Category Learning Through Active Sampling

Doug Markant (doug.markant@nyu.edu)
Todd M. Gureckis (todd.gureckis@nyu.edu)
New York University, Department of Psychology
6 Washington Place, New York, NY 10003 USA

Abstract

Laboratory studies of human category learning tend to emphasize passive learning by limiting participants' control over the information they experience on every trial. In contrast, we explore the impact that *active* data selection has on category learning. In our experiment, participants attempted to learn categories under either entirely passive conditions, or by actively selecting and querying the labels associated with particular stimuli. We found that participants generally acquired categories faster in the active learning condition. Furthermore, this advantage depended on learners actually making decisions about which stimuli to query themselves. However, the effectiveness of active sampling was modulated by the particular structure of the target category. A probabilistic rule-learning model is proposed that explains the results in terms of a strong prior bias towards uni-dimensional rules which impairs learning of alternative category boundaries. Active learners appear to be able to bootstrap their own learning, but this ability may be strongly constrained by the space of hypotheses that are under consideration. **Keywords:** categorization, active learning, information sampling, rule learning, decision-bound models

Despite the widely held view that people learn better by *doing* than simply *observing*, there have been surprisingly few detailed accounts of the impact that "active" information acquisition has on the learning process. In particular, theoretical models which explain how people learn new concepts from examples usually treat learners as passive accumulators of evidence about the structure of categories. For example, the standard procedure in most category learning experiments is to exhaustively and randomly sample the set of training stimuli. However, in everyday life, human learners can often control their own learning by selectively "sampling" particular observations they estimate to be useful or informative. The goal of the present paper is to understand the cognitive consequences of this type of learning.

There are at least two explanations for why active sampling might result in better learning than passive observation. First, rather than being limited by the flow of information from passive experience, active learners are free to select which information they want to learn about. For example, by making directed queries that take into account their current uncertainty, the learner may be able to optimize their experience (e.g., avoiding redundant data). Research in machine learning has shown that the principle of uncertainty sampling (selectively querying data that is expected to be informative) can have a dramatic impact on the amount of training needed to reach a performance criterion (Settles, 2009).

Independent of the advantage of better data, active learners may also benefit from greater engagement in the learning task. For example, the very act of planning interventions or deciding which samples to take may necessitate deeper evaluation of the problem structure and of how observed experience relates to different hypotheses (c.f., Bruner, 1961). In a study of active intervention during a causal learning task, Sobel and Kushnir (2006) showed that active learners were more likely to learn a hidden causal structure than participants that were "yoked" to their interventions (i.e., a group with the same data but who did not independently make sampling decisions). Similar concerns are often used to support educational practices that emphasize "inquiry" or "discovery"-based instruction (Kuhn et al., 2000).

The aims of the present study were two-fold. First, we were interested if participants could adaptively structure their own learning experiences when acquiring new concepts. Second, we were interested in how the effectiveness of active sampling might interact with the specific structure of categories. While a number of recent studies have explored how learners make information sampling decisions to support their own learning (Castro et al., 2008; Kruschke, 2008; Gureckis & Markant, 2009; Steyvers et al., 2003), there has not yet been a systematic evaluation of how this ability might vary across different category structures.

Overview of the present experiment

Our experiment adapts a well-studied paradigm for perceptual category learning using multidimensional, continuous-valued stimuli. In the task, participants learned to classify perceptual stimuli into different abstract groups. Two types of category structures were used: *rule-based* (RB), in which the decision rule is defined as a criterion along a single dimension, and (2) *information-integration* (II), in which the decision rule is a function of at least two dimensions (see Figure 1). Participants in the experiment were further divided into three training conditions. In the *passive-normal* condition, participants observed training stimuli that were generated from two bivariate normal distributions (i.e., a standard training procedure). In the *active* condition, participants were able to "design" stimuli for which they received feedback about the category label. In the *passive-yoked* condition, each participant was linked to an active learner, passively observing the samples they made and receiving the same feedback.

