Title
Dynamic Information Sampling via Rapid Sequential Storage and Recurrence

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 45(45)

Authors
Gao, Mengcun
Ralston, Robert
Sloutsky, Vladimir

Publication Date
2023

Peer reviewed

Dynamic Information Sampling via Rapid Sequential Storage and Recurrence

Mengcun Gao[1] **(gao.643@buckeyemail.osu.edu)**

Robert Ralston[1] **(ralston.123@osu.edu)**

Vladimir Sloutsky[1] **(sloutsky.1@osu.edu)**

[1]The Ohio State University Department of Psychology, 1835 Neil Avenue, Columbus, Ohio 43210

Abstract

When making category decisions, humans sample features following their dynamic informativeness. Attention optimization models successfully predict these categorization behaviors, but optimization is not the only solution. Alternatively, categorization can be viewed as a Reinforcement Learning (RL) task in which learners sample information based on its expected utility. However, RL models of information sampling have high computational load, even though human learners solve this problem on the millisecond timescale. Therefore, we propose ATHENA-RSS, a model that implements reward-based information search in a more computationally efficient way, via the rapid sequential storage of memories and recurrent retrieval. To test the model, we conducted an experiment where participants ($N = 99$) learned hierarchically structured categories by uncovering stimulus features. We then conducted a simulation study, where ATHENA-RSS successfully reproduced all search patterns exhibited by participants. We conclude that rapid sequential storage and recurrent memory retrieval were sufficient to achieve human-like information sampling in this task.

Keywords: Categorization; Reinforcement Learning; Attention Optimization

Introduction

In formal theories of categorization, human category learning has been explained by a mechanism of **attention optimization** (Kruschke, 1992; Love, Medin, & Gureckis, 2004; Nosofsky, 1986; Weichart, Galdo, Sloutsky, & Turner, 2022). For example, the Generalized Context Model (GCM) and the Attention Learning COVEring map model (ALCOVE) (Nosofsky, 1986; Kruschke, 1992) assumes that humans learn categories by learning to attend to stimulus features to optimize performance. However, GCM, ALCOVE, and many classical models which use these attention weights share one critical limitation: they lack the flexibility to solve problems where the importance of dimensions is dynamic.

In classical category-learning models, during one trial, a common set of attention weights is used for all exemplars and stimuli. However, a one-size-fits-all set of attention weights neglects the fact that the importance of stimulus features can vary based on within-trial context, where the identity of one dimension can impact the importance of others.

This limitation has been addressed by more recent models (Aha & Goldstone, 1992; Braunlich & Love, 2022; Kruschke, 2001; Nosofsky & Hu, 2022; Weichart et al., 2022). For example, models assuming stimulus-specific or region-specific attention weights address the problem by incorporating different sets of attention weights for each exemplar or region of psychological space where categories are embedded.

In addition, the Sampling Emergent Attention Model (SEA) views categorization tasks as Reinforcement Learning (RL) tasks where decisions about sampling are driven by the expected utility of sampling each feature (Braunlich & Love, 2022). SEA incorporates two components: a Bayesian concept learning component and an information-utility component that uses a forward-search process (preposterior analysis) to sample features based on their expected utility. Specifically, a full preposterior analysis considers the consequences of selecting each candidate feature several steps in the future. This procedure is computationally costly and seems to imply that participants in category learning experiments are doing extensive computations on the millisecond timescale to determine which feature to sample next (i.e., where to look).

In contrast, the Adaptive Attention Representation Model (AARM) showed that dynamic information sampling could be accomplished by attention optimization with a less-intensive, confirmatory search mechanism. According to AARM, instead of considering all possible sampling paths, learners sample features that are most likely to further support the most probable category based on the information currently available. As an attention optimization model, AARM achieves dynamic information search by allowing attention weights to update between trials and fluctuate within a trial (Weichart et al., 2022).

Although attention optimization can account for category learning, it is unlikely to be the only mechanism. As described above, SEA provides an alternative way to understand information sampling. Instead of optimizing their attention through supervised or unsupervised learning, learners may sample information based on its expected information gain using a reinforcement learning mechanism. This mechanism seems promising; however, are there simpler ways to achieve it without embracing the high computational demands?

