

How goals affect information seeking

Gili Karni (gili@princeton.edu)

Princeton Neuroscience Institute, Princeton University

Nathaniel D. Daw (ndaw@princeton.edu)

Princeton Neuroscience Institute & Department of Psychology, Princeton University

Yael Niv (yael@princeton.edu)

Princeton Neuroscience Institute & Department of Psychology, Princeton University

Abstract

This study investigates how goals influence information sampling strategies in active learning. Previous work in this area compared different sampling heuristics while holding constant participants' goals (e.g., the final incentivized test). In a behavioral experiment, we examine the effect of generation-driven versus discrimination-driven goals on information-seeking sampling strategies by manipulating the (pre-declared) test condition across subjects. Our results suggest that goals affect information sampling, with discrimination-driven tasks leading to more sampling around class borders ("Label-Margin" sampling strategy) and generation-driven tasks leading to more sampling around class centers ("Most-Certain" sampling strategy). Moreover, we show that strategies evolve over time and are related to the performance of participants. These findings highlight the importance of considering goals in understanding human information-seeking behavior.

Keywords: information-seeking; active-learning; goal-directed learning;

Introduction

How do our goals shape the way we search for information? Although naive accounts of information-seeking envision that people simply seek generically to reduce uncertainty, more refined decision-theoretic analyses of the value of information emphasize that different information will be more or less valuable depending on how it serves the agent's goals (Yarbus, 1967; Gittins, 1979; Kaelbling, Littman, & Cassandra, 1998; Yang, Wolpert, & Lengyel, 2016). This study investigates the critical role of goals in shaping active learning strategies, focusing on how goals influence information sampling. By exploring the interaction between different learning objectives (i.e., goals) and sampling behaviors, we aim to illuminate a key facet of human information seeking and, more broadly, active learning.

Consider two individuals engaged with the art world: a painter and an art critic. A painter motivated by the desire to create original and high-quality works might study the specific characteristics of the explored styles and techniques to refine their craft. An art critic, in contrast, might focus on identifying subtle contrasts between works, such as differences in technique, emotion, or historical context. This intuitive difference highlights a fundamental question: How do our objectives shape the way we seek and process information? In other words, how do we direct our active learning?

Active learning is the process of acquiring new knowledge in a self-directed manner, e.g. typically by the agent

choosing which examples to investigate. Active learning has been demonstrated to be highly effective in various contexts (Gureckis & Markant, 2012; Jha, Ashwood, & Pillow, 2024). However, much of the existing research on active learning by humans involves open-ended goals or a single fixed objective (Coenen et al., 2019; Wilson et al., 2014). Although these approaches provide valuable information about exploratory behavior, generally they do not directly address the hypothesis that sampling is driven by the value of information in serving the agent's task or goals. To address this gap, we draw on statistical principles and leverage the distinction of generation-driven versus discrimination-driven approaches in category learning (Ng & Jordan, 2001) and on previous results showing that people can use both generative and discriminative approaches in category learning tasks (Hsu & Griffiths, 2010). This framework allows us to systematically capture how sampling strategies evolve under different goals.

We tackle this challenge through a behavioral experiment that augments an active-learning categorization task with a manipulation of the participants' goals, by instructing them that after learning, they will be given different tests. In the experiment (based on Markant, Settles, and Gureckis (2016)), participants first learn about different antennas that can receive one of three channels, they actively choosing which antennas to investigate. Participants were assigned one of two goals, in the form of a pre-announced test to follow the learning: a *generation-driven goal*, where they were required to "install" antennas that receive specific channels, or a *discrimination-driven goal*, where they were required to "inspect" installed antennas to determine what channel they receive. We hypothesized that these different goals would affect their sampling strategy during active learning, with the discrimination test promoting investigation of category boundaries and the generation test favoring sampling near category centers. We further hypothesize that this goal-driven strategy distinction will strengthen over time, as participants learn more about the state space.