There are three key aspects of the design worth highlighting. First, in binary classification tasks, the optimal sampling strategy is simply to make queries close to the current estimate of the category boundary (or margin) — the region of greatest uncertainty. However, we anticipated that participants' ability to do so might vary between the RB and II learning tasks. Previous research has suggested that these two types of category structures may be learned in fundamentally different ways (Ashby, Alfonso-Reese, Turken, & Waldron,
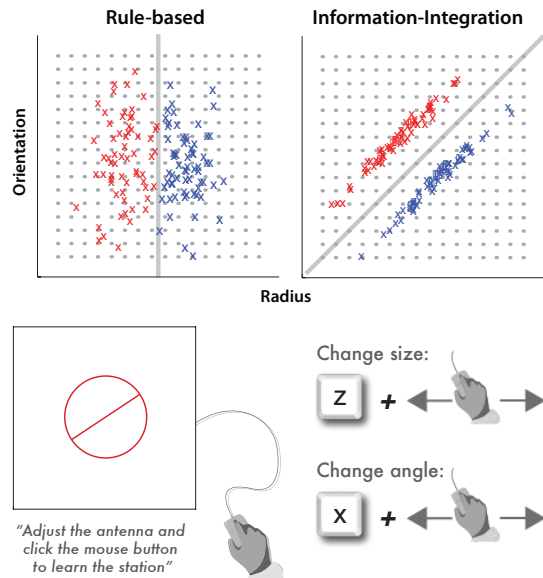
Figure 1: **Top:** Category distributions used in the experiment. 'X's indicate training stimuli shown to participants in the passive-normal condition with color indicating the generating distribution (actual feedback received by participants was probabilistic). The uniform grid of points over the stimulus space indicate the set of unlabeled test stimuli. **Bottom:** An example stimulus (left) and the interface used in the active learning condition.

1998). In particular, RB categories are thought to be learned by reasoning about verbal or explicit hypotheses (which is the default learning mode), while the structure of II categories precludes a simple verbal description and are instead thought to be learned via implicit or procedural learning. To the degree that effective sampling relies on explicit reasoning about uncertainty, people may perform better in the RB condition where this uncertainty may be better articulated. Similarly, active learning may be more effective in the RB case because the category aligns with default biases people bring to the task (Ashby et al., 1999; Kruschke, 1993).

Second, the comparison of active learners with the passive-normal group allowed us to test if active learning could lead to a performance advantage above and beyond the typical training procedure in such tasks. We expected that if active learners were able to make useful queries, they would be faster at learning the correct category distinction than the passive-normal participants. Again, if active learners are less successful at making useful queries in the II task, any learning advantage may be attenuated. Moreover, since successful learning in the II task may be contingent on abandoning rule-based strategies in favor of a more procedural type of learning, active learning might even lead to a learning *impairment* by encouraging perseveration in the search for a sub-optimal rule.

Finally, the inclusion of the passive-yoked training group allowed us to separately evaluate the impact of selecting samples from the statistical information contained in those samples (since the distribution of training data is identical for both groups). While previous research (in causal learning settings)

suggests that active or intervention-driven learning may lead to advantages over comparable yoked conditions (Lagnado & Sloman, 2004; Sobel & Kushnir, 2006; Steyvers et al., 2003), it is unknown how these results generalize to other tasks.

An Experiment

Participants One hundred eighty undergraduates at New York University participated in the study. The experiment was run on standard Macintosh computers in a single 40 min session. Each participant was assigned to either the rule-based (RB) or information-integration (II) task condition, and to one of three training conditions: active (A), passive-normal (P), or passive-yoked (PY).

Stimuli Stimuli were defined by a two-dimensional continuous-valued feature space, where one dimension corresponded to the size (radius) of a circle and the second dimension corresponded to the angle of a central diameter (see example in Figure 1, bottom). One-hundred and twenty-eight training stimuli were created for the passive-normal training condition using bivariate normal distributions (see Figure 1, top) with mean and covariance parameters slightly modified from Ashby et al. (2002). Test stimuli were drawn from a uniform grid of samples over the feature space (depicted by the gray dots in Figure 1). Thirty-two stimuli were presented in each test block, amounting to a total of 256 test trials.

