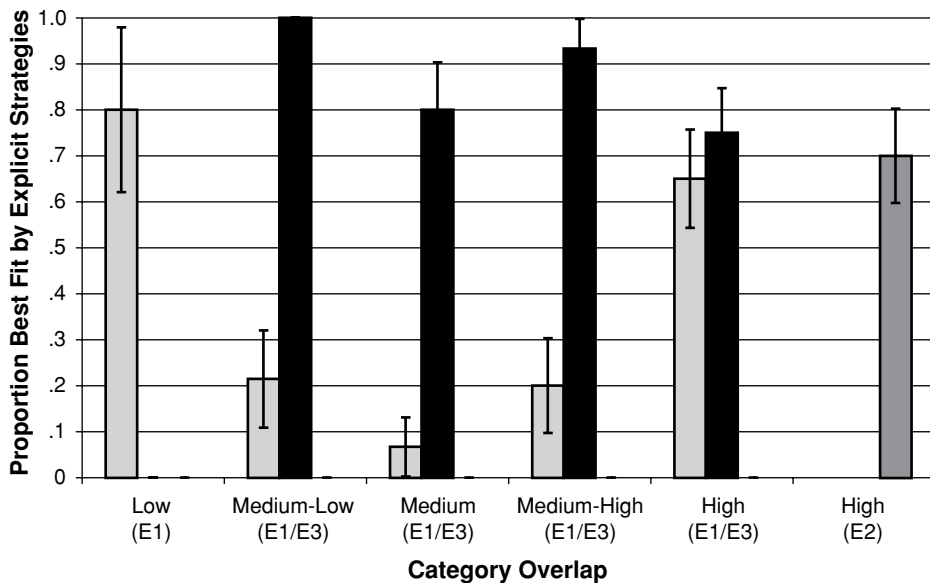
Each model was fit separately to the data for every participant from every training session (excluding the first block in the session, because it was considered practice). Analyzing the data in this way resulted in the following number of data sets per condition: low, 5 (1 training session  $\times$  5 participants), medium-low, 14 (3 training sessions  $\times$  4 participants, 2 training sessions  $\times$  1 participant who completed only 2 training sessions), medium, 15 (3 training sessions  $\times$  5 participants), medium-high, 15 (3 training sessions  $\times$  5 participants), high, 20 (4 training sessions  $\times$  5 participants).

The proportion of data sets best accounted for by explicit strategies across experimental condition is plotted in Figure 4. For simplicity, the data from all three experiments have been plotted in Figure 4, and the Experiment 1 data is plotted in light gray. For moderate levels of category overlap (the medium-low, medium, and medium-high conditions), perceptual-integration strategies were found to dominate, whereas for the two extreme overlap conditions (the low and high conditions), explicit strategies dominated.<sup>4</sup> More-

over, the prevalence of explicit strategies was approximately equal in the low and high conditions.

The results of the model-based analysis are only valid to the extent that the models themselves provide adequate accounts of the observed data. One practical method of assessing goodness-of-fit is to compute the percentage of responses the best-fitting model was able to reproduce. The average percentage of responses accounted for is listed in Table 2. First consider the data from the low condition. Not surprisingly, the best-fitting models provide a very good account of these data. Central to the question of whether or not these data were best fit by explicit or perceptual-integration strategies is how well the best-fitting model was able to reproduce the probe data. As can be seen in Table 2, the best-fitting models did quite well in this regard. In fact, across the 4 participants best fit by explicit strategies, the best-fitting model failed to reproduce only one response. The best-fitting model failed to reproduce five responses for the participant whose data were best fit by a perceptual-integration model. Thus, we can be confident that the dominance of explicit strategies in the low condition cannot be attributed to the undue influence of the probe stimuli.

Given the probabilistic nature of the medium, medium-high, and high conditions, it was expected that the best-fitting models would account for far less than 100% of the responses. However, it is potentially problematic that the best-fitting model accounted for fewer than 60% of the responses for 5 of the 20 data sets in the high condition (2



**Figure 4.** The proportion of data sets best accounted for by explicit strategies as a function of category overlap across the three experiments. Data from the information-integration task of Experiment 1 (E1) are plotted in light gray (low, medium-low, medium, medium-high, and high conditions). Data from the information-integration task of Experiment 2 (E2) are plotted in dark gray (high condition). Data from the rule-based task of Experiment 3 (E3) are plotted in black (medium-low, medium, medium-high, and high conditions). Error bars represent the standard error of the proportion.