In accordance with this hypothesis, our findings reveal that participants with discrimination-driven goals concentrated their efforts near category boundaries while participants with generation-driven goals concentrated their efforts around the centers of the categories. We also show that this difference was indeed developing over the course of the task. These results underscore the adaptive nature of human inquiry and

demonstrate that goals play a critical role in shaping how we seek information. By incorporating goal-directed frameworks into active-learning research, we can deepen our understanding of human behavior and enhance the design of educational tools, user interfaces, and machine-learning systems to align more effectively with diverse learning objectives.

Uncertainty sampling strategies

A generic approach to information seeking is to attempt to minimize expected uncertainty. However, all uncertainty is not equal. Research on statistical decision theory and optimal control quantifies how and why information can be valuable: if it helps the agent make more rewarding choices later (Gittins, 1979; Kaelbling et al., 1998; Friston et al., 2015). A treasure map can lead to riches; a weather forecast may help you avoid getting wet. This perspective motivates our central hypothesis, that what types of uncertainty someone is motivated to resolve depends how they expect to later use it.

Meanwhile, even setting aside specific goals, the seemingly simple objective of minimizing uncertainty is a complex task and there are multiple approaches to it. One useful distinction is between falsifying or verifying methods – the former seeks information most likely to invalidate a theory while the latter searches for information to strengthen it (Wason, 1960; Klayman & Ha, 1987; Oaksford & Chater, 1994). Motivated by this dichotomy, investigators have proposed several heuristics for reducing uncertainty and studied to what extent these capture human behavior (Markant et al., 2016; Keller, Taylor, & Brunyé, 2020).

Here, we draw on these proposed heuristics to describe human behavior, but test the hypothesis that which of these strategies people tend to favor depend on their goals. Specifically, following Markant et al. (2016) and Settles (2011), we compare two uncertainty-driven sampling strategies f that compute the predictive informativeness of a new data point based on the agent’s current uncertainty about a classification problem. Both strategies utilize the learner’s posterior distribution over the class label y as a function of items x , $p(y|x)$ given previous data, but differ in how, at each step, they choose the next sample item to label. For each, we define a score function f and maximize it greedily to choose the next sample: $x^* = \arg \max_x f(p(y|x))$. Intuitively, the objectives of these strategies can be summarized as follows: Label-Margin focuses on resolving ambiguities in areas where two classes have high likelihood, that is, in the margins between different classes, and Most-Certain aims to maximize confidence in the top category. Importantly, “Label-Margin” sampling differs from a “Label Entropy” strategy that targets overall uncertainty. The scores given by each model for different $p(y|x)$ in a 3D simplex (label probabilities of a three-class domain, as in our experiment) are illustrated in figure 1.

In more detail, the “Label-Margin” (LM) strategy is defined as the complement of the difference between the probabilities of the top two most probable class labels for an obser-

vation:

$$f_{LM}(p(y|x)) = 1 - (p(y_1|x) - p(y_2|x)) \quad (1)$$

and is defined only for situations with more than 2 categories. This strategy prioritizes querying data where the margin is smallest, effectively breaking down multiclass problems by sequentially reducing binary classification ambiguities (Nosofsky, 2011). “Label-Entropy” sampling strategy is different as it is defined by Shannon entropy over the class-label belief distribution:

$$f_{LE}(p(y|x)) = - \sum_{i \in C} p(y_i|x) \ln p(y_i|x) \quad (2)$$

where C are the possible category labels. It is a widely used metric in both cognitive psychology and machine learning (Nelson & Movellan, 2000; Kruschke, 2008; Jha et al., 2024; Friston et al., 2015), serving as an approximation for normative information-gain models.

In contrast, the “Most-Certain” (MC) strategy scores each item according to the probability of its most likely label:

$$f_{MC}(p(y|x)) = \max_i (p(y_i|x)) \quad (3)$$

This scoring rule is associated with confirmatory sampling, a method that, in some cases, has been observed to be effective for human decision-making (Klayman & Ha, 1987; Austerweil & Griffiths, 2011).




