How goals affect information seeking

Gili Karni (gili@princeton.edu), Nathaniel D. Daw (ndaw@princeton.edu), Yael Niv (yael@princeton.edu)
Princeton Neuroscience Institute, Princeton University
Princeton, NJ, USA

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 generative versus discriminative goals on information-seeking sampling strategies by manipulating the (pre-declared) test condition across subjects. Our results suggest that goals affect sampling behavior, with discriminative tasks leading to more sampling around class borders ("label-margin sampling"). These findings highlight the importance of considering goals in understanding human information-seeking behavior.

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

Introduction

Gathering information is a crucial aspect of learning. Specifically, deliberating which information is valuable and thus worth sampling makes active learning a more productive method than passively receiving 'random' samples (Jha, Ashwood, & Pillow, 2022). This deliberation, yielding sampling decisions, is a rich cognitive process grounded in our ability to value information (Gottlieb, 2018). In recent work, Markant et al. (2016) showed that people use a variety of search strategies when faced with a category-learning task.

We propose that this variation in strategies is, in part, driven by differences in people's goals, building on the hypothesis that the value of information is instrumental, and affects sampling. This is because although the value of information is often computed via heuristics, normatively it should follow from expectations about how the information gained will help obtain rewards. Thus, it depends not just on the information but also on one's goals. Hence, we predict that we should be able to modulate exploratory active-sampling behavior by manipulating participants' goals.

Methods

Experiment

Extending Markant et al. (2016)'s task, we designed a behavioral experiment aiming to dissociate generative from discriminative information-seeking behaviors. The experiment is an active category learning task in which participants learned about a stimulus that varied in two dimensions (depicted as an antenna with a bar and a circle at different angles; Fig. 1). In this two-dimensional state space, we associated a category label ("channel received") with each stimulus according to a deterministic, ternary classification rule. In the learning phase, participants designed sample antennas and were given feedback both on the category label (channel received by the antenna) and about distance from the centroid (reception strength; Fig. 1, top right screenshot).

To manipulate participants' goals, we instructed participants on one of two tasks on which they would be tested after learning, inspired by a well-known dichotomy in category learning: generative versus discriminative learning. The generative-inspired task is to create a new example matching a named category (e.g., "install an antenna that receives channel 1"), while the discriminative-inspired task is to label an example to its correct category (i.e., "inspect this antenna and determine channel received").

N=40 participants (mean age 35(s.d. 11); 2 excluded from analysis due to low accuracy at test) completed one of the two conditions. The experiment has 8 blocks, each consisting of ten 3-step training trials (Fig. 1; top) and 18 testing trials (Fig. 1; bottom). During training, participants were asked to create an antenna of their choosing (Fig. 1; top left), then estimate the likelihood that the antenna would receive each of the channels using a 3-category simplex (Fig. 1; top middle), and Last, they received feedback with the correct category and accuracy (Fig. 1; top right). In the test phase, participants were asked to either install an antenna (generative goal) or inspect an installed antenna (discriminative goal). Participants were told their average test accuracy at the end of each block and were rewarded according to their cumulative accuracy.

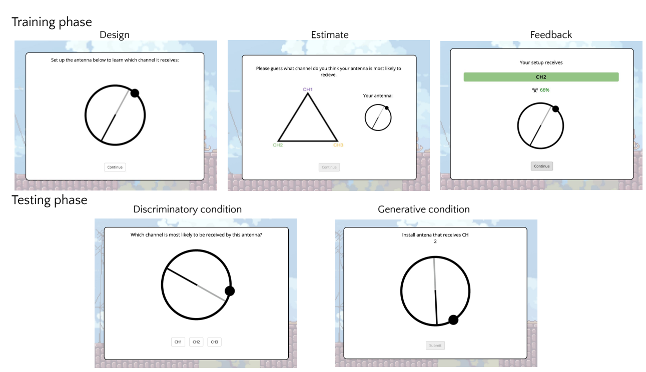

Figure 1: **Experimental design.** Top: Training trials involved antenna design, likelihood estimate, and feedback. Bottom: In the discriminative condition, participants inspected antennas to determine their channel (left); in the generative task they designed antennae to receive specific channels (right).

Theoretical model

We consider a 2-dimensional (input, x), 3-category (output label, y) classification task. We assume that participants model the category labels via a mixture of Gaussians where $p(y = i|x,\theta) \propto Normal(\mu_i, \sigma_i)$ for three unknown Gaussians. On each trial, the model updates its belief using Bayesian updating from the previous trial's observed label, chooses a new sample x via softmax of its utility function, and receives the true label y for its choice.

We compared three value-of-information heuristics that aim to explain the chosen samples by defining different greedy utility functions over the inputs. The first heuristic is Shannon's entropy (ENT; Fig 2A, eq 1). It reflects a naive normative approach to evaluating overall uncertainty in the data. The second is the "label margin" heuristic (LM; Fig 2A, eq 2), which measures the difference between the probability of the predicted class (p1) and the probability of the next most probable class (p2). Last, the "most certain" heuristic (MC; Fig 2A, eq 3) favors samples with the highest classification confidence.