**Table 2**  
**Average Percentage of Responses Accounted for by the Best-Fitting Model Across Levels of Category Overlap for All Experiments**

Condition	E1		E2		E3	
	Avg	SEM	Avg	SEM	Avg	SEM
Low	97.2	1.0	—	—	—	—
Low-probe	92.5	6.1	—	—	—	—
Medium-low	94.7	1.3	—	—	96.8	.4
Medium	84.8	2.3	—	—	89.6	1.6
Medium-high	66.6	1.7	—	—	77.1	1.7
High	61.8	2.5	69.4	1.5	75.4	1.7

Note—Low-probe, probe stimuli from the low condition. E1, Experiment 1; E2, Experiment 2; E3, Experiment 3.

from Session 1; 3 from Sessions 3 and 4). After excluding these 5 data sets, explicit strategies accounted for 87% of the remaining data sets, and the average percentage of responses accounted for increased to 67%.

**EXPERIMENT 2**

The results of Experiment 1 suggest that the decision strategy in information-integration tasks is nonmonotonically related to category overlap. However, these conclusions are critically dependent on the model-based analyses of the data from the high condition. There is a long tradition of collecting a large amount of data (in this case, 3,000 data points per participant) from a limited number of participants in cognitive psychology. Nevertheless, it could reasonably be argued that the interpretation of the present data would be more compelling if the sample size were increased. Experiment 2 replicated the critical high condition of Experiment 1 with the addition of a monetary incentive intended to increase motivation to maximize accuracy. Note that this manipulation intentionally biases

the data in favor of perceptual-integration strategies, because if the participants are maximizing accuracy (to achieve the greater reward) then the model-based analyses would indicate that the participants were using perceptual-integration strategies.

**Method**

**Participants and Design.** Five participants (4 female), ranging in age from 20 to 24, were recruited from the University of California, Berkeley student community, and paid \$10 per experimental session (plus an accuracy bonus) for participation. All participants reported 20/20 vision or vision corrected to 20/20. Each session was approximately 45 min in duration and consecutive sessions were separated by 24 h, on average.

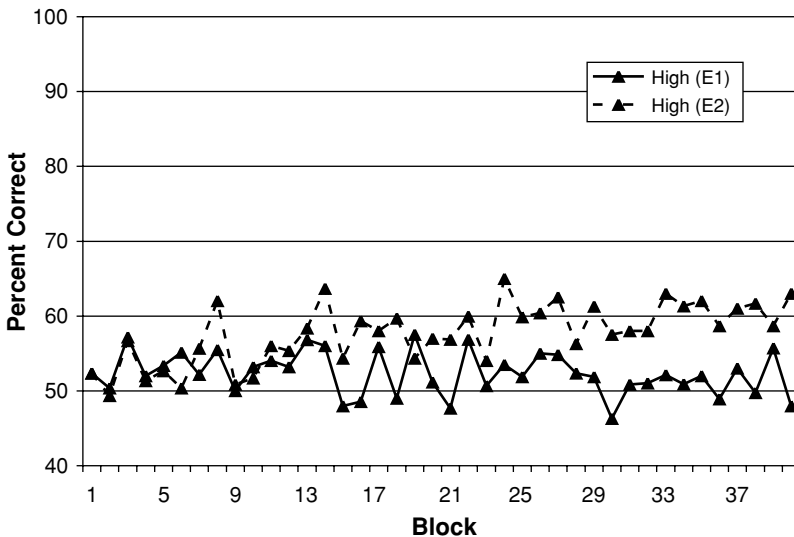
**Stimuli and stimulus generation.** The stimuli were identical to those in the high condition of Experiment 1.

**Procedure.** The procedure was identical to that of Experiment 1, with the exception that the participants were instructed that they would receive a \$1 bonus for every block in which their accuracy exceeded 60% correct. Thus, across the four days of training, it was possible to earn a bonus of \$48.

**Results**

**Accuracy-based analyses.** The learning curve is shown in Figure 5, along with the learning curve from the high condition of Experiment 1, for reference. Inspection of the learning curves suggests that increasing the motivation to maximize accuracy resulted in improved performance by the end of training. A comparison of the average accuracy from the final session of Experiment 1 ( $M = 51.2\%$ ,  $SD = 3.35$ ) and Experiment 2 ( $M = 60.5\%$ ,  $SD = 5.03$ ) supported this claim [ $t(8) = 3.45$ ,  $SE_{diff} = 2.71$ ,  $p = .009$ ].

Importantly, unlike in Experiment 1, there was a clear improvement in accuracy late, relative to early, in training. To verify this claim, a within-subjects ANOVA was conducted on the average accuracy from each of the four sessions of the present experiment (Session 1,  $M = 53.4$ ,



**Figure 5.** Average accuracy in the high condition of Experiment 2 (E2). The data from the high condition of Experiment 1 (E1) are presented for comparison.