Procedure Participants were told that the stimuli in the experiment were "loop antennae" for old televisions, and that each antennae received one of two channels (CH1 or CH2). The channel received by any antenna depended in some way on the two dimensions described above, and participant's goal was to learn the difference between the two types of items. The feedback associated with each item during training was probabilistic and was proportional to the relative likelihood of either category for the ideal observer who knew the true category distributions. Participants were given instruction that the antennas were sometimes "noisy" and would pick up the wrong channel and that it would be beneficial to integrate over a number of trials when learning. The experiment consisted of 8 blocks, with each block divided into a set of 16 training trials followed by 32 (no feedback) test trials.

Training – Active Condition. On each training trial the participant "designed" a TV antenna and learned about its category membership. Each trial began with the presentation of a randomly generated stimulus in the center of the screen. The participant could then alter its size and orientation by moving the mouse from left to right while holding down either the 'Z' or 'X' key, respectively (see Figure 1, bottom). Only one dimension could be changed at a time, but participants could make any number of changes and use as much time as needed. When the stimulus was the desired size and orientation, participants pressed the mouse button to reveal the category label, which appeared above the stimulus and was visible for 1500ms. Querying the category label was not permitted until the participant had made a change to the initial stimulus.

Training Trials – Passive-Normal Condition. In the passive-normal condition, participants were unable to interact with the stimuli in any manner[1]. Instead, in each trial they were presented with a stimulus generated from the category distributions described above. On each trial, a fixation cross was presented, followed by the stimulus (for 250ms), followed by the category label (above the stimulus for 1500ms). When the category label was displayed, the participant was required to press a key corresponding to that category in order to end the trial. This procedure is equivalent to the observational learning condition used in Ashby et al. (2002).

Training – Passive-Yoked Condition. The purpose of the yoked

---

[1] In this design passive participants are not matched to active participants in terms of perceptual-motor task demands (e.g., precisely adjusting the stimulus). However, pilot data suggested that equating this made learning much more difficult for the passive group, potentially playing into any hypothesized active learning advantage.
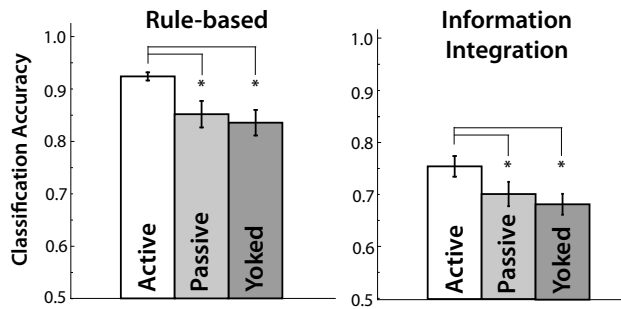
Figure 2: Accuracy in RB (left) and II (right) tasks for the three training conditions. Error bars show the standard error of the mean.

condition was to mimic the passive training experience, but to use a sequence of observations that were selected by a participant in the active condition. Each yoked participant was assigned to a matching participant in the active learning condition that had already completed the study. Training samples from the active participant were used as the set of training items for the yoked participant, and were presented in the identical order as they had been generated by the active participant. All other aspects of the yoked condition were identical to the passive-normal condition.

Test – All Conditions. On each test trial, a single item was presented in the center of the display and participants were asked to classify the item according to the channel the item was most likely to receive. A response was required to complete the trial, and participants responded at their own pace. No feedback was provided on individual test trials. At the end of each block participants were told their cumulative accuracy during the block they just completed, as well as their accuracy during the preceding test block.

Results

Responses during test blocks were scored according to whether the participant identified the correct category of each test item (with respect to the true discriminant function). Overall accuracy across tasks and conditions is shown in Figure 2. A 2-way ANOVA with task type (RB/II) and training condition (A/P/PY) as between subjects factors found significant main effects of both task ($F(1, 174) = 155.97, p < 0.001$) and training condition ($F(2, 174) = 15.34, p < 0.001$), but no interaction ($F(2, 174) = 0.27$). In the RB task, overall accuracy was significantly higher in the active condition than in both the passive-normal ($t(58) = 2.69, p < 0.01$) and passive-yoked ($t(58) = 3.96, p < 0.001$) conditions, while there was no difference between the two passive conditions. Similarly, in the II condition, the active group was more accurate than both passive groups (P: $t(58) = 2.58, p < 0.05$; PY: $t(58) = 4.27, p < 0.001$), while there was no difference between passive-normal and passive-yoked ($t(58) = 1.57, p = 0.12$). Note that while active learners generally outperformed their passive counter-parts, active samplers in the II task only achieved 75% correct on average which may reflect a variety of sub-optimal rule-based strategies.