Figure 1: An illustration of the three uncertainty-sensitive sampling strategies plotted over a 3D simplex where each location represents $p(y|x)$ for each of three possible labels y . The color reflects the value of each score function over this probability $f(p(x|y))$.

In the present study, we predicted that LM and MC would be differentially employed for the discriminative versus generative conditions, respectively, because the former focuses information-gathering efforts on category boundaries while the latter focuses them nearer class prototypes. We consider “Label-Entropy” sampling as a more generic uncertainty-driven exploration strategy that is not specifically tied to the task’s goals.

Methods

Experiment

We designed a behavioral experiment, adapting a task by Markant et al. (2016), with the aim of dissociating generation-

driven from discrimination-driven information-seeking behaviors. The experiment consisted of an active category-learning task, in which participants learn to categorize a stimulus varying along two dimensions into three categories.

Participants A total of $N = 206$ participants, recruited through Prolific (Prolific, 2014), completed an online categorization study. $N=45$ were subsequently excluded from the analyzes due to failing our post-instruction comprehension checks. We present data from $N=161$ (mean age = 35.8 years ($SD = 11.0$); 66 women, 93 men, 2 nonbinary). We paid participants \$15 for their participation and up to \$5 as a bonus for high test accuracy.

Stimuli Stimuli were visualized as a 'loop antenna' varying on two dimensions: angular position and size. In this two-dimensional state space, a channel is associated with each stimulus according to a deterministic, ternary classification rule. The size dimension is linear and the rotation dimension is circular. See figure 2A. Both dimensions were discretized into 100 bins, controlled by either the keyboard or the mouse. We split the space in a rule-based manner (i.e., with class boundaries setting a threshold on one or both dimensions), which has been shown to be simpler to learn (Roark & Chandrasekaran, 2023).

Procedure The study included eight blocks. Each block consisted of 10 training trials and 36 test trials. During training, participants repeatedly created an antenna to receive information about, by adjusting its two feature dimensions using the keyboard or mouse. Next, they reported the likelihood they believed their antenna would receive each of the three channels on a 3D simplex. Importantly, during the instructions, participants had been taught how to use the simplex and were tested on their understanding of it (via an instructed sequence of desired probabilities to be marked). Participants who failed to correctly use the simplex were excluded from the analysis. Finally, they were shown the channel that the antenna they had designed receives, alongside their antenna (figure 2B, left). Participants could adjust their antennae or revise their likelihood estimates as many times as they wished before submitting their final answer within the allotted time. The response time was set to 10 seconds for the design of an antenna and 5 seconds for reporting its channel likelihood. Feedback was presented for three seconds.

Testing depended on the condition: In the Generation condition, participants were asked to 'install' antennas that would receive a specific channel, and in the Discrimination condition participants were asked to 'inspect' which channel an installed antenna receives (figure 2B, right). Participants were randomly assigned to one of the two conditions, and notified of their 'role' and that success in this role would earn them a monetary bonus in the instructions. Test stimuli in each block, and their associated class, were uniformly drawn from

each of the equally sized 36 bins of the state space and presented to participants in pseudo-randomized order. In the Discrimination condition, participants were presented with each test antenna and asked to report what channel they receive, while in the Generation condition they were asked to generate an antenna that would receive the channel associated with each test antenna. Each of the categories comprised a third of the state space, such that test stimuli sampled each of the three channels equally often. Participants were not presented with feedback at every trial, but received their average accuracy score at the end of each block. Like in training, participants could adjust their test answers as many times as they liked before submitting their final answer and were given 10 seconds to do so.




Figure 2: **Experimental Design:** [A] The stimuli state space was 2 dimensional with one linear axis that controlled the size of the antenna (x-axis) and one circular axis that controlled the angular rotation (y-axis). The state space was divided into 3 equal-sized parts, each receiving one channel. An example stimulus for each category is illustrated on the right. [B] The experiment included 8 blocks, each consisting of a training phase (left) and a testing phase (right). See **Procedure**.