$SD = 1.04$ ; Session 2,  $M = 57.59$ ,  $SD = 4.40$ ; Session 3,  $M = 59.43$ ,  $SD = 3.29$ ; Session 4,  $M = 60.53$ ,  $SD = 5.04$ ). This analysis revealed a significant effect of session [ $F(3,12) = 7.29$ ,  $MS_e = 6.71$ ,  $p = .005$ ] with a marginally significant increase in accuracy from Block 1 to Block 2 ( $p = .08$ ), and significant increases in Blocks 3 ( $p = .01$ ) and 4 ( $p = .03$ ), relative to Block 1. The average accuracy was also significantly greater than chance performance in all four sessions [Session 1,  $t(4) = 7.32$ ,  $p = .002$ ; Session 2,  $t(4) = 3.86$ ,  $p = .018$ ; Session 3,  $t(4) = 6.41$ ,  $p = .003$ ; Session 4,  $t(4) = 4.68$ ,  $p = .009$ ].

**Model-based analyses.** As in Experiment 1, each model was fit separately to the data for every participant from every training session (excluding the first block in the session, because it was considered practice). Analyzing the data in this way resulted in 20 data sets (4 training sessions  $\times$  5 participants). As can be seen in Figure 4 (i.e., the dark gray bar on the far right), the results of the model-based analyses are consistent with the data from Experiment 1 in suggesting that suboptimal, explicit decision strategies dominated. Importantly, the best-fitting model accounted for more than 60% of the responses in 17 of the 20 data sets and more than 70% of the responses in 11 of the 20 data sets. In sum, even when the participants have incentive to use perceptual-integration decision strategies, they seem limited in their ability to do so.

Given the relatively low accuracy in the high conditions of Experiments 1 and 2, and the less than perfect ability of the best-fitting models to reproduce the observed data, it is important to consider whether these data truly reflect an intact decision process. One approach to answering this question is to investigate the pattern of data that would be expected from participants known to be using some specific decision strategy. We simulated such data from two different models in the high overlap condition. One assumed the optimal decision strategy (i.e., the optimal perceptual-integration model in Appendix B) and one assumed the most accurate explicit strategy (i.e., the dashed line in the low condition of Figure 1). In both models, we set the internal noise to the median noise estimate obtained from the model-based analyses in the high-overlap conditions of Experiments 1 and 2 ( $Mdn = 22.5$ ). The optimal decision bound accounted for a mean of only 62% of the responses generated from the optimal decision bound model ( $SD = 2.2\%$ , averaged over 100 replications). Similarly, the most accurate explicit bound accounted for a mean of 67% of the responses generated from the most accurate explicit model ( $SD = 1.7\%$ , averaged over 100 replications). Note that these percentages are similar to the percentages of responses accounted for by the models that best fit the data from the high-overlap condition. If participants were simply guessing in this condition, each of the models we fit would account for only 50% of the observed responses. Thus, the model-based analyses are consistent with the hypothesis that participants did use an intact decision process in the high-overlap condition, and that they were not simply guessing. Although accuracy was quite low in the high condition of Experiment 1, the addition of a monetary incen-

tive in Experiment 2 encouraged learning and resulted in accuracy rates well above chance.

### EXPERIMENT 3

The results of Experiments 1 and 2 showed that the decision strategy varied nonmonotonically with category overlap in an information-integration task. Specifically, explicit strategies were used at extreme levels of category overlap, and perceptual-integration strategies at intermediate levels. Recall, however, that the primary question is whether the effect of overlap on the decision strategy will vary as a function of the categorization task. Alternatively, it may be the case that the nonmonotonicity observed across Experiments 1 and 2 is a generic feature of all categorization tasks. Experiment 3 tests this hypothesis by replicating Experiment 1 on a rule-based task. The rule-based task is plotted at four levels of category overlap in Figure 2. The low condition was omitted from Experiment 3 because of its similarity to the medium-low condition (and because there are no predicted theoretical differences between these conditions in the present experiment).

#### Method

**Participants and Design.** Twenty-three participants, ranging in age from approximately 18 to 26, were recruited from the University of California, Santa Barbara student community, and paid \$15 per experimental session for their participation. There were a total of four experimental conditions, varying in category overlap. Eight participated in the medium-low condition, whereas 5 participated in each of the remaining conditions. The number of training sessions varied across conditions as follows: medium-low, 1; medium, 3; medium-high, 3; high, 4. No one participated in more than one experimental condition. All of the participants reported 20/20 vision or vision corrected to 20/20. Each session was approximately 45 min in duration and consecutive sessions were separated by 24 h, on average.

**Stimuli and stimulus generation.** The stimuli were generated by rotating the Experiment 1 category structures from the medium-low, medium, medium-high, and high conditions 45° counterclockwise. Specifically,

$$\mathbf{X}' = \mathbf{X} \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix},$$

where  $\mathbf{X}'$  is the new Experiment 3 stimuli,  $\mathbf{X}$  is the matrix of stimuli from Experiment 1, and  $\theta = .7854$  radians. The low condition was omitted from Experiment 3, because the rotated low and medium-low conditions were redundant with respect to the theoretical predictions. A scatterplot of the stimuli used in the experiment is shown in Figure 2.