For participants in the II task, a 2-way ANOVA on average accuracy revealed a main effect of condition ($F(2, 609) = 8.74, p < 0.001$), a main effect of block ($F(7, 609) = 3.92, p < 0.001$), and a significant condition-by-block interaction ($F(14, 609) = 1.74, p < 0.05$). Examination of this

interaction suggested that it was driven by an early learning advantage for the active learners which was reduced later in the task. A similar analysis in the RB condition found only a main effect of training condition ($F(2, 87) = 6.65, p < 0.005$) and block ($F(7, 609) = 17.31, p < 0.001$).

Sampling behavior. Figure 3A shows the distribution of queries for active participants in the RB and II tasks for the final training block. In both tasks, participants begin by widely distributing their samples over the stimulus space, but over time make samples that are closer to the true category boundary. We measured the orthogonal distance of each sample to the true category boundary and computed the average distance within each block. Figure 3B shows that in the RB task average distance was significantly smaller than the null hypothesis of a random sampling strategy by the second training block (one-sample t-test, $t(29) = 4.33, p < 0.001$). This shift toward margin sampling was slower and less extreme in the II task, with average distance reliably smaller than expected from a random strategy starting around the sixth training block ($t(29) = 4.53, p < 0.001$).

Relating sampling behavior and learning. We found that overall sample distance from the boundary (averaged across blocks) was significantly correlated with active learners' overall test performance in both the RB ($r = -0.42, p < 0.05$) and II ($r = -0.8, p < 0.001$) tasks (see Figure 3D, blue line). One question is if being yoked to a high-performing active participant leads to a similar learning advantage for the passive-yoked participants. In contrast to active learners, average sample distance was not strongly correlated with performance in either task condition (RB: $r = 0.36, p = 0.051$, II: $r = -0.05, p = 0.4$, see Figure 3D, orange line). In fact, there was even a trend toward the reverse relationship in the RB task; that is, passive-yoked learners who received the most objectively useful training data were among the worst performers in the group for that task.

One objection to measuring sample "quality" by its distance from the true category boundary is that people may instead evaluate samples relative to their subjective belief about the boundary at any point in time. Using logistic regression we found the best-fit linear decision boundary for subjects' response data on each test block. We then computed the average "subjective" sample distance from that boundary in the following training block, and computed the average over blocks for all active and passive-yoked participants. We found that this distance was smaller in the active group than passive-yoked group in both tasks highlighting the divergence in inference between the two groups (RB: $t(29) = -4.07, p < 0.001$, II: $t(28) = -4.94, p < 0.001$). In addition, subjective distance measure was negatively correlated with overall accuracy in all conditions (RB(A): $r = -0.54, p < 0.005$, RB(PY): $r = -0.47, p < 0.05$, II(A): $r = -0.79, p < 0.001$, II(PY): $r = -0.41, p < 0.05$, see Figure 3E).