Memory Storage in Categorization Models

Many prior category-learning models rely on a memory store to hold entire examples or summary representations of previous stimuli (Braunlich & Love, 2022; Kruschke, 1992; Love et al., 2004; Nosofsky, 1986; Weichart et al., 2022). The idea that episodic memory is important to categorization is reinforced by studies showing that categorization behavior is re-

In M. Goldwater, F. K. Anggoro, B. K. Hayes, & D. C. Ong (Eds.), *Proceedings of the 45th Annual Conference of the Cognitive Science Society*.

lated to recognition memory (Deng & Sloutsky, 2015, 2016; Jacoby, Wahlheim, & Coane, 2010; Unger & Sloutsky, 2023), as well as fMRI studies showing activation in the hippocampus during category learning (Davis, Love, & Preston, 2011).

One characteristic shared by essentially all memory-based category learning models is that memory is updated on the trial timescale (i.e., a single trace is added at the end of a learning trial). However, in contrast to the models, humans have been observed to store rapid sequential memories, recalling order information even when visual items are presented rapidly (Paivio & Csapo, 1971; Ghirlanda, Lind, & Enquist, 2017). Moreover, storing memories with low temporal resolution may discard important temporal structure; if individuals retain a record of the order in which features are sampled, as well as an estimate of the information gained by sampling each feature, this could provide a basis for information sampling via the same inference mechanisms already used in models of categorization. In other words, through memory alone, individuals could decide to sample features which were previously informative in similar situations. In the following section, we develop this reasoning into a full model of category learning.

ATHENA with Rapid Sequential Storage

In this section, we introduce AuToassociative and HEteroassociative Neural Attention with Rapid Sequential Storage (ATHENA-RSS), a model that accomplishes category learning and simple information sampling via instance-based memory. The name comes from the fact that ATHENA is a more general system for memory retrieval derived from the dot product attention layers used extensively in machine learning (Ralston & Sloutsky, 2022; Vaswani et al., 2017). Then, to apply ATHENA to the problem of information sampling, we further assume that individuals engage in *Rapid Sequential Storage* during category learning, acquiring many temporally-ordered memory traces during the same trial.

The inference mechanism in ATHENA can be stated simply. Say that we wish to infer unobserved values of a set of stimulus features x_{t+1} from a set of known features x_t. This can be obtained by the following expression.

$$Athena(x_t) = softmax(\beta x_t K^T)V = x_{t+1} \qquad (1)$$

Here, K and V represent matrices where each row corresponds to a memory trace and each column to a stimulus feature. K represents the features of an item which are being cued, while V represents the features of an item which are being retrieved. Furthermore, $\beta \geqslant 0$ is a constant related to the specificity of the stored memory traces. In this paper, our goal is to explain information sampling using memory alone. Thus, we will omit the use of attention weights, though they could be added to these equations.

Inferences in ATHENA are accomplished by comparing the current item (x_t) to stored items (rows of K) in memory (via the dot product). Previous work has shown that, under suitable conditions, the inference mechanism of ATHENA results in rational inferences (Ralston & Sloutsky, 2022). Specifically, if v_x represents the unobserved dimensions corresponding to the cue x_t, and x_t is assumed to be drawn from one of the items stored in memory with added angular noise following a von Mises distribution, then $x_{t+1} \approx E[v_x|x_t]$. Since x_{t+1} represents an expectation, it can be used to infer many aspects of stimuli. This will be important, as we use this mechanism to determine both the category label of an item and the expected value of sampling a dimension.

As suggested by their subscripts, ATHENA can be used recursively to generate predictions about temporal sequences when the dimension of x_t and x_{t+1} are the same. Consider the case where a participant has experienced a sequence of items, $s_1, s_2, \ldots, s_n$, many times, and has formed memory traces binding each item to its successor; i.e., an s_1, s_2 binding, an s_2, s_3 binding, and so on. If we represent the first item in each binding as a row in K, and the second item as a row in V, then Eq. 1 will provide the expectation of each feature for the item expected to occur after x_t, i.e., $x_{t+1} \approx E[s_{t+1}|s_t]$. Then, by recursively substituting $x_{t+1} \rightarrow x_t$, Eq. 1 can be used to infer s_{t+2}, s_{t+3} and so on. In brief, sequential binding allows ATHENA to recursively infer the state of a system in the future from 2-place bindings stored in memory.