Approach

Simulation To compare the performance of the three sampling strategies in this task, we used each of them together with ideal observer model to learn from observations and simulate active category learning. The model maintained a belief state over the category boundaries mapping the three categories in a discretized 100×100 2D grid over the state space.

The model assumed a uniform distribution over all plau-

sible category boundaries in the belief state. On each trial, the model computed the marginal class probability for every point in the grid given its belief state. This probability distribution was then passed into one of the sampling strategy functions (“Label-Entropy”, “Label-Margin”, or “Most-Certain”; Equations 1–3) to select a new data point. The data point was sampled from the temperature-scaled softmax-transformed output of the function. Upon receiving an observation (the true class label of the chosen sample), the model updated the list of plausible category boundaries by discarding those that were inconsistent with the new observation. This model is a simple instance of Bayesian inference, where the likelihood function acts as a binary filter that eliminates category boundaries inconsistent with the observed data. The posterior remains uniform over the surviving category boundaries.

At the end of each block, the model performed both Generation and Discrimination tasks. To calculate the discrimination test score, we computed the model’s accuracy in classifying a test set consisting of $N = 36$ uniformly sampled data points. The generation score reflected the class accuracy of $N = 36$ samples drawn from marginal class probabilities according to the belief state (12 samples for each of the three classes).

Each sampling strategy was simulated with $N = 250$ independent repetitions. One set of parameters, including a learning rate for the belief update ($\eta = 0.25$) and an inverse-temperature for softmax sampling ($\beta = 14$), was optimized to maximize overall performance.

Empirical sampling strategy scores We computed sampling strategy scores for each participant and each learning trial based on participants’ self reported class-probability likelihoods $p(y|x)$. That is, for each antenna chosen by a participant, we used their report on the simplex to calculate the probability that sample resulted from each of the sampling strategies as per figure 1. Using the likelihood rather than the raw sample coordinates allowed us to incorporate participants’ beliefs while inferring their sampling strategies. We used these likelihoods to calculate three strategy scores, one for each sampling strategy as described above (see Equations 1-3).

Next, we computed the relative difference between the “Label-Margin” sampling and “Most-Certain” sampling scores, which we refer to as the sampling strategy index:

$$SSI = \frac{LM - MC}{LM + MC}. \quad (4)$$

The participant-level sampling strategy index (SSI) was calculated as the mean index for each participant across blocks. Similarly, at the block level, it represents the mean index per block.

The SSI ranges from -1 to 1, where positive values indicate greater reliance on “Label-Margin” sampling and negative values reflect a tendency towards “Most-Certain” sampling. An SSI near 0 suggests no consistent preference for

either strategy. Thus, higher SSI values correspond to participants favoring margin-focused sampling, while lower values correspond to more center-focused choices.

We focused our analysis on the LM and MC strategies because they highlight the key distinction between the two test conditions, as per our hypothesis. Using a normalized difference ensured comparability between conditions, participants, and time (blocks). More importantly, this metric captured the competing nature of the two strategies, providing a clear and interpretable summary.

Results

All participants included in the analysis showed high accuracy on test trials, with participants in the Generation condition reaching higher accuracy than participants in the Discrimination condition ($t(161) = 11.3, p < 0.001$; figure 3).




Figure 3: **Behavioral data:** Mean (shading: 95% CI) test accuracy of $N=75$ in the Discrimination condition (left) and $N=86$ participants in the Generation condition (right) across 8 task blocks. All participants performed well above chance (33%), with participants in the Generation condition achieving higher asymptotic performance.