Discussion

There are three key behavioral findings from the experiment. First, active learners were more accurate than passive ob-
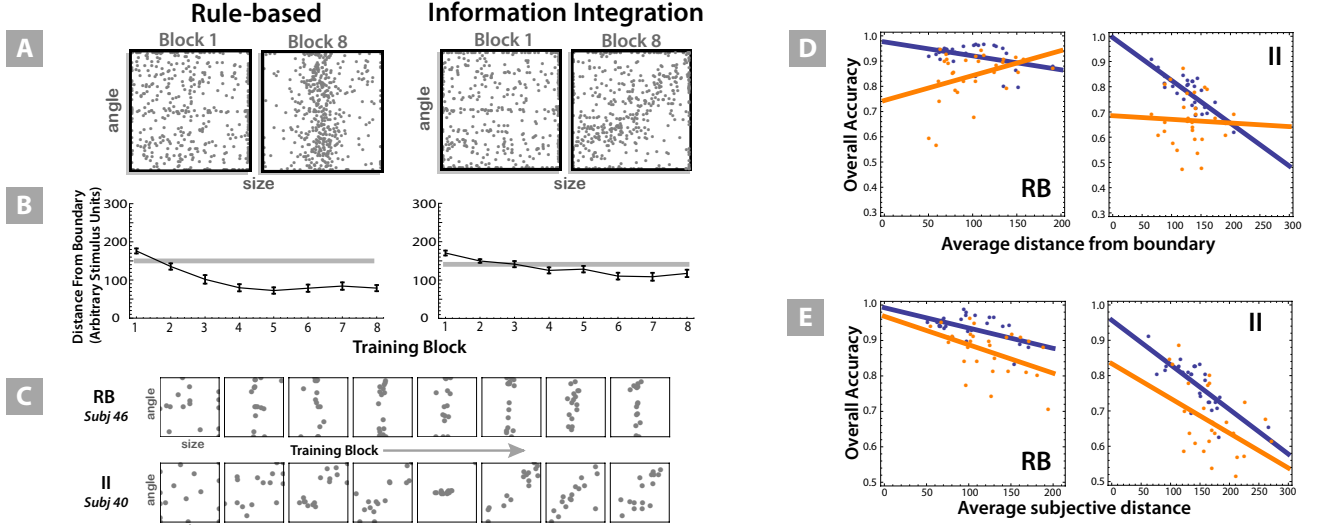
Figure 3: **A**: Composite of samples chosen by active participants in the first (left) and last (right) training block. **B**: Average distance of participants' samples from the true category boundary (black) as compared to average distance expected from random sampling (gray line). **C**: Examples of active participants in both tasks that successfully sample close to the true category boundary. **D**: Samples closer to the true boundary are associated with higher accuracy in active but not passive-yoked learners, while low "subjective" distance from a best-fit response boundary is predictive of higher accuracy in both groups (**E**).

servers in both tasks. One explanation is that active learners are able to query regions in the stimulus space where they are most likely to commit classification errors (i.e., the margin of the category boundary). Since participants in the passive-normal condition received samples from a "true" category distribution, they may be at a disadvantage because they were less likely to observe test items close to the boundary. Nevertheless, it is extremely interesting that näive participants could intuitively identify what information would most useful to support their own learning in an abstract problem space.

However, the advantage for active learners cannot be explained by a difference in training data alone. Most striking is the finding that yoked participants showed no improvement over the passive-normal group despite learning from the exact same observations as the active group. Indeed, the passive-yoked participants that observed the most objectively useful training data were among the worst performers, particularly in the RB task. If active and passive-yoked learners are assumed to update their beliefs through a common process (as would be predicted by all existing models of human categorization) then this strong pattern of divergence is unexpected.

Finally, we found a main effect of category structure. Overall, participants in the II task performed more poorly at the task. Also, even though active learners in the II condition out-performed their passive counterparts, they were unable to boost performance near to RB levels. In addition, their sampling behavior suggests that (outside of a few surprising exceptions, see Figure 3C) most participants were unable to sample near the diagonal category margin, as would be predicted by an optimal information selection strategy (Oaksford & Chater, 1994). In the following section, we present a simple model-based analysis of each of these effects.

A Probabilistic Model of Decision-Bound Learning

While there have been a number of models proposed for how people classify items using rules in continuous dimension spaces, there have been fewer attempts to articulate an inference procedure for such models (c.f., Nosofksy & Palmeri, 1998). As a result, there were two key properties that guided the development of our modeling framework. First, we wanted a way to specify a strong inductive bias towards uni-dimensional rules along either stimulus dimension (similar to the default verbal system in Ashby et al., 1998). Most existing models can specify a prior bias towards a particular dimension (e.g., based on salience), but not a more general preference for arbitrary uni-dimensional rules (Heller et al., 2009). Second, analysis of the decision rules that participants use from one block to the next suggested that these were updated in a rather rapid fashion characteristic of serial hypothesis testing.