For the current implementation, we use Eq. 1 for three inferences, two of which are recursive. Based on the information currently available to an individual, we infer a) the value of sampling each feature, b) the expected feature values after sampling each feature, and c) the category label. This is possible under Rapid Sequential Storage, which postulates that a memory trace is formed after sampling each new feature of a stimulus. We define a memory trace - i.e., a row of K and V - as a compound representation binding two items:

$$\begin{aligned} k_t &= [\text{stimulus at } t \quad | \quad \text{IS at } t] \\ v_t &= [\text{stimulus at t+1} \quad | \quad \text{IS at } t+1 \quad | \\ &\qquad | \quad \text{value obtained from } t \text{ to } t+1 \quad | \quad \text{label at } t] \end{aligned}$$

Here, IS represents the *information state*, denoting what information is available at the current time (see our implementation below), and the vertical bars indicate that these vectors are concatenated together. Unsampled features receive a value of 0 in the trace and do not contribute to the inference in Eq. 1, and the label is only nonzero during feedback.

One important choice to make when storing memory traces is whether to impose any constraints on the value obtained by sampling a feature. In our simulations below, we use an individual's assessment of the information gained by sampling a feature as the reward for sampling that feature. One possible constraint is the principle of premise monotonicity, which states that acquiring more information can never make an individual less certain about an outcome (Osherson, Smith, Wilkie, López, & Shafir, 1990). This constraint can be instantiated by clamping stored values at 0, so that information gain could never be negative. However, in empirical studies, human participants often violate monotonicity, suggesting that such constraints may be unwarranted (Voorspoels, Navarro, Perfors, Ransom, & Storms, 2015; Ransom, Perfors, & Navarro, 2016). Below, we report one simulation with the monotonicity assumption and one without.

We will now outline how ATHENA-RSS accomplishes categorization and information sampling. The decision mechanisms below are chosen to be simple in order to show that recurrent memory retrieval is sufficient for information sampling without a complex decision component.

Category Decisions ATHENA-RSS makes category decisions only considering information currently available to it. Let y_{t+1} represent the dimensions of x_{t+1} which are associated to the category label, and $Pr(C|x_t)$ represent a vector where each place corresponds to the probability that an item is in one of the categories. We can use Luce's choice rule to model decision making (Nosofsky, 1986):

$$Pr(C|x_t) = softmax(d_{cat} y_{t+1}) \tag{2}$$

Here, $d_{cat} \geqslant 0$ corresponds to a *response determinism* parameter, which determines how sensitive an individual is to small differences in evidence for each category.

Information Sampling The model determines which features to sample by assessing the expected *future* value of each feature using the process shown in Figure 1. For each feature which can be sampled, the model initializes an instance of Eq. 1 where the information state obtained by sampling the candidate feature is substituted into its current representation. In each instance, the predicted features and IS of x_{t+1} are fed back into Eq. 1, forming a recurrent architecture, which can predict the state of the system after sampling many features.

To decide what to sample, the model attempts to maximize the time-discounted value which can be obtained using a variant of the Bellman equation (Bellman, 1957). To determine $V(d)$, the value of sampling feature d, we use:

$$V(d) = E[v_0|x_0+d] + \sum_i \gamma^i E[v_i|x_t = E[x_i|x_0+d]] \tag{3}$$

$$= \hat{v}_0 + \sum_i \gamma^i \hat{v}_i \tag{4}$$

Here, $0 \leqslant \gamma \leqslant 1$ is a discounting parameter, $x_0 + d$ is the current stimulus with the information state obtained by sampling d, and v_i and $\hat{v}_i$ are the true and estimated immediate value at time i respectively. In other words, this equation says that the value of sampling a feature is given by the immediate value expected from sampling that feature, plus the discounted value of sampling further features given that the features revealed take their expected values. At each time step, $\hat{v}$ is estimated by the value dimension of x_{t+1}, which represents $E[v_t|x_t]$. Therefore, during the recurrent process, time-discounted reward is accumulated during each iteration.

The estimation of $\hat{v}$ with the value dimension of x_{t+1} at each time-step is imperfect, and will struggle when projected into the future if

$$\sum_{x_{t+1}} P(x_{t+2}|x_{t+1})P(x_{t+1}|x_t) \not\approx Athena(E[x_{t+1}|x_t]) \tag{5}$$

However, it has the benefit of being immediately calculable in the recurrent computation described above. Notice that, unlike the computationally intense tree search that is typically needed to determine future value (Braunlich & Love, 2022), the computation described here determines future value using one instantiation of Eq. 1 per feature.

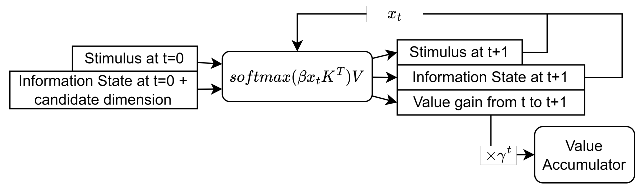

Figure 1: One instantiation of Equation 1 to estimate the expected value of a candidate dimension.