Following the same logic as in Experiment 1, a number of probe stimuli were included in the medium-low condition to aid in the identifiability of the participant's decision strategy. A total of eight probe stimuli were generated (the black, filled circles in Figure 2) and were included in the final block of training. The coordinates of the probe stimuli are listed in Appendix A.

**Procedure.** The procedure was identical to that of Experiment 1.

#### Results

**Accuracy-based analyses.** The learning curves for each of the four levels of category overlap are shown in Figure 6. Visual inspection of the learning curves indicates






Figure 6. Average accuracy across levels of category overlap in the rule-based task of Experiment 3.

that accuracy improved with training in all conditions. As expected, average accuracy was negatively correlated with category overlap. A one-way ANOVA conducted on the average accuracy during the final day of training was generally consistent with this observation [ $F(3,22) = 75.88$ ,  $MS_e = 27.36$ ,  $p < .001$ ], but only the medium-low ( $M = 95.37$ ,  $SD = 5.51$ ) and medium ( $M = 90.20$ ,  $SD = 3.97$ ) conditions were significantly different from the medium-high ( $M = 69.09$ ,  $SD = 2.99$ ) and high ( $M = 66.04$ ,  $SD = 1.75$ ) conditions ( $p < .001$ ). None of the remaining pairwise comparisons were significant ( $p > .05$ ).

**Model-based analyses.** As in the previous experiments, each model (described in Appendix B) was fit separately to the data for every participant from every training session (excluding the first block in the session, because it was considered practice). Analyzing the data in this way resulted in the following number of data sets in each condition: medium-low, 8 (1 training session  $\times$  8 participants), medium, 15 (3 training sessions  $\times$  5 participants), medium-high, 15 (3 training sessions  $\times$  5 participants), high, 20 (4 training sessions  $\times$  5 participants).

The results of the model-based analysis are summarized in Figure 4 (black bars). In contrast to the results from the information-integration task of Experiments 1 and 2, inspection of Figure 4 reveals that explicit strategies consistently dominated across levels of category overlap in the rule-based task. Furthermore, on average, the best-fitting model accounted for more than 70% of the responses in all conditions (see Table 2). At the level of the individual participant, the best-fitting model accounted for at least 60% of the responses for all but 1 participant.

## DISCUSSION

The results of these experiments demonstrate that, contrary to the traditional view, category overlap can be used

as an effective means to constrain decision strategy. The success of this approach, however, varies depending on the category-learning task. Specifically, the degree of category overlap constrains the type of decision strategy used in information-integration tasks, but not rule-based tasks. For information-integration tasks, participants were constrained to use perceptual-integration strategies at moderate amounts of category overlap, whereas participants used explicit strategies at extreme amounts of overlap. In contrast, for rule-based tasks, explicit strategies were used consistently, regardless of the amount of category overlap.

One possible explanation of these results is that different category-learning systems, specialized to use explicit and perceptual-integration strategies (respectively), are in competition throughout learning. Such assumptions are found in the COVIS theory of category learning (competition between verbal and implicit systems; Ashby et al., 1998). COVIS hypothesizes that category learning is a competition between separate explicit and implicit systems. The explicit system is a logical reasoning system that uses explicit strategies and is assumed to dominate learning in rule-based tasks. The implicit system is a procedural-learning system that uses perceptual-integration strategies and is assumed to dominate in information-integration tasks. In COVIS, the two systems operate in parallel, and both systems compete for control of the observable categorization response. Initially, the system weight, which reflects the relative dominance of the explicit system, strongly favors the explicit system (Ashby et al., 1999). The system weight is adjusted up and down during learning, based on the relative success of each system.<sup>5</sup> For example, in the medium condition of Experiment 1, the system weight would eventually shift in favor of the implicit system.

In the low condition of Experiment 1, a perceptual-integration strategy yields perfect accuracy, but the most accurate explicit strategy (i.e., the dashed line in Figure 1)

does just as well. Because the COVIS explicit system initially dominates and there is a simple explicit strategy that will consistently be rewarded, COVIS predicts that the explicit system will dominate performance in the low condition. In the remaining conditions, the most accurate explicit strategies predict poorer performance than the optimal perceptual-integration strategy. Despite the initial dominance of the explicit system, COVIS predicts that the implicit system will be rewarded more frequently and that participants will eventually learn to use perceptual-integration strategies. In contrast, when category overlap is high, the accuracy of the best explicit strategy is low, but so is the accuracy of the best perceptual-integration strategy. In this case, neither system will be consistently rewarded, so COVIS predicts that the implicit system will not be able to overcome the initial advantage of the explicit system and that responding will be dominated by explicit strategies.