These concerns led us to a probabilistic model of classification which assumes that the goal of learning is to discover the latent parameters of a simple linear decision boundary. In our model, the probability that an observation, o^t, on trial t falls in category A is assumed to depend on a set of latent model parameters $\{\mathbf{w}, b, \sigma\}$:

$$P(o^t = A|\mathbf{w}, b, \sigma) = (1 + exp(-\sigma(\sum_i w_i o_i^t) - b))^{-1} \quad (1)$$

where o_i^t is the stimulus value of dimension i. Since the classification is binary, $P(o^t = B|\mathbf{w}, b, \sigma) = 1 - P(o^t = A|\mathbf{w}, b, \sigma)$. The weight vector, $\mathbf{w}$, contains the decision weight assigned to each dimension. The bias term, b, allows fine adjustments to the position of the decision rule in the stimulus space. Finally, the slope of the sigmoid function is controlled by σ
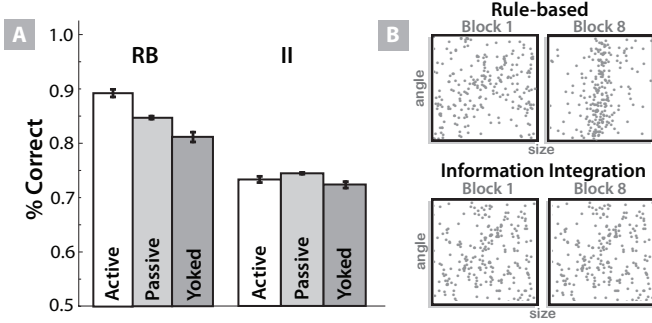
251

Figure 4: **A**: Expected accuracy of models trained on subject data. Active models which consider multiple hypotheses show an improvement in accuracy. **B**: Samples generated by the models during learning reflect the distribution of hypothesized rules; RB models are initially random but converge on the correct rule by the final block, while II models remain widely dispersed at both points in training.

which reflects how deterministic the decision rule is. Thus, each parameter combination $\{\mathbf{w}, b, \sigma\}$ reflects a unique decision rule or hypothesis about the category. The likelihood of a particular set of labeled observations $\mathcal{D} = \{o^1, .., o^t\}$ is given by $P(\mathcal{D}|\mathbf{w}, b, \sigma) = \prod_t P(o^t|\mathbf{w}, b, \sigma)$ (see Courville et al., 2003 for a similar approach). This basic model is equivalent to an equal variance Gaussian mixture model with two components.

We assume that learners are strongly biased toward unidimensional rules along either dimension. Accordingly, we defined a prior over the decision weights $\mathbf{w} = \{w_1 = cos(\theta), w_2 = sin(\theta)\}$, where θ is the angle of the vector corresponding to the decision boundary. We created a piece-wise scheme for translating θ into relative distances (bound between 0 and 1) from the horizontal axis:

$$r = \begin{cases} (2\theta)/\pi & : 0 < \theta \le \frac{\pi}{2} \\ (2(\pi-\theta))/\pi & : \frac{\pi}{2} < \theta \le \pi \\ (2(\theta-\pi))/\pi & : \pi < \theta \le \frac{3\pi}{2} \\ (2(2\pi-\theta))/\pi & : \frac{3\pi}{2} < \theta \le 2\pi \end{cases} \quad (2)$$

with $r \sim Beta(\alpha, \beta)$. Using this form, α and β act as a type of abstract attention weight (i.e., $\alpha = \beta < 1$ result in a general preference for rules along a single dimension. $\alpha, \beta < 1$ but $|\alpha - \beta| > 0$ results in a slight preference for one stimulus dimension over the other. $\alpha = \beta = 1$ implies no preference for rules of a particular orientation). The prior over the bias term was a Gaussian centered in the middle of the stimulus space, $b \sim N(0, 75)$, and the prior on the noise parameter was $\sigma \sim Beta(2, 1)$ (implying a mild preference for deterministic rules). Given these priors and the likelihood given in Equation 1, it is possible to infer the posterior distribution over the model parameters using Bayes rule. However, since full Bayesian updating in such a model is intractable, we assume that participants maintain an impoverished representation of the posterior distribution using a small set of point estimates from the posterior (similar to Sanborn et al., 2006).