To determine which feature to ultimately sample, ATHENA-RSS uses the same decision rule as above. If $d_1,\ldots,d_k$ are the features under consideration, the probability of sampling each feature $P(D|x_t)$ is given by:

$$Pr(D|x_t) = softmax(d_{samp}[V(d_1),\ldots,V(d_k)]) \tag{6}$$

Analogously to d_{cat} above, $d_{samp} \geqslant 0$ is a response determinism parameter. Note that, when d_{samp} is low, this allows for more exploration, while a large d_{samp} will result in always sampling the feature with the largest expected value.

Search Termination Finally, the model needs a mechanism to decide when to terminate the search and give a response. To keep this simple, we used a threshold on the pre-response category evidence. Let e_c, the pre-decision evidence in favor of category c, be given by:

$$e_c = \frac{y_{t+1}^{(c)}}{\sum_i y_{t+1}^{(i)}} \tag{7}$$

Here, $y_{t+1}^{(c)}$ represents the component of y_{t+1} corresponding to category c. If τ is the model's threshold, then information search will be terminated if $e_c \geqslant \tau$ for at least one category c.

Current Study

To investigate whether ATHENA-RSS is able to predict human learners' dynamic information sampling, the current study used a category learning paradigm with hierarchical category structure (Blair, Watson, Walshe, & Maj, 2009; Meier & Blair, 2013). This structure was useful because it involves stimulus features whose informativeness changes depending on the identities of other sampled features (see Table 1 and our description below).

In addition, our category learning paradigm incorporated explicit feature sampling. At the beginning of each trial, all stimulus features were covered, and learners needed to uncover features throughout the course of a trial. This allowed us to examine which features were sampled by participants as well as the order in which features were sampled. By comparing observed data from human learners to model simulations, we show that ATHENA-RSS can produce the qualitative patterns present in human information sampling without using attention weights (as in AARM) or a computationally costly information search (as in SEA).

Experiment

Participants

Ninety-nine young adults (50 females, 48 males, one did not report gender; Age: M = 19.28, SD = 1.52) from the Ohio State University participated in the experiment for course credit. Prior to the study, all procedures were approved by the appropriate Institutional Review Board.

Methods

Stimuli Stimuli were alien creatures that were constructed according to a hierarchical category structure as in Blair et al. (2009) (see Table 1). All the stimuli could be first classified into two higher-level categories (Category A and B), and within each higher-level category, stimuli could be further classified into two lower-level categories (A1, B2, etc.). All stimuli had three features with two binary values. Among the three features, one feature was assigned as the Higher-level feature, and the other two were assigned as Lower-level features (Category A Relevant feature and Category B Relevant feature). The value of Higher-level feature determined the higher-level category membership of the stimulus and which of the two Lower-Level features was relevant for lower-level category judgment. The assigned role (higher level vs. lower level) of each feature was counterbalanced across participants to control for any prior bias towards certain features.

Table 1: Category Structure.

Higher-Level	Lower-Level	F1	F2	F3
A	A1	1	1	1
	A1	1	1	0
	A2	1	0	1
	A2	1	0	0
B	B1	0	1	1
	B1	0	0	1
	B2	0	1	0
	B2	0	0	0

Procedure In the experiment, participants were told that their job was to help alien friends get home. These alien friends lived on two different planets on which their homes were located in different places (e.g., trees or grass on the green planet). Then participants were trained on the lower level categories for 64 trials with an equal number of trials for each lower-level category. On each trial, all the features of the stimulus were covered by grey boxes. Participants were instructed to click the boxes to uncover features and respond when they knew the answer. Feedback was provided after they made decisions. During feedback, participants were presented with an entire stimulus that had all features uncovered and the correct category labels (both higher-level and lower-level category labels). After training, participants received 32 Phase 1 Testing trials that were visually the same as training trials. Finally, participants received Phase 2 Testing in which they were tested on the higher level categories for 32 trials. No feedback was provided during testing phases.

