

# The effects of positive versus negative feedback on information-integration category learning

F. GREGORY ASHBY AND JEFFREY B. O'BRIEN  
University of California, Santa Barbara, California

A number of studies have shown that in category learning, providing feedback about errors allows faster learning than providing feedback about correct responses. However, these previous studies used explicit, rule-based tasks in which the category structures could be separated by a simple rule that was easily verbalized. Here, the results of the first experiment known to compare the efficacy of positive versus negative feedback during information-integration category learning are reported. Information-integration tasks require participants to integrate perceptual information from incommensurable dimensions, and evidence suggests that optimal responding recruits procedural learning. The results show that although nearly all of the full-feedback control participants demonstrated information-integration learning, participants receiving either positive-only or negative-only feedback generally used explicit, rule-based strategies. It thus appears that, unlike rule-based learning, consistent information-integration learning requires full feedback. The theoretical implications of these findings for current models of information-integration learning are discussed.

To categorize is to respond differently to objects or events in separate classes or categories. Humans typically learn thousands of categories in their lifetimes. One critical feature that enables much of this learning is feedback. Positive feedback signals that a categorization response was correct, whereas negative feedback signals an incorrect categorization response. In some cases, positive and negative feedback are provided by a teacher (e.g., mother, schoolteacher), but many times, feedback is provided automatically by the environment. For example, we know whether we have correctly categorized a white, powdery substance as sugar as soon as we place it on our tongues.


Despite the prevalence of feedback-mediated learning in everyday life, the relative efficacy of positive versus negative feedback has been studied only in rule-based category learning tasks—that is, tasks in which the category structures can be separated by a simple rule that is easily verbalized (e.g., a one-dimensional rule). The consensus from these studies is that negative feedback is more effective than positive feedback. Difficult real-life categorization tasks, however, typically cannot be mastered via simple rules. For example, no simple verbal rule allows a radiologist to separate x-rays that display a tumor from those that do not. Category-learning tasks that require participants to integrate perceptual information from incommensurable dimensions are called *information-integration* (II) tasks (Ashby & Ell, 2001). Although such tasks occur frequently in real life, no studies have compared the efficacy of positive versus negative feedback in II category learning. Here, we report the results of the first such study. Nineteen participants completed a total of 57,000 trials

over five experimental sessions. Our results suggest that unlike rule-based category learning, both positive and negative feedback are critical for II category learning.

## Two Category-Learning Tasks

Figure 1 shows the category structures and some of the stimuli used in the experiment reported in this article. Each stimulus was a circular sine-wave grating whose dark and light bars varied in orientation and width across trials. On each trial of the experiment, a single disk was shown to the observer, whose task was to assign that disk to category A or B by pressing the appropriate response key. The category boundary is denoted by the solid diagonal line. This is an II category-learning task (Ashby & Ell, 2001) because accuracy is maximized only if information from the two incommensurable stimulus dimensions is integrated at some predecisional stage. Note that the category bound is difficult or impossible to describe verbally (Ashby, Alfonso-Reese, Turken, & Waldron, 1998). Even so, many previous studies have shown that with full (i.e., both positive and negative) and immediate feedback, healthy young adults can reliably learn such categories (see, e.g., Ashby & Gott, 1988; Ashby & Maddox, 1990, 1992; Ashby & Waldron, 1999).

The same Figure 1 stimuli could be used to create a rule-based category-learning task simply by rotating the category bound to either vertical or horizontal. For example, with a vertical bound, the optimal strategy would be to respond “B” for thick bars and “A” for thin bars. More difficult rule-based tasks require attention to two or more dimensions. For example, the correct rule might be a



**Figure 1.** Category structures that might be used in an information-integration category-learning task. The diagonal line indicates the optimal decision bound.

conjunction of the type “the stimulus belongs in category A if the bars are thick and the orientation is steep.” (Note that the diagonal bound shown in Figure 1 cannot be described as a simple conjunction.) The key requirement for rule-based tasks is that the correct categorization rule can be discovered via an explicit reasoning process. Virtually all category-learning tasks used in neuropsychological assessment are rule based, including the widely known Wisconsin Card Sorting Test (Heaton, 1981).

Many recent studies suggest that learning in rule-based and II tasks is largely mediated by functionally separate category-learning systems. Evidence supporting this hypothesis comes from behavioral dissociations between the two tasks in healthy young adults (Ashby, Ell, & Waldron, 2003; Ashby, Maddox, & Bohil, 2002; Ashby, Queller, & Berretty, 1999; Maddox, Ashby, & Bohil, 2003; Maddox, Ashby, Ing, & Pickering, 2004; Waldron & Ashby, 2001; Zeithamova & Maddox, 2006), studies of various neuropsychological patient groups (Ashby, Noble, Filoteo, Waldron, & Ell, 2003; Filoteo, Maddox, & Davis, 2001; Filoteo, Maddox, Salmon, & Song, 2005; Smith, Tracy, & Murray, 1993), and neuroimaging studies (Nomura et al., 2006; Seger & Cincotta, 2002; Smith, Patalano, & Jonides, 1998). For example, evidence suggests that in rule-based tasks, abstract category labels are learned (Ashby et al., 2003), and such learning requires working memory and executive attention (Waldron & Ashby, 2001; Zeithamova & Maddox, 2006). In addition, although feedback processing in rule-based tasks requires attention and effort (Maddox, Ashby, et al., 2004), the nature and timing of the feedback are not critical (Ashby et al., 2002; Maddox et al., 2003). In contrast, evidence suggests that because response positions are learned in II tasks (Ashby et al., 2003; Maddox, Bohil, & Ing, 2004), such learning places few demands on working memory or executive attention (Waldron & Ashby, 2001; Zeithamova & Maddox, 2006), whereas feedback processing, although largely automatic (Maddox, Ashby,

et al., 2004), is quite sensitive to the nature and timing of the feedback (Maddox et al., 2003).

### Feedback Studies

A long history of research has investigated the relative efficacy of positive and negative feedback. For example, early two-choice discrimination-learning studies with rats found that punishment-only training caused faster learning than reward-only training (see, e.g., Hoge & Stocking, 1912; Warden & Aylesworth, 1926). The first human studies, which used simple two-choice rule-based category learning tasks, also found that negative feedback was more effective than positive feedback (see, e.g., Buss & Buss, 1956; Buss, Weiner, & Buss, 1954; Meyer & Offenbach, 1962). More recently, however, Frank, Seeberger, and O'Reilly (2004) reported that dopamine replacement medications reversed this effect in Parkinson's disease patients (i.e., positive feedback became more effective than negative feedback). Several researchers hypothesized that the more commonly observed negative feedback advantage occurs because positive feedback is less informative than negative feedback, at least in two-choice tasks (Buchwald, 1962; Jones, 1961; Meyer & Offenbach, 1962). The idea is that negative feedback informs the participant that his or her hypothesis was incorrect and also signals which response was correct (i.e., the other response), whereas positive feedback signals only that the response was correct (i.e., the hypothesis might have been incorrect, but, by chance, the response was correct). With more than two categories, negative feedback loses some of this advantage. This information asymmetry hypothesis was supported by results from a four-category study that found no difference between negative and positive feedback (Buss & Buss, 1956).

In all of these studies, however, every error received feedback in the negative feedback conditions, and every correct response received feedback in the positive feedback conditions. Thus, participants might have inferred

that the absence of feedback in the negative conditions signaled a correct response and that feedback absence in the positive conditions signaled an error. Waldron (2000) addressed this problem by designing a negative (positive) condition in which feedback was provided on only a random proportion of error (correct) trials, and no feedback was given on correct (error) trials. Thus, for trials on which no feedback was given, no inference about the accuracy of the response was possible (and participants were told this). Waldron used an II task, but, unfortunately, because optimal accuracy was very low (70%), all participants used rule-based strategies—that is, none of them learned the category structures. This was true even for the full-feedback control group. As a result, Waldron’s results tell us little about the relative efficacy of positive versus negative feedback in II category learning.

The COVIS model of category learning (Ashby et al., 1998; Ashby & Waldron, 1999) hypothesizes that category learning is a competition between separate rule-based and procedural-learning systems. The rule-based system is a logical reasoning system that uses explicit strategies and is assumed to control performance in rule-based tasks. The procedural-learning system uses perceptual integration strategies and is assumed to dominate in II tasks. COVIS predicts, and recent evidence suggests, that rule-based strategies dominate in II tasks if optimal accuracy is either very high or very low (Ell & Ashby, 2006). Consistent II learning requires an intermediate level of category separation that is higher than that used by Waldron (2000). In the present study, we adjusted the category overlap so that optimal accuracy was 86%. This corresponds to  $d' = 2.16$  (i.e., the distance between category means divided by the common standard deviation), which is theoretically important because a plot of accuracy versus  $d'$  is steepest at exactly this value. In other words, small changes in

strategy have the greatest effect on performance (i.e., accuracy) when  $d' = 2.16$ . For this reason, COVIS predicts that II learning is most likely when  $d' = 2.16$ .


An abstract representation of the categories that we used is shown in Figure 2. Each plus symbol denotes the bar width and orientation of an exemplar from category A, and each circle symbol denotes these values for the exemplars from category B. Note that the categories overlap, so perfect accuracy is impossible. In fact, as mentioned above, optimal accuracy is 86%.