The Generation conditions may present an inherently easier task in our setup because participants can use easy, low-ambiguity antennas during the test (Markman & Ross, 2003; Markant & Gureckis, 2014), whereas in the discrimination test, participants do not get to influence the antennas presented, which are distributed uniformly across the space. This imbalance between ‘easy’ and ‘hard’ examples likely contributes to the observed accuracy gap. Indeed, while overall accuracy in the Discrimination condition increased from 55% in block 1 to 79% in block 8 (as seen in figure 3), accuracy on the less ambiguous antennas—those farther from category boundaries—was higher, ranging from 60% to 89%, suggesting that the difficulty of test antennas plays a critical role in the observed accuracy.

Additionally, there was a tendency toward using the “Most-Certain” strategy across all participants in both conditions. This outcome is inconsistent with the “Label-Margin” sampling preference reported by Markant et al. (2016). However, although our study design was adapted from Markant et al., there were important differences in how participants

reported their probability judgments in the two experiments, which may explain the divergent behavior. In our task, participants used a probability simplex to report category judgments directly, whereas Markant et al. employed three independent response scales presented one at a time, which were not required to sum to 1 and were later normalized for analysis. As such, participants in Markant et al.’s study may not have realized that marking “100%” on all three channels corresponds to 33% on all channels, etc. Since these reports were the main data analyzed in both studies, such differences in response format can significantly influence how participants report uncertainty, potentially contributing to the divergent results (Fischhoff, Slovic, & Lichtenstein, 1978).

Simulations showed different sampling strategies are better for different tasks. Simulations of the sampling strategies demonstrated that the “Label-Margin” sampling strategy performs best in the Discrimination task, while the “Most-Certain” sampling strategy performs best in the Generation task (figure 4).




Figure 4: **An ideal observer model showed that different sampling strategies achieve higher accuracy depending on the task.** Performance is shown across blocks. In the Discrimination task (left), the LM strategy attained a higher test accuracy than MC. In the Generation task (right), the MC sampling strategy achieved the higher test accuracy.

Participants’ sampling strategies are congruent with their goals. We characterized each sampling choice by how well it was predicted by the three active sampling strategies, using the probabilities participants reported for each of the classes, and the strategies’ score functions (figure 1). As our main focus was a hypothesized trade-off between “Label-Margin” and “Most-Certain” sampling, we computed an index summarizing relative reliance on these two strategies. Indeed, participants in the Discrimination condition were significantly more likely to use the “Label-Margin” sampling strategy and less likely to use the “Most-Certain” sampling strategy compared to those in the Generation condition. This effect was reflected in a significantly larger sampling strategy index for the Discrimination condition ($t(161) = 2.1; p = 0.037$; figure 5).

The differences between the sampling index score in the Discrimination and Generation conditions was driven by significant differences in both LM ($t = 2.15, p < 0.05$) and MC ($t = -1.97, p < 0.05$) strategies. No significant difference was observed between the two groups in their “Label-Entropy” sampling score.




Figure 5: **People’s sampling strategies are congruent with their goals.** Participants in the Discrimination condition showed a significantly higher mean sampling strategy index as compared to participants in the Generation condition, reflecting increased use of LM and decreased use of MC strategies in the former as compared to the latter group. Light dots: individual participants. Dashed line: overall mean index score. * denotes $p < .05$

Participants’ initially used both MC and LM strategies equally, with use of the goal-appropriate strategy increasing across blocks. We further examined how strategy usage changed with training on the task. The sampling strategy index started around 0, suggesting use of both strategies to a similar extent. The index then decreased across blocks, with a notably faster decline in the Generalization condition than in the Discrimination condition. This pattern suggests that participants gradually shifted toward more goal-aligned strategies as they gained experience with the task. A mixed-effects model confirmed this pattern: there was a significant overall decline in the sampling strategy index over blocks ($z = -6.97; p < 0.001$), and a significant interaction between condition and block ($z = -8.31; p < 0.001$), indicating a sharper decline in the Generation condition relative to the Discrimination condition. This sharper decline indicates a stronger tendency to shift toward the “Most-Certain” sampling strategy over time, as opposed to the “Label-Margin” sampling strategy (figure 6A).