COVIS assumes that, throughout learning, explicit and implicit systems compete for control of the categorization response and that the system that is best suited for the particular task eventually dominates. Indeed, many of the observed dissociations between rule-based and information-integration tasks could be interpreted as evidence for competition. However, we would argue that the present data go one step further by beginning to characterize conditions under which this competition facilitates the dominance of one system or the other. At extreme levels of category overlap, the explicit system wins the competition, which is reflected by the high frequency of explicit strategies, regardless of the task. In contrast, at moderate levels of category overlap, the implicit system is able to overcome the initial bias to use explicit strategies in the information-integration task. Consistent with this claim, in the conditions with the greatest overlap (i.e., high conditions of Experiments 1 and 2), there was very little change in the average proportion of data sets best fit by an explicit strategy: .70, .70, .60, and .70 across Sessions 1 through 4, respectively. In contrast, at moderate levels of category overlap (medium-low, medium, and medium-high conditions), there was a trend for the average proportion of data sets best fit by an explicit strategy to decrease across the three experimental sessions from .27 to .13 to .07. Further research will be needed to better characterize the dynamics of this competitive process.

In sum, COVIS successfully predicts that for an information-integration task, performance should be dominated by explicit strategies at extreme levels of category overlap, whereas performance should be dominated by implicit, perceptual-integration strategies at moderate levels of category overlap. COVIS predicts the present results because it assumes that (1) the explicit system initially dominates responding, and (2) it is only in situations in which the implicit system is rewarded more frequently than the explicit system that participants will be able to overcome this initial dominance and learn perceptual-integration strategies (Ashby et al., 1998; Ashby et al., 1999).

The present results add to the growing number of dissociations between rule-based and information-integration tasks that have been predicted a priori by COVIS (see Ashby & Maddox, 2005; Maddox & Ashby, 2004, for reviews). Multiple systems arguments have also been made in such diverse fields as reasoning (Sloman, 1996), motor learning (Willingham, 1998), discrimination learning (Kendler & Kendler, 1962), function learning (Hayes & Broadbent, 1988), and identification (Ashby, Waldron, Lee, & Berkman, 2001), as well as by other category learning researchers (e.g., Brooks, 1978; Erickson & Kruschke, 1998; Nosofsky, Palmeri, & McKinley, 1994). Nevertheless, many recent categorization articles have argued for a single system that mediates all category learning (Nosofsky & Kruschke, 2002; Pothos, 2005; Zaki, Nosofsky, Jessup, & Unversagt, 2003; Zaki, Nosofsky, Stanton, & Cohen, 2003). Although we cannot rule out the possibility that some single system model could account for the present results, a significant challenge for single system theorists is to account for the growing number of observed dissociations between rule-based and information-integration tasks within the same unified model.

Our results suggest that one strategy for designing a categorization task that encourages the use of implicit, perceptual-integration strategies is to use an information-integration task with moderate category overlap. In past applications of information-integration tasks, the primary design criterion for encouraging participants to use a perceptual-integration strategy, rather than an explicit strategy, was to ensure that the accuracy of the optimal (perceptual-integration) strategy substantially exceeded that of the most accurate explicit strategy. The present results suggest that this criterion, by itself, is insufficient. For example, consider the medium-low, medium, medium-high, and high conditions of Experiment 1, in which perceptual-integration strategies outperform explicit strategies by 8%, 21%, 16%, and 12%, respectively. In the medium-low, medium, and medium-high conditions, perceptual-integration strategies clearly dominated. If ensuring an accuracy advantage for the best perceptual-integration strategy was the sole condition for recruiting the implicit system, then perceptual-integration strategies should have dominated in the high condition, but this was not the case. Future research will focus on characterizing the necessary conditions for recruiting systems that utilize explicit and perceptual-integration strategies.

Increasing category overlap correlated highly with task difficulty in the present experiments. To be clear, we are not claiming that any method of varying task difficulty would be sufficient to reproduce these results. Instead, the value of manipulating category overlap was that it produced category structures that varied in asymptotic accuracy. Some manipulations of task difficulty will lower asymptotic accuracy, which should therefore favor explicit strategies. Other manipulations would not necessarily reduce asymptotic accuracy. Instead, they might simply delay the time it takes participants to reach such a level of performance (e.g., when the optimal decision bound is

quadratic vs. linear). In these cases, COVIS predicts that varying task difficulty will not cause a shift to explicit strategies.

Increasing category overlap also reduces the validity of the feedback. In the low and medium-low conditions, each exemplar is unambiguously a member of one and only one category. Therefore, when a stimulus is presented, feedback indicates that a Category A response was correct, with probability 1 or 0. In the medium, medium-high, and high conditions, however, the categories overlapped, which means that each stimulus in the overlapping region could belong to Category A or B. Thus, for these stimuli, the feedback indicates that a Category A response was correct with some probability between 0 and 1. Tasks such as this, in which the relationship between stimuli and category membership is probabilistic, are known in the literature as probabilistic classification tasks. Although the majority of category-learning studies have used deterministic tasks, probabilistic classification also has a long history (Estes, Campbell, Hatsopoulos, & Hurwitz, 1989; Gluck & Bower, 1988; Kubovy & Healy, 1977).