Each of 19 participants completed 3,000 categorization trials with these categories (over five experimental sessions). Each participant was assigned either to one of two experimental conditions or to one of two control conditions. Following Waldron (2000), the positive feedback (PFB) participants were given feedback only about some of their correct responses, whereas the negative feedback (NFB) participants were given feedback only about some of their incorrect responses. We included two separate control conditions. In the control (partial) condition, participants received feedback about both correct and incorrect responses, but only on some trials, whereas in the control (full) condition, feedback was given after each trial. In each of the three partial feedback conditions, an adaptive algorithm (described below) was used to ensure that every participant received about the same number of feedback trials each session, regardless of condition.

**METHOD**

**Participants**

Nineteen young adults from the University of California at Santa Barbara community participated in this experiment (15 graduate and 4 undergraduate students); 4 in the control (full) condition, 5 in the control (partial) con-



**Figure 2.** Category structures used for this experiment. The diagonal line indicates the optimal decision bound.

dition, and 5 each in the PFB and NFB conditions. All participants had normal or corrected-to-normal vision. Each participant completed 5 sessions of approximately 45 min duration each. Participants were paid \$15 per session.

### Stimuli and Stimulus Generation

Each exemplar was a circular sine-wave grating of the type shown in Figure 1. The stimuli varied across trials on two dimensions: bar width and bar orientation. Category exemplars were generated using a randomization technique (Ashby & Gott, 1988) in which each category is defined as a bivariate normal distribution along the two stimulus dimensions. On each trial, a random sample  $(x, y)$  was drawn from either the category A or the category B distribution (each with probability .5). A stimulus was constructed from each sample by defining the stimulus orientation (in degrees counterclockwise from horizontal) as  $\theta = x - 30$  degrees, and the stimulus spatial frequency (in cycles per degree) as  $f = (y \cdot .001) + .01$ . The scaling factors were chosen in an attempt to roughly equate the salience of each dimension. Category A had an orientation mean of 34 and a bar width mean of 56, whereas category B had means on these two dimensions of 56 and 34, respectively. In both categories, the variance on each dimension was 235, and the covariance between dimensions was 19.

Three hundred random samples were drawn from each of these category distributions. For each category, the resulting sample was then linearly transformed so that the sample means, variances, and covariance exactly equaled the population means, variances, and covariance. The transformed samples are shown in Figure 2. The optimal category bound is denoted by the diagonal line. A participant responding "A" to any exemplar above this bound and "B" to any exemplar below it would achieve 86% accuracy. The overlap also served to guarantee incorrect responses, allowing for the administration of equivalent numbers of feedback trials in each of the restricted feedback conditions.

### Procedure

Each participant was randomly assigned to one of four feedback conditions: PFB (feedback for correct responses only), NFB (feedback for incorrect responses only), control (partial) (feedback for both correct and incorrect responses, but only on some trials), and control (full) (feedback on every trial). For the NFB condition, no feedback was given following a correct response. Following each error, feedback was given with probability 0.8. For the PFB condition, no feedback was given following an incorrect response. For correct responses, an adaptive procedure was used that provided about the same number of feedback trials as in the NFB condition. Specifically, if a correct response occurred on trial  $n$ , then feedback was given with probability

$$P(\text{feedback} \mid \text{correct response}) = \frac{.8Q(\text{errors on last 50 trials})}{Q(\text{correct on last 50 trials})}$$

where  $P$  is probability and  $Q$  is proportion. On the first 50 trials, we assumed chance performance; thus, participants

in both restricted feedback conditions received feedback on these stimulus presentations at a rate of 80% on relevant trials. For participants in the control (partial) condition, feedback was given for both correct and incorrect responses but at a rate that provided about the same amount of feedback per session as in the NFB and PFB conditions. Specifically, the probability of feedback on each trial was  $P(\text{feedback}) = .8Q(\text{errors on last 50 trials})$ . Participants in the control (full) condition received feedback on every trial.

Each participant completed five experimental sessions on separate days. Each experimental session included 600 categorization trials. Participants were instructed to categorize disk patterns on the basis of the orientation and bar width of the disk presented on each trial. Participants in the control (full) condition were instructed that they would receive feedback after every trial. Participants in the control (partial) condition were instructed that they would receive feedback on a proportion of both correct and incorrect response trials. In the PFB and NFB conditions, participants were instructed that they would receive feedback on only a proportion of the trials for which they made correct and incorrect responses, respectively. Participants in these latter three conditions (i.e., control [partial], PFB, and NFB) were instructed that on trials in which no feedback was provided, they could be correct or incorrect and that any assumption about their performance on the no-feedback trials would not help them learn the categories.

On each trial, a single stimulus was presented at fixation and the participant was instructed to make a category assignment by pressing one of two response keys (labeled A or B) with the index finger. The fixation point was presented on a gray background at the center of the screen for 500 msec, then the stimulus appeared and remained on the screen until the participant responded. Feedback was immediate (on relevant trials) and took the form of a 500-Hz tone lasting 500 msec for correct responses or a 200-Hz tone lasting 500 msec for incorrect responses. For blank trials in which no feedback was given, nothing was presented during the 500-msec feedback interval. After the feedback interval, a blank screen was presented for 1,000 msec, then the next trial began. Each session consisted of twelve 50-trial blocks, between which a participant-controlled rest period was provided. During each session, the same 600 stimuli were presented in different random orders.

## RESULTS

Two types of analyses were performed: analyses of the accuracy data using standard statistical techniques and model-based analyses to determine the decision strategies used by participants. The model-based analyses are important because participants may have improved their accuracy in a condition, but not by using a decision strategy that integrated information from the two stimulus dimensions. For example, a one-dimensional rule of the type "respond 'A' if the bars are narrow and 'B' if the bars are wide" can cause accuracy to be well above chance (i.e., 77.8% correct). Such a rule does not require information integration, and there is considerable evidence that the learning of explicit rules of this type is mediated by a dif-

ferent system than the one that mediates the learning of strategies that require information integration (see Ashby & Maddox, 2005, for a review).

**Accuracy-Based Analyses**

The adaptive feedback procedure used in this study was designed to provide approximately equal amounts of feedback in the various partial feedback conditions, under the null hypothesis that no accuracy differences existed among the groups. Specifically, in all partial feedback conditions, feedback was given at a rate equal to 80% of the current error rate. However, this algorithm guaranteed that participants who performed better received less overall feedback. Therefore, our first analysis was to determine whether participants in different feedback groups received a different number of feedback trials.

We conducted a mixed-design ANOVA to determine whether participants in different conditions received a different number of feedback trials. Not surprisingly, given the performance-based nature of the feedback algorithm, there were significant main effects of session [ $F(4,48) = 12.64, p < .001$ ] and condition [ $F(2,12) = 4.40, p < .05$ ], but the interaction between condition and session was not significant [ $F(8,48) = .695, p > .05$ ]. Additional analyses showed that participants in the control (partial) group received significantly fewer feedback trials than participants in the PFB group [ $F(1,8) = 12.72, p < .01$ ] and fewer feedback trials than participants in the NFB group, although this latter result was not significant [ $F(1,8) = .95, p > .05$ ]. Importantly, there was no significant difference in the number of feedback trials between participants in the PFB and NFB groups [ $F(1,8) = 2.86, p > .05$ ].

Figure 3 shows average learning curves for all participants. The learning curves suggest that learning occurred for each group over the course of the five training sessions. Pairwise sign tests conducted on the thirty 100-trial learning blocks showed no difference between the PFB and NFB groups (sign test:  $S = 19$  of 30 blocks,  $p > .05$ ). As expected, the control (full) group performed significantly better than both the PFB (sign test:  $S = 28$  of 30 blocks,  $p < .001$ ) and NFB groups (sign test:  $S = 28$  of 30 blocks,  $p < .001$ ). The control (partial) group also outperformed both the PFB (sign test:  $S = 30$  of 30 blocks,  $p < .001$ ) and NFB groups (sign test:  $S = 28$  of 30 blocks,  $p < .001$ ), and there was no significant difference between the two control groups (sign test:  $S = 14$  of 30,  $p > .10$ ). Despite being outperformed by participants in both control conditions, the PFB and the NFB groups both performed significantly above chance on the final session (PFB:  $Z = 12.25, p < .001$ ; NFB:  $Z = 11.76, p < .001$ ), suggesting that learning did occur for these groups over the course of the experiment.

Since the average number of feedback trials for each participant in the partial feedback conditions was 750 (150 feedback trials per session  $\times$  5 sessions), a separate analysis of only the first 750 trials of the control (full) group was conducted. In all of these comparisons, we ignored no-feedback trials. Thus, a comparison of the control (full) and PFB groups asked whether these groups differed in accuracy after each had completed  $n$  feedback




Figure 3. Group learning curves by condition.

trials, for values of  $n$  from 1 to 750. It is important to keep in mind that for the PFB group, these 750 feedback trials were intermixed with many no-feedback trials over the course of five sessions, whereas for the control (full) group, these were the first 750 trials during Sessions 1 and 2. With this in mind, sign tests showed that the performance of the control (full) group on the first 750 trials was not significantly different from the performance of the PFB (sign test:  $S = 18$  of 30 blocks,  $p > .05$ ) or NFB group (sign test:  $S = 17$  of 30 blocks,  $p > .05$ ). Interestingly, however, the performance of the control (partial) group was significantly better than the control (full) group (sign test:  $S = 20$  of 30 blocks,  $p < .05$ ).