To examine how the effectiveness of these sampling strategies varied with the agent's incentivized goals, we simulated samples using each of the three utility functions and used them to form three sets of parameter estimates for the Gaussian moments. We then compared these estimates in terms of their success in accomplishing the two terminal tasks. In the discrimination task, we estimated the category of 100 random samples and computed a classification accuracy score. In the generation task, we computed the mean squared error of samples generated from each model from the true means. We compared the two task scores between the models.

Results

Our simulations showed that an active learning model with a label margin (LM) utility function outperformed both entropy (ENT) and most certain (MC) models in the discrimination task. An active learning model with a most certain utility function outperformed both entropy and label margin models at the generation task. (Fig 2B). These results support our hypothesis that different goals favor different sampling heuristics.

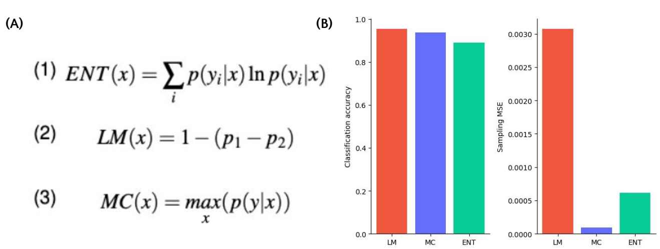

Figure 2: Simulated data highlight the differential success of sampling heuristics. (A) Utility functions (B) Left: The classification accuracy score of the LM model is higher than that of the MC model, indicating it is better in the discrimination task. Right: The MSE score of the MC model is lower than that of the LM model, indicating it is better in the generation task.

Following Markant et al. (2016), to assess what model each participant used, we evaluated the three sampling heuristics on the participant-reported estimates of channel probabilities for each sampled antenna. Consistent with previous work, the data suggested that various sampling strategies were used to complete the task. The model that explained the majority of participants' sampling choices was MC, that is, participants rated most samples as highly likely to be one of the three channels. In accordance with our hypothesis, there was a trend for participants in the discrimination task to use the LM model more often than those in the generation task (Fig 3A,B). We defined an information-seeking index $IS = MC - LM$ as the difference between the samples that were best fit by MC and those that were best fit by LM. A t-test comparing the IS of the two conditions revealed a difference (although not significant) between the two groups: $t = -1.81, p = 0.068$ (Fig 3C).

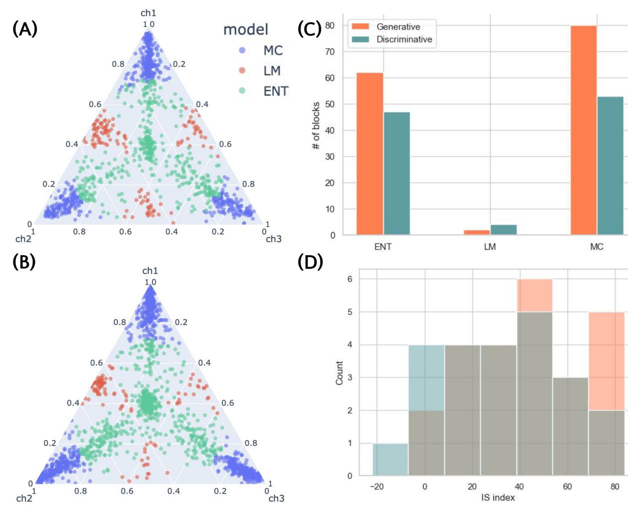

Figure 3: **Increased LM sampling in the discriminative condition.** Participant probability estimates plotted within the three-category simplex colored according to the model best fit to each estimate for participants in the discriminative condition (A) and the generative one (B); note relatively more red (LM) in the top. (B) Number of blocks fit to each one of the models. Notice higher LM in the discriminative condition than the generative one. (C) t discriminative task drives a higher use of LM

Discussion

We showed that all else equal, manipulating the test tended to shift people's behavior in an active learning clustering task. These findings highlight the importance of considering one's goals when evaluating sampling strategies. We intend to replicate these findings with a larger sample and a within-participants design.

We recognize several limitations of our design, in particular, our reliance on participants' self-report estimates. In addition, our results are not statistically significant and we intend to collect more data. Our participants exhibited a strong bias towards MC sampling, which is inconsistent with previous results (Markant et al., 2016).

Acknowledgments

This work was supported by CRCNS grant number 10013459

References

Gottlieb, J. (2018). Understanding active sampling strategies: Empirical approaches and implications for attention and decision research. *Cortex*, *102*, 150–160.

Jha, A., Ashwood, Z. C., & Pillow, J. W. (2022). Bayesian active learning for discrete latent variable models. *arXiv preprint arXiv:2202.13426*.

Markant, D. B., Settles, B., & Gureckis, T. M. (2016). Self-directed learning favors local, rather than global, uncertainty. *Cognitive Science*, *40*(1), 100–120.