Participants perform better when they rely more on the goal-appropriate strategy. Finally, we tested whether the

adoption of the appropriate sampling strategy actually predicted better performance on either task. Individuals' sampling strategies were indeed strongly associated with their performance in the categorization task. The sampling strategy index, which quantifies the extent to which participants use "Label-Margin" sampling strategy versus "Most-Certain" sampling strategy, was positively correlated with accuracy in the Discrimination condition ($r = 0.23; p = 0.05$). In contrast, a negative correlation was observed in the Generation condition ($r = -0.29; p = 0.007$). A Fisher r -to- z transformation confirmed that this difference was statistically significant ($z = 3.31; p < 0.001$). These results suggest that participants who prioritized goal-appropriate sampling strategies tended to achieve higher task accuracy (figure 6B).




Figure 6: [A] **Sampling strategy index declines with blocks.** The decline in the Generation condition was faster than in the Discrimination condition, suggesting more (and more rapid) reliance on the MC strategy in the latter. Shading: 95% CI. [B] **Sampling strategy index correlates with performance,** with a positive correlation in the Discrimination condition (red) and a negative correlation in the Generation condition (blue). Each dot is a participant.

Discussion

Our findings reveal that task goals significantly shape sampling strategies in active learning. Participants with discrimination-driven goals preferred resolving ambiguity near class boundaries ("Label-Margin" sampling), while those with generation-driven goals focused on prototypical examples ("Most-Certain" sampling). These results highlight how human learners adaptively align strategies with task demands.

Participants in the Discrimination condition in our antenna task were more likely to use "Label-Margin" sampling when choosing samples to learn the label for, emphasizing learning to differentiate between categories. In contrast, participants in the Generation condition focused their active sampling on prototypical examples, suggesting greater use of the "Most-Certain" strategy. Simulations indeed confirmed that "Label-Margin" sampling was preferred for the discrimination task, while "Most-Certain" sampling excelled in the generation task, aligning empirical and theoretical insights.

Participants initially used both "Label-Margin" and "Most-Certain" strategies to a similar extent, but shifted toward goal-aligned approaches with experience. The Generation condition showed faster adaptation, indicating a clearer alignment with "Most-Certain" sampling. This demonstrates the flexibility of human learning in optimizing exploration based on task requirements. Additionally, participants who adopted goal-specific strategies achieved higher accuracy. The sampling strategy index was correlated positively with performance in the Discrimination condition, and negatively in the Generation condition, underscoring the importance of aligning strategies with task goals for effective learning. These findings challenge task-agnostic heuristics in active learning, emphasizing goal-driven adaptability.

While our study provides valuable insights, there are some limitations to address. First, the pre-declared test condition might have influenced participants' sampling strategies by creating explicit expectations about task demands. Future research could explore whether similar patterns emerge in contexts where goals are more implicit, or only revealed gradually through feedback. Additionally, it remains unknown how participants would adjust their sampling behavior if task goals were to change mid-task; whether they would flexibly shift strategies or persist in previously learned behaviors. Investigating the flexibility of goal-specific sampling behaviors, as well as their generalization to more complex tasks with higher-dimensional stimuli, represents a promising direction for future research.

In conclusion, our study highlights the critical role of goals in shaping information-seeking behavior, revealing how human learners dynamically adjust their strategies to optimize performance.