The advantage of the control (partial) group over the PFB and NFB groups suggests that receiving both positive and negative feedback, as opposed to only one feedback type, may have improved performance. The superior performance of the control (partial) group is especially impressive since that group received fewer feedback trials than either the PFB or NFB group. On the other hand, the superior performance of the control (partial) group compared with that of the control (full) group (on their first 750 trials) suggests that participants in the control (partial) condition received some advantage from the no-feedback trials. There are several possibilities. First, the feedback received by the control (partial) group was distributed over 5 different days, whereas the feedback received by the control (full) group on their first 750 trials was distributed over only 2 days. Thus, the control (partial) group had more opportunity for between-session consolidation than did the control (full) group. Second, the control (partial) group may have learned on no-feedback trials. Previous research has found no evidence that II learning is possible in the absence of feedback (Ashby et al., 1999). Even so, unsupervised rule-based learning is sometimes possible (Ashby et al., 1999), so one possibility is that the extra no-feedback trials experienced by the control (partial) group allowed them to give up on their search for a satisfactory explicit rule more quickly. Note that the advantage of the control (partial) group over the

control (full) group (on the first 750 trials) is unlikely to be the result of a spaced versus massed learning effect, since participants in both groups gave a categorization response on every trial (and they could not predict which responses would receive feedback). Thus, neither group experienced spaced training.

### Model-Based Analyses

Our accuracy-based analyses suggest that the two control groups [control (full) and control (partial)] outperformed the PFB and NFB groups, but that the accuracy of the PFB and NFB groups was significantly above chance. Therefore, participants did learn from either positive-only or negative-only feedback. However, it is impossible to tell from these analyses exactly what they learned. In particular, before we can conclude that II learning is possible with positive-only or negative-only feedback, we must be certain that participants in the PFB and NFB conditions used II strategies. To answer this question, a variety of different decision-bound models (see, e.g., Ashby & Gott, 1988; Maddox & Ashby, 1993) were fit to the data from each participant's last session. Decision-bound models assume that 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 emits the associated response.

Two different types of models were fit to each participant's responses (see the Appendix for details). One type was compatible with the assumption that participants used a rule-based strategy and one type assumed an II strategy. These models make no detailed process assumptions, since a number of different process accounts are compatible with each of the models (see, e.g., Ashby, 1992; Maddox & Ashby, 1993). For example, if an II model fits significantly better than a rule-based model, then we can be confident that participants did not use a rule-based strategy, but we cannot specify which specific II strategy was used (e.g., a weighted combination of the two dimensions vs. more holistic processing).

Model parameters were estimated using the method of maximum likelihood, and the statistic used for model selection was the Bayesian information criterion (BIC; Schwarz, 1978), which is defined as  $BIC = r \ln N - 2 \ln L$ , where  $r$  is the number of free parameters,  $N$  is the sample size, and  $L$  is the likelihood of the model given the data. The BIC statistic penalizes models for extra free parameters. To determine the best-fitting model within a group of competing models, the BIC statistic is computed for each model, and the model with the smallest BIC value is chosen.

Figures 4–7 show the final-session response plots for all participants in the PFB, NFB, control (partial), and control (full, first 750 trials) groups, along with the decision bounds from the best-fitting models. The BIC values across all tested models and for all participants are shown in Table 1, along with the percentage of responses accounted for by the best-fitting bounds. These model fits illustrate several important results. First, all but one set of control data (i.e., from the three control groups) was best fit by a model that assumed information integration (i.e., the general linear classifier). In fact, in each of these 12

control data sets, the best-fitting model assumed a bound of the same type as the optimal bound shown in Figure 1. Figures 6 and 7 and Table 1 make it clear that the best-fitting bounds account for the vast majority of each control participant's responses, which suggests a high degree of consistency in response strategy. In summary, when feedback was given about both correct and incorrect responses, participants consistently used II strategies that were nearly optimal, even when this feedback was given only on some trials.

Second, the data from 4 of the 5 participants in both the PFB and NFB groups were best fit by a model that assumed explicit rule use—i.e., either participants simply set a criterion on one stimulus dimension and ignored the other, or they used a conjunction rule of the type “respond ‘A’ if the bar is thin and the orientation is steep; otherwise, respond ‘B.’” In the latter case, decisions about each dimension are made separately, and information is not perceptually integrated across these dimensions. Considerable evidence suggests that the use of conjunction rules of this type is mediated by the same explicit, rule-based system responsible for the selection of one-dimensional rules (Maddox & Ashby, 2004; Maddox, Bohil, & Ing, 2004; Maddox, Filoteo, Hejl, & Ing, 2004; Maddox & Ing, 2005).

Table 1 and Figures 4 and 5 reveal another subtle difference between the PFB and NFB groups. Consider the single participant in each group whose data were best fit by a model that assumed information integration. In the PFB group, approximately 93% of Participant 3's responses were accounted for by the decision bound assumed by the best-fitting II model. This suggests consistent application of an II strategy and provides strong evidence of II learning by this participant, which suggests that II learning is possible in the presence of positive feedback only, at least by some participants. In contrast, many of the responses of the NFB participant whose data were best fit by an II model (Participant 4) were not accounted for by the best-fitting II strategy (23.5%). Although some II learning by this participant cannot be ruled out, the poor fit of the best-fitting II model is also consistent with the hypothesis that the participant relied on a mixture of rule-based strategies.

Model fits of the first four sessions support the reliability of the final-session model fits—in general, the percentage of responses accounted for by the best-fitting models increased over sessions for all participants. Data from participants in the partial feedback conditions that were best fit by a rule-based model in the final session either were best fit by rule-based models in the preceding sessions or were poorly accounted for by any model (i.e., they had a low percentage of responses accounted for by the best-fitting model). Importantly, all participants in the control conditions were best fit by II models during at least one of the first four sessions.

## DISCUSSION

This study was designed to investigate the relative efficacies of positive and negative feedback in an II category-learning task in which the amount and type of feedback