Recently, a probabilistic classification task called the weather-prediction task has become especially popular in the neuropsychology literature (Knowlton, Squire, & Gluck, 1994). On each trial of the weather-prediction task, participants are presented with one, two, or three out of four possible cards and are asked to use this information to predict the weather (rain or shine; Knowlton et al., 1994). Each card displays a highly discriminable pattern, which, by itself, predicts sunshine 75%, 57%, 43%, or 25% of the time (in the original application; see, e.g., Knowlton, Mangels, & Squire, 1996). As is the case in information-integration tasks, optimal accuracy (i.e., approximately 76% correct) can only be achieved by integrating the information across the different cards. Even so, participants can achieve almost optimal accuracy with an explicit strategy in which they simply respond on the basis of the presence or absence of the most informative card (i.e., approximately 73% correct). Because the weather-prediction task is probabilistic, and optimal accuracy requires information integration, it has often been assumed that learning in the task is mediated by a single system that recruits some implicit, incremental process (Knowlton et al., 1996; Weickert et al., 2002; Witt, Nuhman, & Deuschl, 2002). However, our results suggest that explicit strategies should be common in the weather-prediction task, because a simple explicit strategy is nearly optimal. In fact, recent strategy analyses indicate that, at least initially, learning in the weather-prediction task is dominated by explicit strategies (Gluck, Shohamy, & Myers, 2002). This result, together with the results described here, suggests that knowing whether a category-learning task is deterministic or probabilistic, by itself, provides little information about how people will learn the task. Instead, our results show that it is critical to know whether a rule-based or information-integration task was used, and in the latter case, whether there exists some simple explicit strategy that is nearly optimal.

## REFERENCES

- AKAIKE, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, **19**, 716-723.
- ALLEN, S. W., & BROOKS, L. R. (1991). Specializing the operation of an explicit rule. *Journal of Experimental Psychology: General*, **120**, 3-19.
- ASHBY, F. G. (1992a). Multidimensional models of categorization. In F. G. Ashby (Ed.), *Multidimensional models of perception and cognition* (pp. 449-483). Hillsdale, NJ: Erlbaum.
- ASHBY, F. G. (1992b). Multivariate probability distributions. In F. G. Ashby (Ed.), *Multidimensional models of perception and cognition* (pp. 1-34). Hillsdale, NJ: Erlbaum.
- ASHBY, F. G., ALFONSO-REESE, L. A., TURKEN, A. U., & WALDRON, E. M. (1998). A neuropsychological theory of multiple systems in category learning. *Psychological Review*, **105**, 442-481.
- ASHBY, F. G., & ELL, S. W. (2001). The neurobiology of human category learning. *Trends in Cognitive Sciences*, **5**, 204-210.
- ASHBY, F. G., & GOTT, R. E. (1988). Decision rules in the perception and categorization of multidimensional stimuli. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **14**, 33-53.
- ASHBY, F. G., & MADDOX, W. T. (1992). Complex decision rules in categorization: Contrasting novice and experienced performance. *Journal of Experimental Psychology: Human Perception & Performance*, **18**, 50-71.
- ASHBY, F. G., & MADDOX, W. T. (2005). Human category learning. *Annual Review of Psychology*, **56**, 149-178.
- ASHBY, F. G., QUELLER, S., & BERRETTY, P. M. (1999). On the dominance of unidimensional rules in unsupervised categorization. *Perception & Psychophysics*, **61**, 1178-1199.
- ASHBY, F. G., & WALDRON, E. M. (1999). On the nature of implicit categorization. *Psychonomic Bulletin & Review*, **6**, 363-378.
- ASHBY, F. G., WALDRON, E. M., LEE, W. W., & BERKMAN, A. (2001). Suboptimality in human categorization and identification. *Journal of Experimental Psychology: General*, **130**, 77-96.
- BRAINARD, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, **10**, 433-436.
- BROOKS, L. R. (1978). Nonanalytic concept formation and memory for instances. In E. Rosch & B. B. Lloyd (Eds.), *Cognition and categorization* (pp. 169-211). Hillsdale, NJ: Erlbaum.
- ERICKSON, M. A., & KRUSCHKE, J. K. (1998). Rules and exemplars in category learning. *Journal of Experimental Psychology: General*, **127**, 107-140.
- ESTES, W. K., CAMPBELL, J. A., HATSOPoulos, N., & HURWITZ, J. B. (1989). Base-rate effects in category learning: A comparison of parallel network and memory storage-retrieval models. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **15**, 556-571.
- GARNER, W. R. (1974). *The processing of information and structure*. Potomac, MD: Erlbaum.
- GLUCK, M. A., & BOWER, G. H. (1988). From conditioning to category learning: An adaptive network model. *Journal of Experimental Psychology: General*, **117**, 227-247.
- GLUCK, M. A., SHOHAMY, D., & MYERS, C. (2002). How do people solve the "weather prediction" task?: Individual variability in strategies for probabilistic category learning. *Learning & Memory*, **9**, 408-418.
- HAYES, N. A., & BROADBENT, D. E. (1988). Two modes of learning for interactive tasks. *Cognition*, **28**, 249-276.
- KEMLER NELSON, D. G. (1993). Processing integral dimensions: The whole view. *Journal of Experimental Psychology: Human Perception & Performance*, **19**, 1105-1113.
- KENDLER, H. H., & KENDLER, T. S. (1962). Vertical and horizontal processes in problem solving. *Psychological Review*, **69**, 1-16.
- KNOWLTON, B. J., MANGELS, J. A., & SQUIRE, L. R. (1996). A neostriatal habit learning system in humans. *Science*, **273**, 1399-1402.
- KNOWLTON, B. J., SQUIRE, L. R., & GLUCK, M. A. (1994). Probabilistic classification learning in amnesia. *Learning & Memory*, **1**, 106-120.
- KUBOVY, M., & HEALY, A. F. (1977). The decision rule in probabilistic categorization: What it is and how it is learned. *Journal of Experimental Psychology: General*, **106**, 427-446.
- MADDOX, W. T., & ASHBY, F. G. (1993). Comparing decision bound and exemplar models of categorization. *Perception & Psychophysics*, **53**, 49-70.