Behavioral Analyses Categorization accuracy in the last sixteen training trials ($M = 0.72$, $SD = 0.27$) was significantly above chance (0.25), $t(98) = 17.188$, $p < 0.001$, $d = 1.727$. Their categorization accuracy in Phase 1 Testing (lower-level categorization task; $M = 0.71$, $SD = 0.27$) was significantly above chance (0.25), $t(98) = 17.213$, $p < 0.001$, $d = 1.730$, and categorization accuracy in Phase 2 Testing (higher-level level categorization task; $M = 0.86$, $SD = 0.22$) was also above chance (0.5), $t(98) = 15.976$, $p < 0.001$, $d = 1.616$.

Moreover, to examine whether participants explored features based on their relevance, we conducted A 2 (True Higher-Level Category: Category A vs. Category B) × 3 (Feature Type: Higher-level vs. Category A Relevant vs. Category B Relevant) mixed ANOVA for Phase 1 and Phase 2 Testing trials. The results for Phase 1 Testing revealed a significant main effect of feature type, $F(2, 490) = 34.037$, $ps < 0.001$, $\eta^2 = 0.12$, and a significant interaction between feature type and true Higher-level category, $F(2, 490) = 58.618$, $ps < 0.001$, $\eta^2 = 0.19$. Furthermore, the results for Phase 2 Testing found only a significant main effect of feature type, $F(2, 490) = 354.039$, $ps < 0.001$, $\eta^2 = 0.58$. These findings indicated that participants were able to explore features following their dynamic importance. Specifically, during Phase 1 Testing, participants explored the Higher-level feature and Category A Relevant feature more than Category B Relevant feature when the current stimulus belonged to Category A, but explored the Higher-level feature and Category B Relevant feature when the current stimulus belonged to Category B. However, during Phase 2 Testing, because the only relevant feature in categorizing stimuli at a higher level was the Higher-level feature, participants showed reduced exploration of both Category A and Category B Relevant features. Together, our results found strong evidence that human learners were able to sample features based on their relevance.

Simulation Study

Our goal in using ATHENA-RSS is to show that a model with recurrent memory and rapid sequential storage produces the same sampling patterns as participants. To do this, we simulated participant behavior during the Training and Phase 1 Testing phases of the experiment, as optimal performance in these phases required dynamic information sampling.

Model Implementation To implement the model, we used a one-hot code for each feature of the stimulus, while the in-