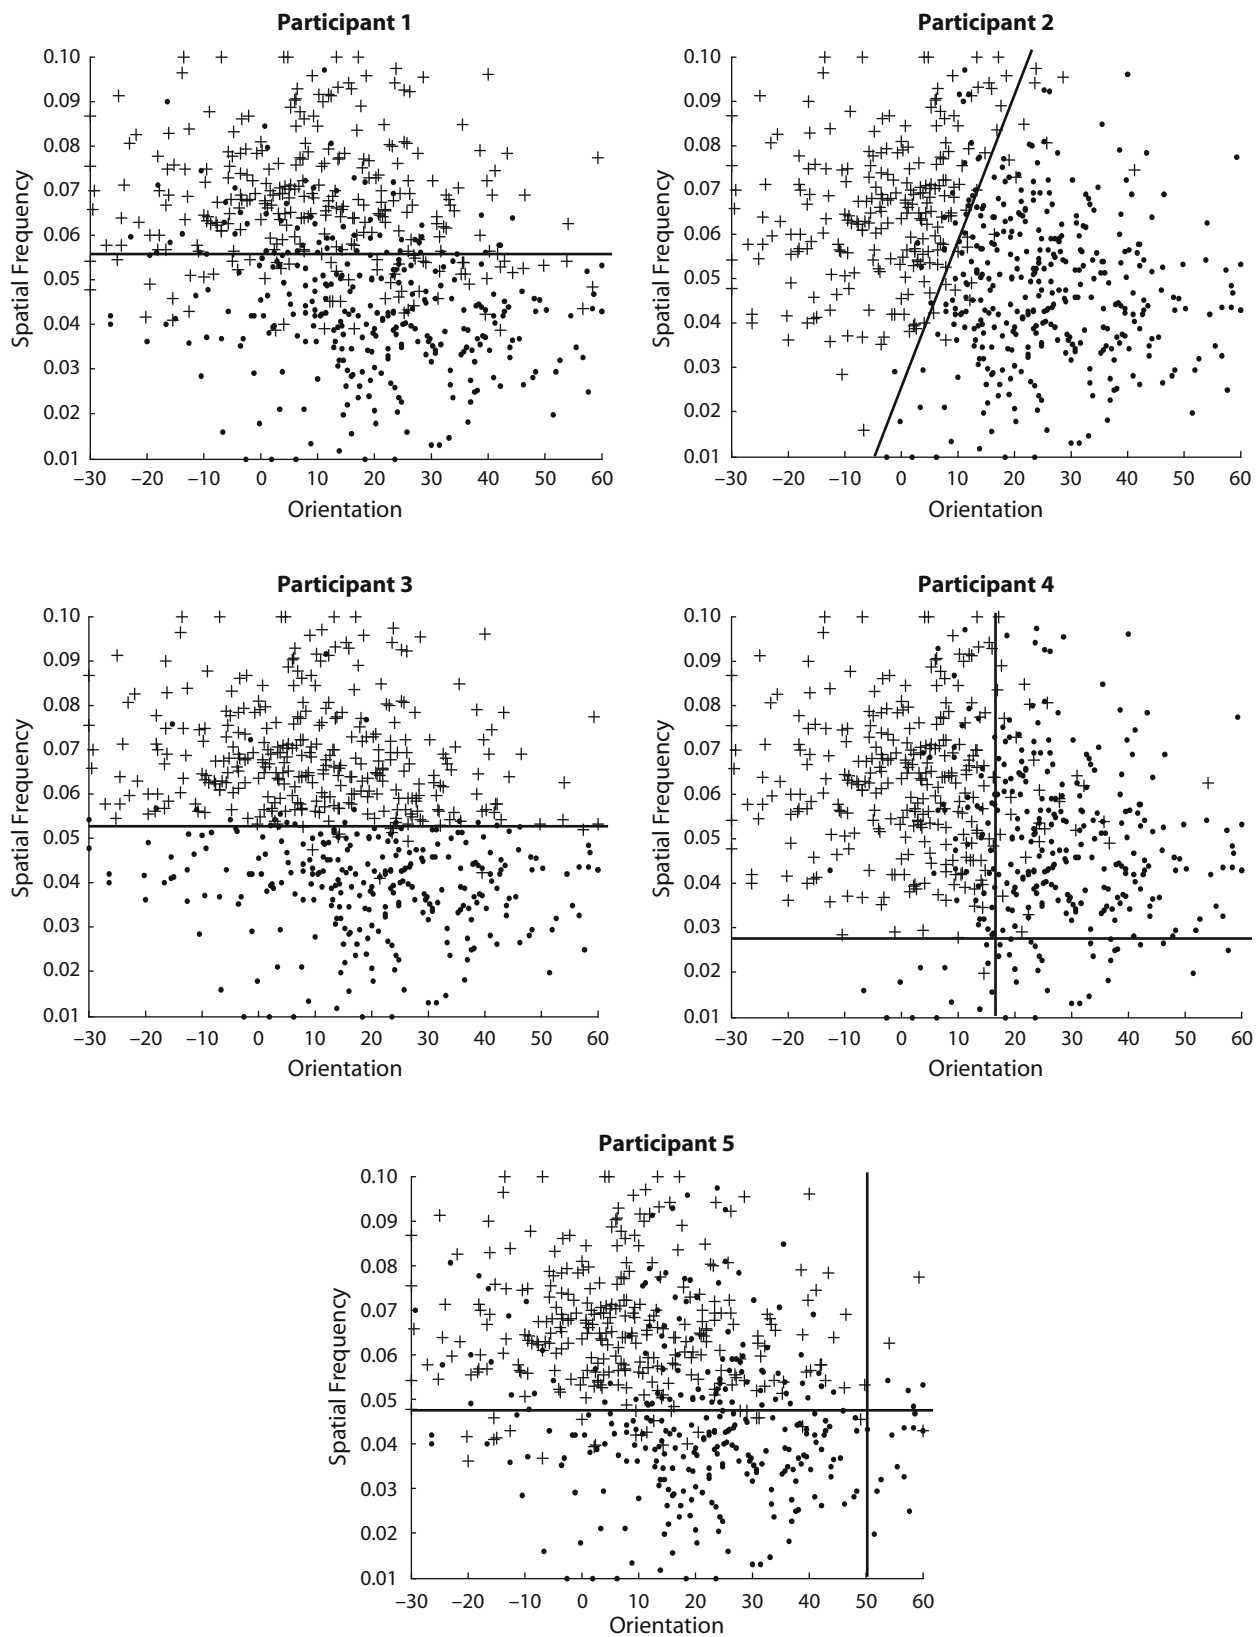


Figure 4. Response plots for the 5 participants in the positive feedback condition. Also shown are the decision bounds of the best-fitting models.

**Table 1**  
**Bayesian Information Criterion (BIC) Scores for All Participants and the**  
**Percentage of Responses Accounted for by the Best-Fitting Bounds**

Condition	Participant	GLC	GCC	DIM-W	DIM-O	% Accounted for
Positive feedback	1	533	528	806	<b>528</b>	79.0
	2	<b>222</b>	313	323	675	93.0
	3	210	216	799	<b>208</b>	95.2
	4	409	<b>408</b>	437	749	86.7
	5	575	<b>542</b>	759	602	79.0
Negative feedback	6	543	<b>535</b>	816	537	81.0
	7	<b>615</b>	634	645	759	76.5
	8	417	428	<b>412</b>	790	87.5
	9	328	<b>319</b>	545	567	87.8
	10	548	<b>544</b>	652	647	77.3
Control (partial)	11	<b>284</b>	348	558	512	90.3
	12	<b>448</b>	488	687	522	85.8
	13	472	<b>451</b>	782	472	84.1
	14	<b>233</b>	340	524	498	92.3
	15	<b>291</b>	351	623	397	90.5
Control (full, 3,000 trials)	16	<b>303</b>	398	614	441	89.5
	17	<b>287</b>	360	632	403	90.5
	18	<b>437</b>	473	720	488	84.7
	19	<b>343</b>	417	428	684	88.8
Control (full, 750 trials)	16	<b>68</b>	85	176	83	92.7
	17	<b>78</b>	99	170	90	91.3
	18	<b>113</b>	125	193	116	85.3
	19	<b>109</b>	134	141	159	86.7

Note—GLC, general linear classifier; GCC, general conjunctive classifier; DIM-W, one-dimensional (width) classifier; DIM-O, one-dimensional (orientation) classifier. Numbers in boldface indicate the best-fitting BIC for each participant.

received by participants was controlled. Participants who received feedback about both correct and incorrect responses (i.e., those in the control [full] and control [partial] groups) were more accurate than participants in groups that received feedback only about correct responses (PFB) or only about errors (NFB). Additionally, the model-based analyses revealed significant qualitative differences in the types of decision strategies used by participants in these conditions. In particular, participants in the PFB and NFB groups generally used explicit, rule-based strategies, despite categorizing 3,000 separate stimuli each and despite receiving feedback on 750 trials. In contrast, participants in all control groups used II strategies of the same type as the optimal decision bound. More subtly, one PFB participant showed II learning, but the evidence for any II learning among the NFB participants was weak. One NFB data set was best fit by an II model, but the fit was poor, suggesting that this participant made many responses that were inconsistent with an II strategy.

Overall, these results support the hypothesis that consistent II learning requires full feedback and that negative-only feedback offers no advantage over positive-only feedback in II tasks. Of course, more research is needed to determine the generalizability of these conclusions. For example, in the present experiment, the best one-dimensional rule could achieve an accuracy of 77.8%, only 8.2% below the accuracy of the best II strategy of 86%. This was a great enough difference to cause participants receiving feedback about both correct and incorrect responses to favor an II strategy over a rule-based strategy, but our results do not rule out the pos-

sibility that had this difference been greater, participants in the PFB and NFB groups might have been more likely to use II strategies.

The conclusions that consistent II learning requires full feedback and that negative-only feedback has no advantage over positive-only feedback in II tasks are in sharp contrast with the results that have been reported for explicit, rule-based learning. Many studies have reported that consistent rule-based learning is possible with either positive-only or negative-only feedback, and that in two-category tasks like the one used here, negative-only feedback is more effective than positive-only feedback (see, e.g., Buchwald, 1962; Jones, 1961; Meyer & Offenbach, 1962). These differences between the role of positive and negative feedback in rule-based and II category-learning tasks add to the growing list of empirical dissociations that have been reported between these two types of tasks (see, e.g., Ashby et al., 1999, 2002, 2003; Maddox et al., 2003; Maddox, Ashby, et al., 2004; Waldron & Ashby, 2001; Zeithamova & Maddox, 2006). Together, this evidence suggests that learning in rule-based and II category-learning tasks is mediated by functionally separate systems.