References

- Austerweil, J. L., & Griffiths, T. L. (2011). Seeking confirmation is rational for deterministic hypotheses. *Cognitive Science*, 35(3), 499–526.
- Coenen, A., Nelson, J. D., & Gureckis, T. M. (2019). Asking the right questions about the psychology of human inquiry: Nine open challenges. *Psychonomic Bulletin & Review*, 26, 1548–1587.
- Fischhoff, B., Slovic, P., & Lichtenstein, S. (1978). Fault trees: Sensitivity of estimated failure probabilities to problem representation. *Journal of Experimental Psychology: Human Perception and Performance*, 4(2), 330.
- Friston, K., Rigoli, F., Ognibene, D., Mathys, C., Fitzgerald, T., & Pezzulo, G. (2015). Active inference and epistemic value. *Cognitive neuroscience*, 6(4), 187–214.
- Gittins, J. C. (1979). Bandit processes and dynamic allocation indices. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 41(2), 148–164.
- Gureckis, T. M., & Markant, D. B. (2012). Self-directed learning: A cognitive and computational perspective. *Perspectives on Psychological Science*, 7(5), 464–481.

- Hsu, A. S., & Griffiths, T. E. (2010). Effects of generative and discriminative learning on use of category variability. In *32nd annual conference of the cognitive science society*.
- Jha, A., Ashwood, Z. C., & Pillow, J. W. (2024). Active learning for discrete latent variable models. *Neural computation*, 36(3), 437–474.
- Kaelbling, L. P., Littman, M. L., & Cassandra, A. R. (1998). Planning and acting in partially observable stochastic domains. *Artificial intelligence*, 101(1-2), 99–134.
- Keller, A. M., Taylor, H. A., & Brunyé, T. T. (2020). Uncertainty promotes information-seeking actions, but what information? *Cognitive Research: Principles and Implications*, 5(1), 42.
- Klayman, J., & Ha, Y.-W. (1987). Confirmation, disconfirmation, and information in hypothesis testing. *Psychological review*, 94(2), 211.
- Kruschke, J. K. (2008). Bayesian approaches to associative learning: From passive to active learning. *Learning & Behavior*, 36, 210–226.
- Markant, D. B., & Gureckis, T. M. (2014). Is it better to select or to receive? learning via active and passive hypothesis testing. *Journal of Experimental Psychology: General*, 143(1), 94–122.
- Markant, D. B., Settles, B., & Gureckis, T. M. (2016). Self-directed learning favors local, rather than global, uncertainty. *Cognitive science*, 40(1), 100–120.
- Markman, A. B., & Ross, B. H. (2003). Category use and category learning. In *Psychology of learning and motivation* (Vol. 43, pp. 307–353). Elsevier.
- Nelson, J., & Movellan, J. (2000). Active inference in concept learning. *Advances in neural information processing systems*, 13.
- Ng, A., & Jordan, M. (2001). On discriminative vs. generative classifiers: A comparison of logistic regression and naive bayes. *Advances in neural information processing systems*, 14.
- Nosofsky, R. M. (2011). The generalized context model: An exemplar model of classification. *Formal approaches in categorization*, 18–39.
- Oaksford, M., & Chater, N. (1994). A rational analysis of the selection task as optimal data selection. *Psychological review*, 101(4), 608.
- Prolific. (2014). *Prolific*. Retrieved from <https://www.prolific.com>
- Roark, C. L., & Chandrasekaran, B. (2023). Stable, flexible, common, and distinct behaviors support rule-based and information-integration category learning. *npj Science of Learning*, 8(1), 14.
- Settles, B. (2011). From theories to queries: Active learning in practice. In *Active learning and experimental design workshop in conjunction with aistats 2010* (pp. 1–18).
- Wason, P. C. (1960). On the failure to eliminate hypotheses in a conceptual task. *Quarterly journal of experimental psychology*, 12(3), 129–140.
- Wilson, R. C., Geana, A., White, J. M., Ludvig, E. A., & Cohen, J. D. (2014). Humans use directed and random exploration to solve the explore–exploit dilemma. *Journal of experimental psychology: General*, 143(6), 2074.
- Yang, S. C.-H., Wolpert, D. M., & Lengyel, M. (2016). Theoretical perspectives on active sensing. *Current opinion in behavioral sciences*, 11, 100–108.
- Yarbus, A. L. (1967). *Eye movements and vision*. Plenum Press.