At a given point in time we assume the learner has in mind a decision rule which can be characterized by param-

eter set $p^t = \{\mathbf{w}^t, b^t, \sigma^t\}$. On each trial, a new set of parameters p^{t+1} is proposed (or generated) which represents a change to the current rule. The learner is assumed to compare this new hypothesis to the old one and "accept" it as the new hypothesis if it provides a better account of the data (weighted by the prior belief in that parameter combination). If the new hypothesis results in a worse account of past data, it is accepted in proportion to the relative posterior likelihood of the new hypothesis compared to the old, otherwise the current parameter estimate remains unchanged. This procedure is similar to the Metropolis-Hastings algorithm (a form of Markov-Chain Monte Carlo) with an additional parameter k dictating the likelihood of accepting a proposal with a lower posterior estimate, giving the acceptance function $P(\mathcal{D}|p^{t+1})/(P(\mathcal{D}|p^t) + k)$. Proposals were generated from independent Normal distributions centered on the current parameter estimates: $\mathbf{w}^{t+1} \sim N(\mathbf{w}^t, \pi/2)$; $b^{t+1} \sim N(b^t, 20)$; $\sigma^{t+1} \sim N(\sigma^t, .05)$. The computational demands of this procedure are low: the learner is assumed to maintain a single hypothesis at any point in time. On each trial they must simply generate a new hypothesis and judge its relative quality. While we began with the simplification of assuming that the learner considers a single hypothesis on every trial, it is also possible that participants consider multiple hypotheses which are simultaneously updated in the same way.

Finally, we assume that the learner only stores n recent observations in memory, and evaluates the likelihood of a hypothesis over this limited set. This limitation results in ongoing shifts in the estimated decision rule, consistent with the variability in participants' response behavior throughout the task. Given the strong prior favoring rules along a single dimension, the estimate of the decision weights $\mathbf{w}$ will tend to bounce between these different modes of the hypothesis space, and convergence on the correct mode will be sensitive to the usefulness of recent training samples. This incremental, top-down hypothesis search may explain divergences between training conditions seen in our empirical results.

Evaluation of the Model The first goal of the simulation was to reproduce the difference in performance between the RB and II tasks. Individual models were trained with the data from passive-normal and passive-yoked participants in our experiment (the active group is addressed below). For this initial simulation the following parameter settings were used: $\alpha = \beta = 0.001, k = 1, n = 4$. Expected accuracy on each test item was calculated using the predicted likelihood that the item belonged to the correct category. Expected accuracy was averaged over test blocks and across 100 runs. The comparison of passive models (Figure 4A) shows a strong difference in accuracy between RB and II tasks as seen in our behavioral results. Due to the strong prior bias toward single-dimensional rules, in the RB task the model quickly converges on a rule similar to the true boundary, despite only retaining a small number of recent observations. In the II task, however, the model alternates between single-dimensional rules on different dimensions.

One way that the model can account for differences between active and passive-yoked groups is by assuming that active participants represent more than one hypothesis at any given time (consistent with the generalized "engagement" hypothesis described in the Introduction). In the model, this might correspond to an increase in the number of point estimates of the posterior maintained in working memory. To evaluate this idea, active participants were modeled using a set of 5 posterior samples per run (in contrast to one sample used for the passive groups), with the additional assumption that learners classify items according to the most likely hypothesis from the set they are considering. As seen in Figure 4A, the greater number of samples leads to higher accuracy over the passive groups in the RB task, but not in the II task. While a change in the number of particles maintain considered is consistent with the idea that active learners are more cognitively engaged in the task (and thus search the hypothesis space more effectively), further work is needed to directly test this representational hypothesis. At the very least, the potential divergence between the sequence of data observed in the task and the sequence of hypotheses considered by the learner provide a potential mechanism for explaining the active/passive-yoked distinction.

Finally, we were interested if samples generated by the active models show the same pattern as produced by our participants. Simulated samples were generated using *margin sampling*, in which an observation is most likely to occur when its predicted likelihood of belonging to category A and B are equal (i.e., the likelihood of making an observation o^t was proportional to $1 - |P(o^t = A|\mathbf{w}, b, \sigma) - P(o^t = B|\mathbf{w}, b, \sigma)|$). As seen in Figure 4B, the predict sampling distribution qualitatively matches the behavioral results. In the first block of both tasks, the model produces samples that are widely dispersed throughout the feature space. By the final block, RB models have converged on the true boundary, querying the margin of the boundary where uncertainty is greatest. In the II task, the diffuse distribution of samples reflects the variability in the hypotheses under consideration.