### Implications for Models of II Category Learning

Early models of II category learning were all purely cognitive. Most of these were derived from prototype, exemplar, or decision bound theories. None of these models incorporate enough structure to make specific predictions about the present experiment. For example, exemplar models assume that participants compare the stimulus with the memory representation of every previously seen



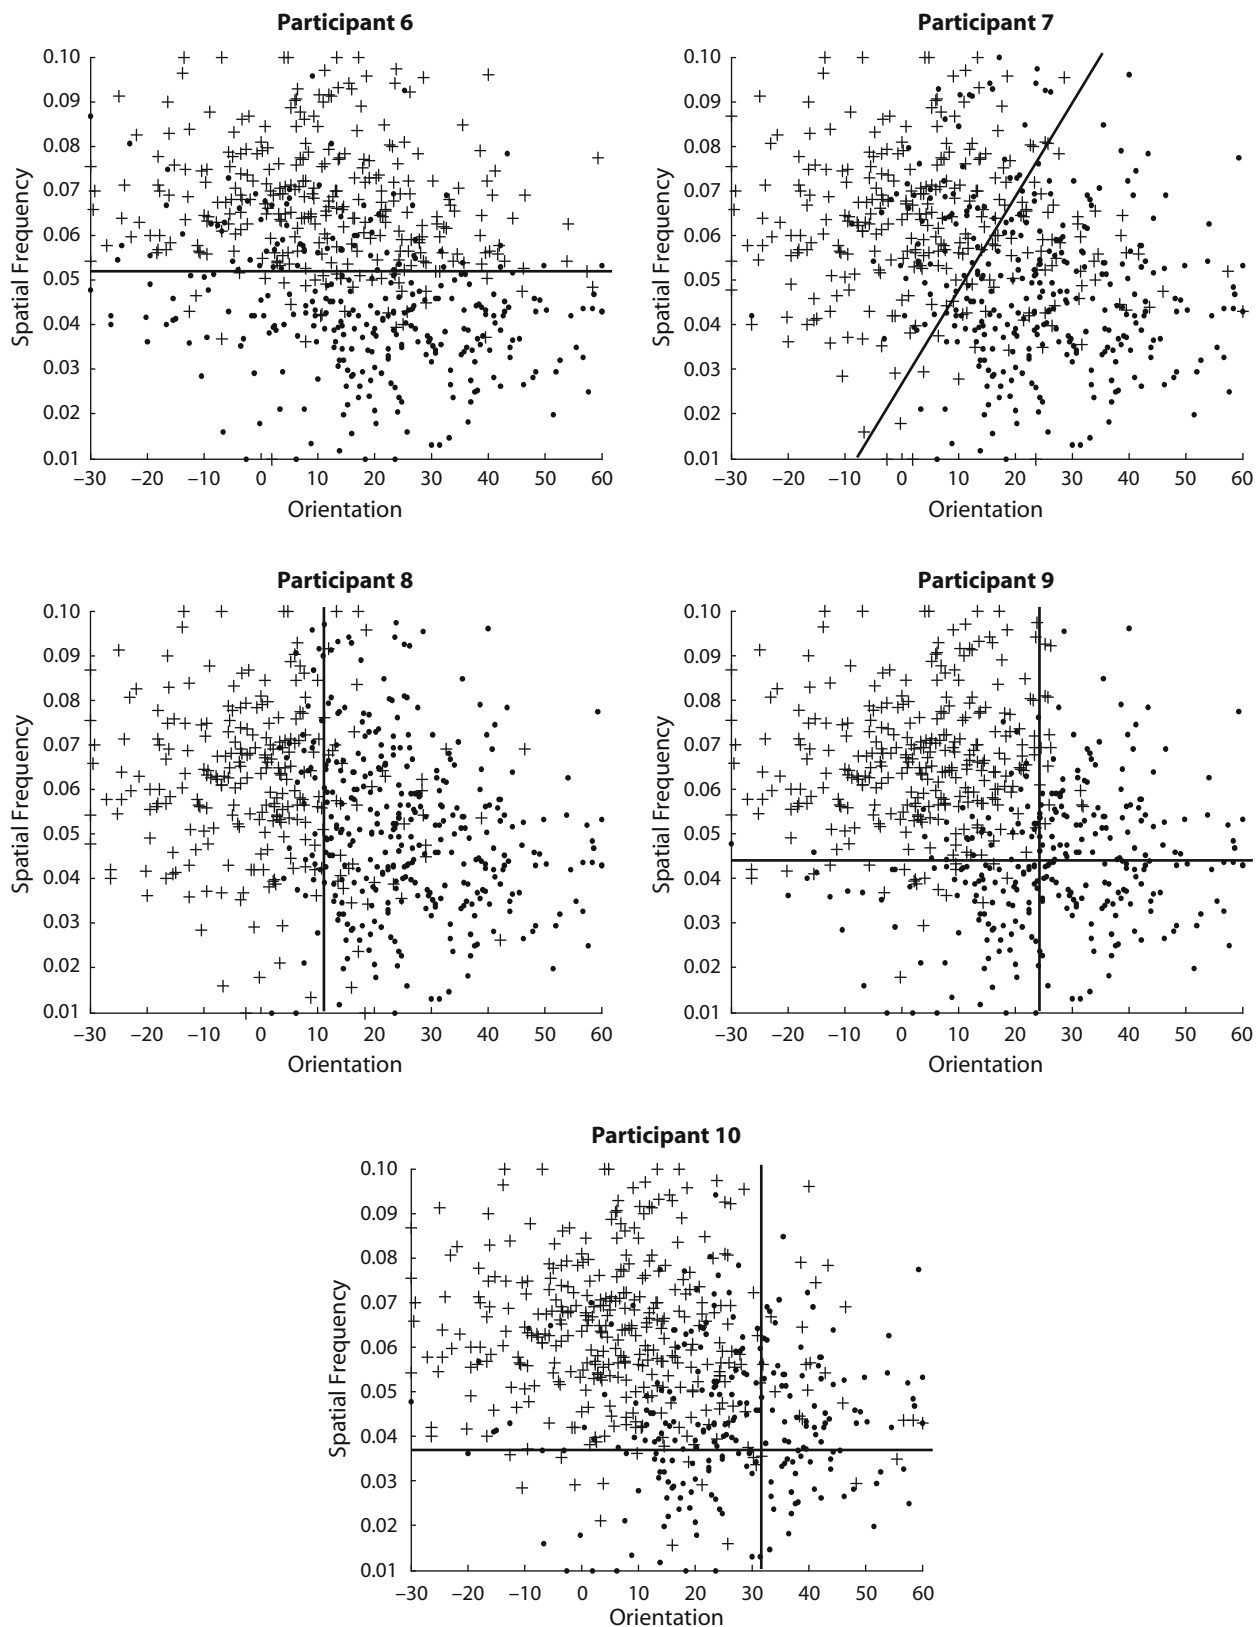


Figure 5. Response plots for the 5 participants in the negative feedback condition. Also shown are the decision bounds of the best-fitting models.