- MADDOX, W. T., & ASHBY, F. G. (2004). Dissociating explicit and procedural-learning based systems of perceptual category learning. *Behavioural Processes*, **66**, 309-332.
- MADDOX, W. T., BOHIL, C. J., & ING, A. D. (2004). Evidence for a procedural-learning-based system in perceptual category learning. *Psychonomic Bulletin & Review*, **11**, 945-952.
- MEDIN, D. L., & SCHAFER, M. M. (1978). Context theory of classification learning. *Psychological Review*, **85**, 207-238.
- MEDIN, D. L., WATTENMAKER, W. D., & HAMPSON, S. E. (1987). Family resemblance, conceptual cohesiveness, and category construction. *Cognitive Psychology*, **19**, 242-279.
- NOSOFKY, R. M., & KRUSCHKE, J. K. (2002). Single-system models and interference in category learning: Commentary on Waldron and Ashby (2001). *Psychonomic Bulletin & Review*, **9**, 169-174.
- NOSOFKY, R. M., PALMERI, T. J., & MCKINLEY, S. C. (1994). Rule-plus-exception model of classification learning. *Psychological Review*, **101**, 53-79.
- PELLI, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, **10**, 437-442.
- POTHOS, E. M. (2005). The rules versus similarity distinction. *Behavioral & Brain Sciences*, **28**, 1-49.
- SCHWARZ, G. (1978). Estimating the dimension of a model. *The Annals of Statistics*, **6**, 461-464.
- SHAW, M. L. (1982). Attending to multiple sources of information: I. The integration of information in decision making. *Cognitive Psychology*, **14**, 353-409.
- SLOMAN, S. A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin*, **119**, 3-22.
- WEICKERT, T. W., TERRAZAS, A., BIGELOW, L. B., MALLEY, J. D., HYDE, T., EGAN, M. F., ET AL. (2002). Habit and skill learning in schizophrenia: Evidence of normal striatal processing with abnormal cortical input. *Learning & Memory*, **9**, 430-442.
- WICKENS, T. D. (1982). *Models for behavior: Stochastic processes in psychology*. San Francisco: Freeman.
- WILLINGHAM, D. B. (1998). A neuropsychological theory of motor skill learning. *Psychological Review*, **105**, 558-584.
- WITT, K., NUHSMAN, A., & DEUSCHL, G. (2002). Dissociation of habit-learning in Parkinson's and cerebellar disease. *Journal of Cognitive Neuroscience*, **14**, 493-499.
- ZAKI, S. R., NOSOFKY, R. M., JESSUP, N. M., & UNVERSAGT, F. W. (2003). Categorization and recognition performance of a memory-impaired group: Evidence for single-system models. *Journal of the International Neuropsychological Society*, **9**, 394-406.
- ZAKI, S. R., NOSOFKY, R. M., STANTON, R. D., & COHEN, A. L. (2003). Prototype and exemplar accounts of category learning and attentional allocation: A reassessment. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **29**, 1160-1173.

## NOTES

1. The rule-based categories were generated by rotating the Figure 1 categories 45° counterclockwise—a procedure that guarantees that optimal accuracy, within-category scatter, and category coherence were equivalent across the rule-based and information-integration tasks.

2. Note that because category overlap was increased by simply decreasing the distance between the category means (while equating the respective distances between the category boundary), both the optimal perceptual-integration strategy and the most accurate explicit strategy are constant across conditions. As a result, the regions of stimulus space for which explicit and perceptual-integration strategies differ in their predictions also remain constant across levels of category overlap.