Conclusions

In our experiment, active learners were able to make informative queries to support their own learning, but this ability was more successful for RB categories than for II categories. Our simulation results explain this difference in terms of a bias toward considering rules along a single dimension. In addition, we evaluated one explanation for the divergence between active and passive-yoked participants, namely that active participants consider a greater number of hypotheses about the latent category structure. Our general finding that the effectiveness of active sampling may depend on the structure of the category adds to recent work examining active learning in binary classification tasks (Castro et al., 2008). While a number of theorists have attempted to explain active data selection in terms of optimal information gain (Oaksford & Chater, 1994; Nelson, 2005), our results suggest that the ability to design

useful queries is strongly limited by the hypothesis search process that guides learning. To the degree that participants prefer particular types of rules, their sampling behavior will tend to be sub-optimal when the target rule mismatches these expectations, a similar point made in analyses of active machine learning (Mackay, 1992; Settles, 2009). In summary, active learning may promote learning, but it works best when you have a strong and correct idea of what you are trying to learn.

References

Ashby, F., Alfonso-Reese, L., Turken, A., & Waldron, E. (1998). A neuropsychological theory of multiple system in category learning. *Psychological Review*, *105*(5), 442-481.

Ashby, F., Maddox, W. T., & Bohil, C. J. (2002, Jul). Observational versus feedback training in rule-based and information-integration category learning. *Memory & cognition*, *30*(5), 666–77.

Ashby, F., Queller, S., & Berretty, P. (1999). On the dominance of unidimensional rules in unsupervised categorization. *Perception & Psychophysics*, *61*, 1178-1199.

Bruner, J. (1961). The art of discovery. *Harvard Educational Review*, *31*(21-32).

Castro, R., Kalish, C., Nowak, R., Qian, R., Rogers, T., & Zhu, X. (2008). Human active learning. In *Advances in neural information processing systems* (Vol. 21). Cambridge, MA: MIT Press.

Courville, A., Daw, N., Gordon, G., & Touretsky, D. (2003). Model uncertainty in classical conditioning. *Advances in Neural Information Processing Systems*, *20*.

Gureckis, T., & Markant, D. (2009). Active learning strategies in a spatial concept learning game. *Proceedings of the 31st Annual Conference of the Cognitive Science Society*.

Heller, K., Sanborn, A., & Chater, N. (2009). Hierarchical learning of dimensional biases in human categorization. In J. Lafferty & C. Williams (Eds.), (Vol. 22). Cambridge, MA: MIT Press.

Kruschke, J. (1993). Human category learning: Implications for backpropagation models. *Connection Science*, *5*, 3-36.

Kruschke, J. (2008). Bayesian approaches to associative learning: From passive to active learning. *Learning and Behavior*, *36*(3), 210-226.

Kuhn, D., Black, J., Keselman, A., & Kaplan, D. (2000). The development of cognitive skills to support inquiry learning. *Cognition and Instruction*, *18*(4), 495–523.

Lagnado, D. A., & Sloman, S. (2004, Jul). The advantage of timely intervention. *Journal of experimental psychology Learning, memory, and cognition*, *30*(4), 856–76.

Mackay, D. (1992). Information-based objective functions for active data selection. *Neural Computation*, *4*, 590-604.

Nelson, J. (2005). Finding useful questions: On bayesian diagnosticity, probability, impact, and information gain. *Psychological Review*, *112*(4), 979-999.

Nosofksy, R., & Palmeri, T. J. (1998). A rule-plus-exception model of classifying objects in continuous-dimension spaces. *Psychonomic Bulletin & Review*, *5*(3), 345-369.

Oaksford, M., & Chater, N. (1994). A rational analysis of the selection task as optimal data selection. *Psychological Review*, *101*(4), 608-631.

Sanborn, A., Griffiths, T., & Navarro, D. (2006). A more rational model of categorization. In R. Sun & N. Miyake (Eds.), *Proceedings of the 28th annual meeting of the cognitive science society*. Mahwah, NJ: Erlbaum.

Settles, B. (2009). Active learning literature survey. *Technical Report*.

Sobel, D., & Kushnir, T. (2006). The importance of decision making in causal learning from interventions. *Memory and Cognition*, *34*(2), 411.

Steyvers, M., Tenenbaum, J., Wagenmakers, E., & Blum, B. (2003). Inferring causal networks from observations and interventions. *Cognitive Science*, *27*(3), 453-489.