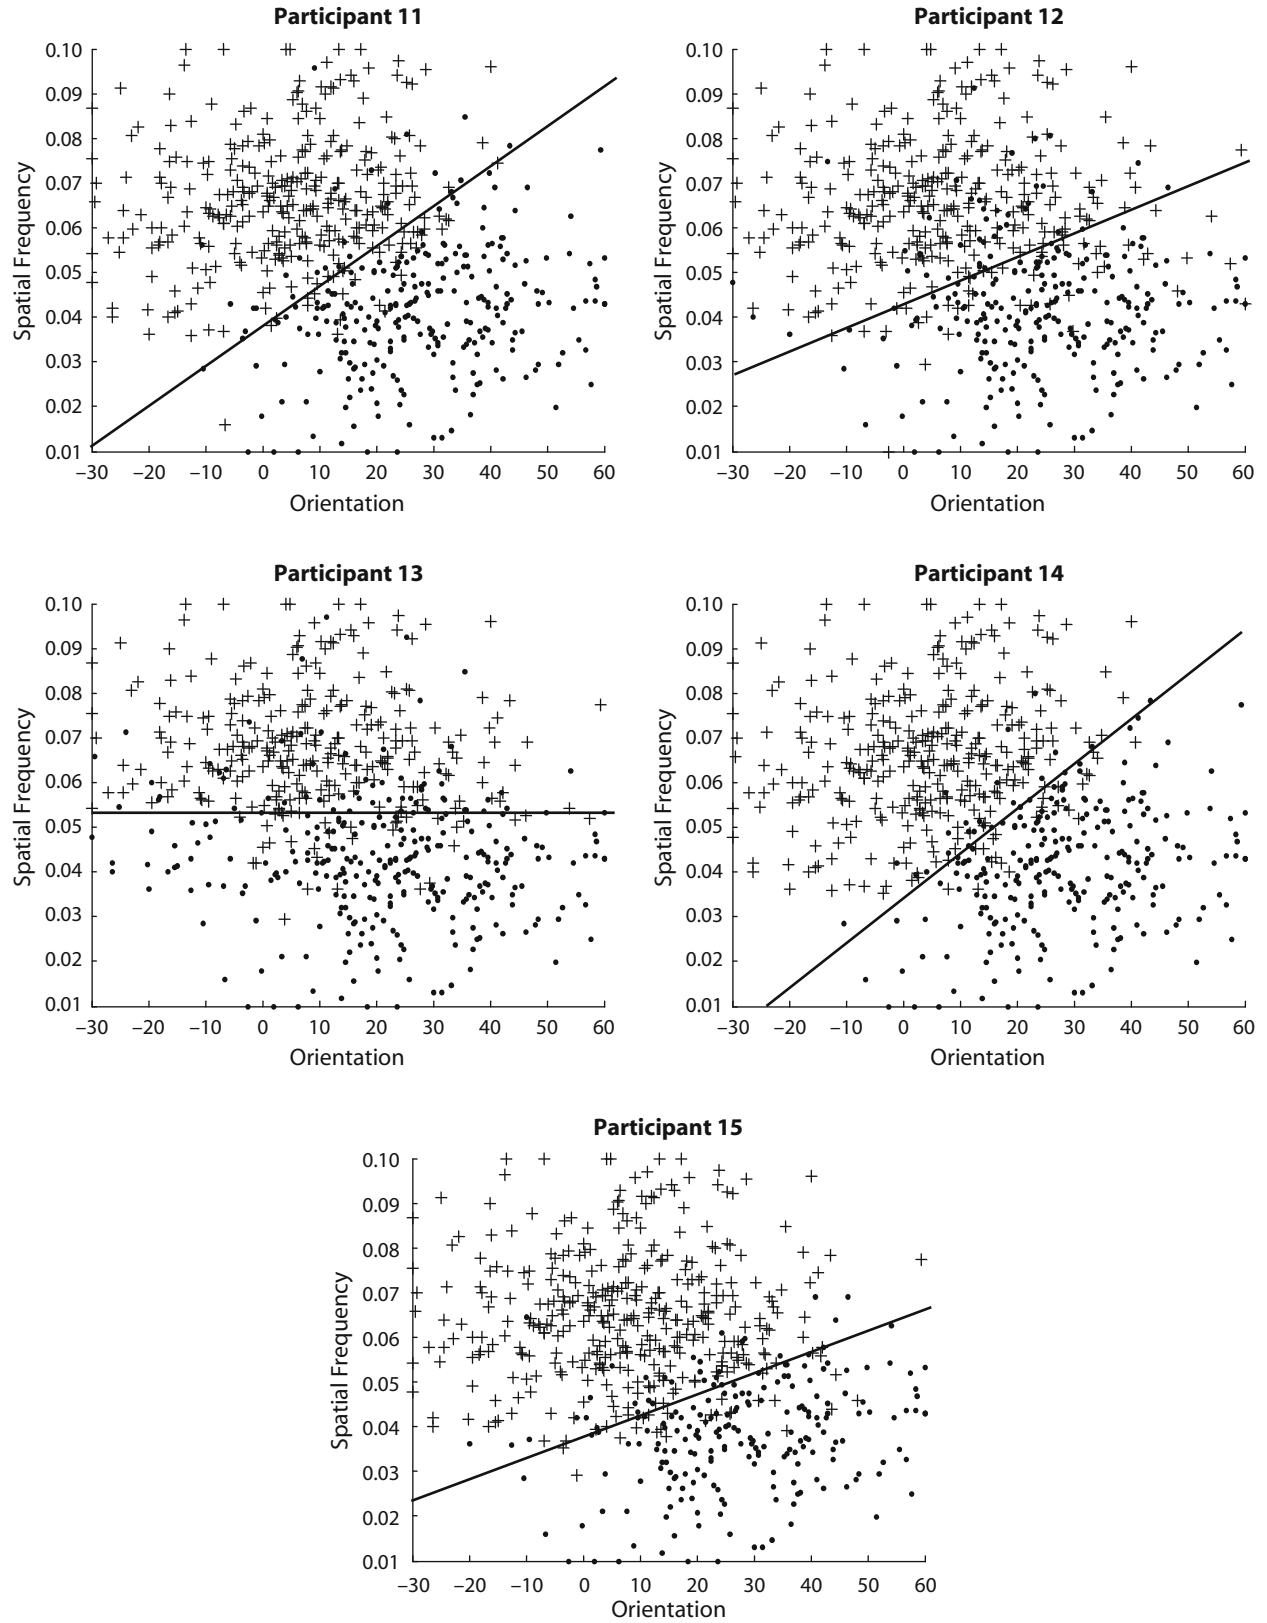
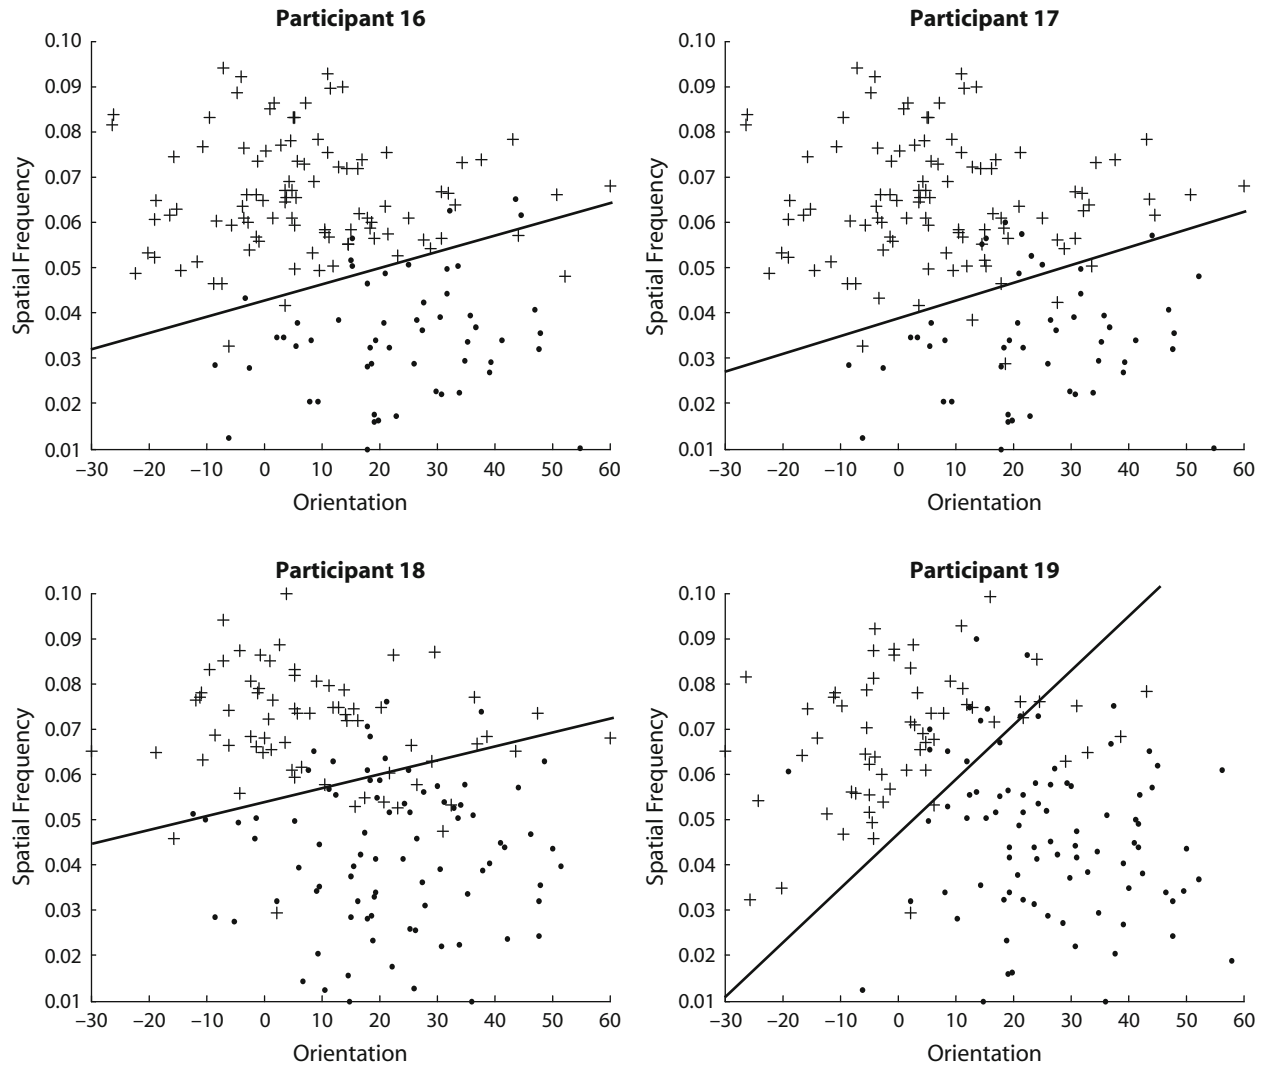


Figure 6. Response plots for the 5 participants in the control (partial) condition. Also shown are the decision bounds of the best-fitting models.



**Figure 7.** Response plots for the 4 participants in the control (full, first 750 trials) condition. Also shown are the decision bounds of the best-fitting models.

exemplar, but exemplar models do not specify how these memory representations might be differently updated following correct or error feedback.

More recently, there has been interest in the cognitive neuroscience of II category learning. Models developed at this level have more structure and are therefore more likely to make specific assumptions about the trial-by-trial effects of positive and negative feedback. Perhaps the most widely tested model of this type is COVIS (Ashby et al., 1998; Ashby & Waldron, 1999). The key structure in the COVIS procedural-learning system is the striatum—the major input structure within the basal ganglia. There is now much evidence from a wide variety of sources (including animal lesion studies, single-unit recording studies, neuroimaging experiments, studies with special neuropsychological patient groups, and traditional cognitive/behavioral experiments) that this is a critical site of II category learning (for a review, see Ashby & Ennis, 2006).

The striatum receives convergent projections from all cortical visual association areas, and then projects via the globus pallidus and thalamus to frontal cortex. COVIS assumes that trial-by-trial feedback facilitates the strengthening of cortical–striatal synapses following correct responses and the weakening of these synapses following errors. The neurotransmitter dopamine is assumed to play a critical role in changing these synaptic strengths. Many studies show that dopamine cells increase their firing levels above baseline following unexpected rewards and decrease their firing levels below baseline following the failure to receive an expected reward (Hollerman & Schultz, 1998; Schultz, 1998; Schultz, Apicella, & Ljungberg, 1993; Schultz, Dayan, & Montague, 1997). Another large body of literature suggests that dopamine has important effects on synaptic plasticity (see, e.g., Calabresi, Pisani, Centonze, Bernardi, 1996; Calabresi et al., 1997; Centonze, Picconi, Gubellini, Bernardi, & Calabresi, 2001; Nairn, Hemmings, Walaas, & Greengard, 1988; Nishi, Snyder, & Greengard, 1997; Wickens, 1993).