3. The accuracy rates were obtained by computing the percentage of correct classifications predicted by the most accurate explicit strategy in the absence of internal noise. Specifically, the most accurate explicit strategy (e.g., the dashed vertical line in Figure 1) assumes that the participant sets a single criterion on the spatial frequency dimension at 110 units and responds A or B for stimuli less than or greater than this criterion, respectively. Note that there exists a similar explicit strategy defined along the orientation dimension that predicts identical accuracy rates.

4. The statistic used for model comparison (BIC; see Appendix B for details) tends to favor less complex models (e.g., optimal models). Repeating the analysis using a statistic that decreases the penalty for model complexity ( $AIC = 2r - 2\ln L$ ; Akaike, 1974) did not alter the distribution of explicit and perceptual-integration models across the five experimental conditions.

5. A full description of COVIS's dynamics is beyond the scope of the discussion, but the interested reader should see Ashby et al. (Ashby et al., 1998).

## APPENDIX A

Spatial Frequency (Arbitrary Units)	Orientation (Degrees From Horizontal)
Experiment 1: Low Condition	
60	100
50	65
80	90
90	105
150	165
120	130
120	160
130	175
100	90
75	55
105	70
85	80
160	120
125	115
150	140
180	160
Experiment 3: Medium-Low Condition	
108	70
112	80
107	90
113	100
108	110
112	120
107	130
113	140

## APPENDIX B

This appendix briefly describes the decision bound models. For more details, see Ashby (1992a) or Maddox and Ashby (1993). The classification of these models as either *perceptual-integration* or *explicit models* is designed to reflect current theories of how these strategies are learned (see, e.g., Ashby et al., 1998) and has received considerable empirical support (see Ashby & Maddox, 2005, and Maddox & Ashby, 2004, for reviews).

**Perceptual-Integration Models**

**The general linear classifier (GLC).** This model assumes that the decision bound between each pair of categories is linear and requires the integration of perceived spatial frequency and orientation. The GLC has three parameters (slope and intercept of the linear bound and  $\sigma^2$ ). In the information-integration task of Experiments 1 and 2, a special case of the GLC assumes that participants use the linear bound that maximizes accuracy (the diagonal bounds shown in Figure 1). This model has only one free parameter ( $\sigma^2$ ) and is referred to as the optimal perceptual-integration model.

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

**Explicit Models**

Two models assumed that participants used an explicit strategy.

**The one-dimensional model.** This model assumes that the participant sets a criterion on a single perceptual dimension and then makes an explicit decision about the level of the stimulus on that dimension (Ashby & Gott, 1988; Shaw, 1982). Two versions of the one-dimensional model were fit to these data: One assumed that participants attended selectively to spatial frequency and the other assumed that participants attended selectively to orientation. The one-dimensional models have two free parameters: a decision criterion on the relevant perceptual dimension and the variance of internal (perceptual and criterial) noise ( $\sigma^2$ ). In the rule-based task of Experiment 3, a special case of the one-dimensional model assumes that participants use the one-dimensional decision bound that maximizes accuracy (the vertical bounds shown in Figure 2). This special case has only one free parameter (i.e.,  $\sigma^2$ ).

---

**APPENDIX B (Continued)**


---

**Conjunction models.** Although the one-dimensional strategy will maximize accuracy, another type of explicit strategy available to participants is a conjunction strategy. As is the case with perceptual-integration strategies, conjunction strategies also require the integration of spatial frequency and orientation information. For example, a participant might set a criterion along the spatial frequency dimension to determine if the stimulus is high or low in spatial frequency and set a separate criterion on orientation, to determine if the angle is shallow or steep. The results of these independent decision processes might then be combined to make a response—for example, “Respond A if the stimulus is low and shallow.” Although conjunction strategies require integration, they differ from perceptual-integration strategies in that the integration is postdecisional. In other words, decisions are made about the stimulus value on each dimension, and the output of these decisions is explicitly integrated to generate a category response. Indeed, recent evidence supports this distinction between conjunction and perceptual-integration strategies (Maddox, Bohil, & Ing, 2004).

Conjunction models have three parameters (a criterion on each dimension, and  $\sigma^2$ ). Based on inspection of the data from the individual participants, two versions of the conjunction model were fit to these data. The first assumed that individuals assigned a stimulus to Category B if it was high in spatial frequency and low in orientation (i.e., the bars are thin and shallow); otherwise the stimulus was assigned to Category A. The second conjunction model assumed that a stimulus was assigned to Category A if it was low in spatial frequency and high in orientation (i.e., the bars are thick and steep); otherwise the stimulus was assigned to Category B.

**Model Fitting**

The model parameters were estimated using maximum likelihood (Ashby, 1992b; Wickens, 1982) and the goodness-of-fit statistic was

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

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

---

(Manuscript received March 25, 2005;  
revision accepted for publication October 7, 2005.)