COVIS assumes that within the procedural-learning system, synapses are strengthened only if the following three conditions are present: (1) strong presynaptic activation, (2) strong postsynaptic activation (i.e., strong enough to activate NMDA receptors), and (3) an elevation in dopamine above baseline levels (which occurs following feedback that indicates a correct response). For trials on which the feedback indicates a correct response, COVIS therefore assumes that some synapses will be strengthened (i.e., those for which all three conditions are present), some synapses will be weakened (i.e., those at which the NMDA activation threshold is not met), and some will not change their synaptic strength (i.e., those at which there is no activation). In contrast, on error trials, COVIS assumes that synapses either will be weakened or will remain unchanged. Thus, there is an asymmetry in the COVIS procedural-learning system predictions. Learning is more flexible following positive versus negative feedback. As such, COVIS predicts that II learning should be better in the PFB condition than in the NFB condition.<sup>1</sup> It is important to note, however, that the present experiment was not designed to test this prediction. Clearly, a satisfactory test would require running many participants. The present data, however, argue strongly against the hypothesis that positive or negative feedback alone is enough to ensure consistent II learning.

The most significant result of the present experiment was that most participants in both restricted feedback conditions showed no evidence of II learning. Instead, these participants appeared to be using explicit, rule-based strategies. COVIS assumes that rule-based and procedural-learning systems each independently learn throughout training (at least under full feedback conditions). Which system controls the observable response depends on a system weight,  $\theta_E$ . Values of  $\theta_E$  greater than 0.5 favor the rule-based system, and values smaller than 0.5 favor the procedural system (see Ashby et al., 1998, for details). At the outset of training,  $\theta_E$  is assumed to be large (typically,  $\theta_E = 0.9$ ), but if feedback favors the procedural system, then  $\theta_E$  gradually decreases until eventually the procedural system begins controlling the response. According to COVIS, then, one possible account of the present results is that the procedural system learned in one or both of the NFB and PFB conditions, but that the incomplete feedback prevented the system weight  $\theta_E$  from decreasing below 0.5. According to this hypothesis, there was some procedural learning, but participants failed to use this knowledge when selecting a response. Further research is needed to test this hypothesis.

Although COVIS is the only current neurobiological model of II category learning, there are other relevant theoretical proposals about learning within the basal ganglia, including, for example, theories concerning the two primary pathways out of the striatum. The direct pathway projects from the striatum to the internal segment of the globus pallidus and from there to the thalamus and then to frontal cortex. The indirect pathway projects from the striatum to the external segment of the globus pallidus, then to the subthalamic nucleus, and then to the internal segment of the globus pallidus (where it rejoins the di-

rect pathway). Current versions of COVIS assign no role to the indirect pathway. Recently, however, Frank (2005) proposed that the indirect pathway mediates NO-GO learning, whereas the direct pathway mediates GO learning. Frank's model of how learning takes place on the direct pathway closely resembles COVIS; however, unlike COVIS, Frank's model predicts that error feedback teaches participants which response *not* to make by strengthening synapses on the indirect pathway. Frank et al. (2004) presented some striking data collected from Parkinson's disease patients that support this model of learning.

Frank's (2005) model was not developed for category learning,<sup>2</sup> so it makes no specific predictions regarding the present experiment. Frank's (2005) hypothesis about the role of the indirect pathway in learning could be incorporated into the procedural-learning system of COVIS.<sup>3</sup> This would make COVIS predict more successful learning under partial (i.e., negative-only) feedback conditions, and so would appear to make the model less rather than more compatible with the present results.

In summary, there appear to be two possible theoretical interpretations of our results. One possibility is that there was II learning in either or both of the PFB and NFB groups, but this learning was masked by an overreliance on explicit rule use. This hypothesis is compatible with the original version of COVIS if the learning was restricted to the PFB condition; it is compatible with a version of COVIS that incorporates Frank's (2005) hypothesis about the indirect pathway if there was II learning in both the PFB and NFB groups. A second possibility, which is inconsistent with either version of COVIS, is that there was no consistent II learning in either restricted feedback group.

## Conclusion

The present results suggest that consistent II category learning requires feedback about both correct and incorrect responses. One of five participants receiving positive-only feedback did show evidence of II learning, suggesting that II learning may be possible for some participants when only positive feedback is provided. In contrast, II models failed to provide a good account of the results of any of the participants who received negative-only feedback.

The present results add to the growing literature showing that II learning is extremely sensitive to the nature and timing of feedback. For example, II learning is impaired if the feedback is delayed by as little as 2.5 sec (Maddox et al., 2003); it is also impaired during observational training in which the category label precedes stimulus presentation (Ashby et al., 2002). Thus, effective II learning seems to require full feedback that is given immediately after the response. In contrast, rule-based category learning is unaffected by feedback delays as long as 10 sec (Maddox et al., 2003), and observational training is nearly as effective as feedback training with rule-based categories (Ashby et al., 2002). In addition, several studies have reported that rule-based category learning is possible with either positive-only or negative-only feedback, although with two categories, negative-only feedback is more efficacious than positive-only feedback (see, e.g., Buchwald, 1962; Jones, 1961; Meyer & Offenhach, 1962). These results suggest

that rule-based learning is less sensitive than II learning to the nature and timing of feedback. On the other hand, II learning does have some advantages over rule-based learning. For example, attention and effort are needed to process the feedback signal in rule-based category learning, whereas feedback processing during II learning requires neither attention nor effort but instead is essentially automatic (Maddox, Ashby, et al., 2004).

Expertise with many real-world categories appears to require some II learning. The present results suggest, for example, that when training a radiologist to identify tumors in an x-ray (a task that likely requires some II learning), it may be necessary to provide both positive and negative feedback.

#### AUTHOR NOTE

This research was supported in part by Public Health Service Grant MH3760-2. The authors are grateful to Amanda Price and two anonymous reviewers for their helpful comments. Correspondence concerning this article should be addressed to F. G. Ashby, Department of Psychology, University of California, Santa Barbara, CA 93106 (e-mail: ashby@psych.ucsb.edu).

#### REFERENCES

- ASHBY, F. G. (1992). Multidimensional models of categorization. In F. G. Ashby (Ed.), *Multidimensional models of perception and cognition* (pp. 449-483). 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 category learning. *Trends in Cognitive Sciences*, **5**, 204-210.
- ASHBY, F. G., ELL, S. W., & WALDRON, E. M. (2003). Procedural learning in perceptual categorization. *Memory & Cognition*, **31**, 1114-1125.
- ASHBY, F. G., & ENNIS, J. M. (2006). The role of the basal ganglia in category learning. *The Psychology of Learning & Motivation*, **46**, 1-36.
- 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. (1990). Integrating information from separable psychological dimensions. *Journal of Experimental Psychology: Human Perception & Performance*, **16**, 598-612.
- 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., MADDOX, W. T., & BOHIL, C. J. (2002). Observational versus feedback training in rule-based and information-integration category learning. *Memory & Cognition*, **30**, 666-677.
- ASHBY, F. G., NOBLE, S., FILOTEO, J. V., WALDRON, E. M., & ELL, S. W. (2003). Category learning deficits in Parkinson's disease. *Neuropsychology*, **17**, 115-124.
- ASHBY, F. G., QUELLER, S., & BERRETTY, P. M. (1999). On the dominance of unidimensional rules in unsupervised categorization. *Perception & Psychophysics*, **61**, 1178-1199.
- ASHBY, F. G., & WALDRON, E. M. (1999). On the nature of implicit categorization. *Psychonomic Bulletin & Review*, **6**, 363-378.
- BUCHWALD, A. M. (1962). Variations in the apparent effects of "right" and "wrong" on subsequent behavior. *Journal of Verbal Learning & Verbal Behavior*, **1**, 71-78.
- BUSS, A. H., & BUSS, E. H. (1956). The effect of verbal reinforcement combinations on conceptual learning. *Journal of Experimental Psychology*, **52**, 283-287.
- BUSS, A. H., WEINER, M., & BUSS, E. H. (1954). Stimulus generalization as a function of verbal reinforcement combinations. *Journal of Experimental Psychology*, **48**, 433-436.
- CALABRESI, P., PISANI, A., CENTONZE, D., & BERNARDI, G. (1996). Role of Ca<sup>2+</sup> in striatal LTD and LTP. *Seminars in the Neurosciences*, **8**, 321-328.
- CALABRESI, P., SAIARDI, A., PISANI, A., BAIK, J., CENTONZE, D., MERCURI, N., ET AL. (1997). Abnormal synaptic plasticity in the striatum of mice lacking dopamine D2 receptors. *Journal of Neuroscience*, **17**, 4536-4544.
- CENTONZE, D., PICCONI, B., GUBELLINI, P., BERNARDI, G., & CALABRESI, P. (2001). Dopaminergic control of synaptic plasticity in the dorsal striatum. *European Journal of Neuroscience*, **13**, 1071-1077.
- ELL, S. W., & ASHBY, F. G. (2006). The effects of category overlap on information-integration and rule-based category learning. *Perception & Psychophysics*, **68**, 1013-1026.
- FILOTEO, J. V., MADDOX, W. T., & DAVIS, J. D. (2001). A possible role of the striatum in linear and nonlinear category learning: Evidence from patients with Huntington's disease. *Behavioral Neuroscience*, **115**, 786-798.
- FILOTEO, J. V., MADDOX, W. T., SALMON, D. P., & SONG, D. D. (2005). Information-integration category learning in patients with striatal dysfunction. *Neuropsychology*, **19**, 212-222.
- FRANK, M. J. (2005). Dynamic dopamine modulation in the basal ganglia: A neurocomputational account of cognitive deficits in medicated and nonmedicated Parkinsonism. *Journal of Cognitive Neuroscience*, **17**, 51-72.
- FRANK, M. J., SEEBERGER, L. C., & O'REILLY, R. C. (2004). By carrot or by stick: Cognitive reinforcement learning in Parkinsonism. *Science*, **306**, 1940-1943.
- HEATON, R. K. (1981). *A manual for the Wisconsin Card Sorting Test*. Odessa, FL: Psychological Assessment Resources.
- HOGUE, M. A., & STOCKING, R. (1912). A note on the relative value of punishment and reward as motives. *Journal of Animal Behavior*, **2**, 43-50.
- HOLLERMAN, J., & SCHULTZ, W. (1998). Dopamine neurons report an error in the temporal prediction of reward during learning. *Nature Neuroscience*, **1**, 304-309.
- JONES, A. (1961). The relative effectiveness of positive and negative verbal reinforcers. *Journal of Experimental Psychology*, **62**, 368-371.
- 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. *Behavioral Processes*, **66**, 309-332.
- MADDOX, W. T., ASHBY, F. G., & BOHIL, C. J. (2003). Delayed feedback effects on rule-based and information-integration category learning. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **29**, 650-662.
- MADDOX, W. T., ASHBY, F. G., ING, A. D., & PICKERING, A. D. (2004). Disrupting feedback processing interferes with rule-based but not information-integration category learning. *Memory & Cognition*, **32**, 582-591.
- MADDOX, W. T., BOHIL, C. J., & ING, A. D. (2004). Evidence for a procedural-learning-based system in category learning. *Psychonomic Bulletin & Review*, **11**, 945-952.
- MADDOX, W. T., FILOTEO, J. V., HEIL, K. D., ING, A. D. (2004). Category number impacts rule-based but not information-integration category learning: Further evidence for dissociable category learning systems. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **30**, 227-235.
- 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, & Cognition*, **31**, 100-107.
- MEYER, W. J., & OFFENBACH, S. I. (1962). Effectiveness of reward and punishment as a function of task complexity. *Journal of Comparative & Physiological Psychology*, **55**, 532-534.
- NAIRN, A. C., HEMMINGS, H. C., JR., WALAAS, S. I., & GREENGARD, P. (1988). DARPP-32 and phosphatase inhibitor-1, two structurally related inhibitors of protein phosphatase-1, are both present in striatonigral neurons. *Journal of Neurochemistry*, **50**, 257-262.
- NISHI, A., SNYDER, G., & GREENGARD, P. (1997). Bidirectional regulation of DARPP-32 phosphorylation by dopamine. *Journal of Neuroscience*, **17**, 8147-8155.
- NOMURA, E. M., MADDOX, W. T., FILOTEO, J. V., ING, A. D., GITELMAN,

- D. R., PARRISH, T. B., ET AL. (2007). Neural correlates of rule-based and information-integration visual category learning. *Cerebral Cortex*, *17*, 37-43.
- SCHULTZ, W. (1998). Predictive reward signal of dopamine neurons. *Journal of Neurophysiology*, *80*, 1-27.
- SCHULTZ, W., APICELLA, P., & LJUNGBERG, T. (1993). Responses of monkey dopamine neurons to reward and conditioned stimuli during successive steps of learning a delayed response task. *Journal of Neuroscience*, *13*, 900-913.
- SCHULTZ, W., DAYAN, P., & MONTAGUE, P. R. (1997). A neural substrate of prediction and reward. *Science*, *275*, 1593.
- SCHWARZ, G. (1978). Estimating the dimension of a model. *Annals of Statistics*, *6*, 461-464.
- SEGER, C. A., & CINCOTTA, C. M. (2002). Striatal activity in concept learning. *Cognitive, Affective, & Behavioral Neuroscience*, *2*, 149-161.
- SMITH, E. E., PATALANO, A. L., & JONIDES, J. J. (1998). Alternative strategies of categorization. *Cognition*, *65*, 167-196.
- SMITH, J. D., TRACY, J. I., & MURRAY, M. J. (1993). Depression and category learning. *Journal of Experimental Psychology: General*, *122*, 331-346.
- WALDRON, E. M. (2000). *Feedback processing in category learning*. Unpublished manuscript, University of California at Santa Barbara.
- 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.
- WARDEN, C. J., & AYLESWORTH, M. (1926). The relative value of reward and punishment in the formation of a visual discrimination habit in the white rat. *Journal of Comparative Psychology*, *7*, 117-127.
- WICKENS, J. (1993). *A theory of the striatum*. New York, Pergamon Press.
- ZEITHAMOVA, D., & MADDOX, W. T. (2006). Dual-task interference in perceptual category learning. *Memory & Cognition*, *34*, 387-398.

## NOTES

1. We conducted informal simulations of the procedural-learning system of COVIS to verify this prediction. As expected, this COVIS subsystem learned the Figure 2 category structures with positive-only feedback under many different parameter settings, whereas conditions in which the model learned with negative-only feedback were not found (although an exhaustive search through all possible parameter combinations was not conducted). Thus, the COVIS procedural-learning system does not account for the difficulties experienced by the PFB group.

2. Frank et al. (2004) argued that their model applies to procedural-learning tasks, thus at first glance it seems that deriving predictions from their model in II category-learning tasks should be straightforward. However, such predictions are inappropriate for two reasons. First, Frank et al. (2004) defined procedural learning as trial-and-error learning (p. 1941). This definition is considerably more general than ours and presumably would include, for example, both II and rule-based learning. Second, there is now overwhelming evidence that human category learning is mediated by multiple systems (e.g., see Ashby & Maddox, 2005), so any single-system model, including the model proposed by Frank (2005), will have difficulty accounting for these recent, significant category-learning results. Since Frank (2005) did not propose a general model of category learning, it seems inappropriate to draw strong conclusions about what his model would predict for the present experiment.

3. Computationally, there are a number of ways to add a Frank-like indirect pathway to the procedural-learning system of COVIS. For example, different versions of such a model could be constructed depending on exactly how neurons in the external segment of the globus pallidus and the subthalamic nucleus were modeled and whether Hebbian or three-factor learning was assumed (i.e., neither structure exists in current versions of COVIS). For this reason, we did not attempt to simulate a version of COVIS that included the indirect pathway.

## APPENDIX

Three different rule-based decision bound models were fit to each participant's data (two unidimensional models and the conjunctive model), along with one II model (general linear classifier). For more details, see Ashby (1992) or Maddox and Ashby (1993). We also fit two models that assumed responses were completely random: an equal-response frequency model that assumed participants were equally likely to respond A or B to every stimulus and a biased-response frequency model that assumed participants responded A with some fixed probability for all stimuli (i.e., unequal to one half). The equal-response frequency model and the biased-response frequency model have 0 parameters and 1 parameter, respectively. Since no participant's data were best fit by either random response model, they were not included in discussion of the results.

## Rule-Based Models

**The one-dimensional classifier (DIM).** The DIM assumes that participants set a decision criterion on a single stimulus dimension. For example, a participant might base his or her categorization decision on the following rule: "Respond 'A' if the bar width is small, otherwise respond 'B.'" Two versions of the DIM were fit to the data. DIM (width) assumed a decision on the basis of bar width, and DIM (orient) assumed a decision on the basis of orientation. These models have two parameters: a decision criterion along the relevant perceptual dimension and a perceptual noise variance.

**The general conjunctive classifier (GCC).** The GCC assumes that the rule used by participants is a conjunction of the type: "Respond 'A' if the bar width is small and the orientation is  $> 45^\circ$ ; otherwise, respond 'B.'" Although several different versions of the GCC could be fit to the present data, only the version that seemed plausible on the basis of a visual inspection of the response data was fit. The GCC has three parameters: one each for the single-decision criterion placed along each stimulus dimension (i.e., orientation and width) and a perceptual noise variance.

## Information Integration Model

**The general linear classifier (GLC).** The GLC assumes that participants divide the stimulus space using a linear decision bound. Categorization decisions are then based upon which region each stimulus is perceived to fall in. These decision bounds require linear integration of both stimulus dimensions, thereby producing an II decision strategy. The GLC has three parameters: the slope and intercept of the linear decision bound and a perceptual noise variance.