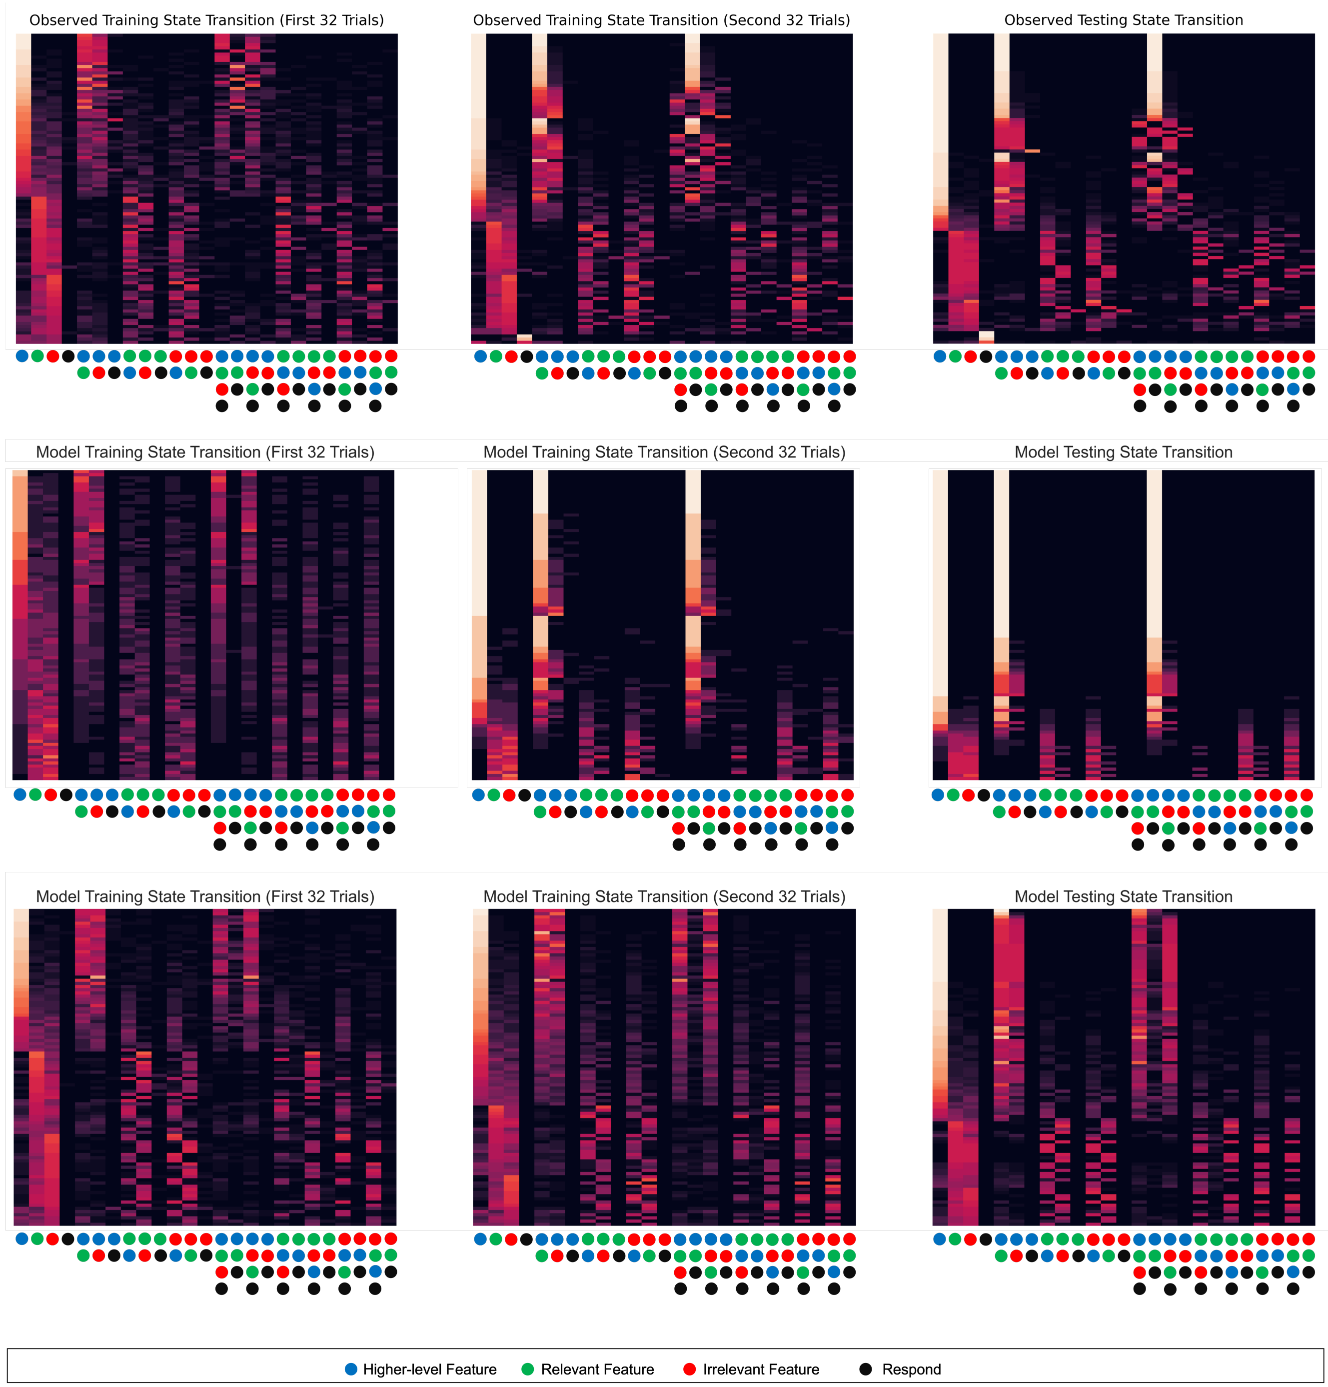

Figure 2: Observed (Top row) and simulated (Middle: Two-Level Simulation and Bottom: One-Level Simulation) proportion of trials in each information state during first-half of training (Left Column), second-half training (Middle Column), and phase 1 testing (Right Column). Within each sub-plot, a row represents one participant (top) or one simulation (middle and bottom), and brighter colors represent a higher proportion of trials. The first four columns represent the four possible choices at the beginning of a trial, and the following columns represent the possible two-step and three-step sampling patterns which can occur. ”High” (blue dots) choices represent choices for F1 (from Table 1), ”Relevant” (green dots) and ”Irrelevant” (red dots) choices are those for the remaining features which are relevant or irrelevant for categorization given F1. The order of dots represents the order of actions taken (i.e., feature sampling and responding). For instance, if a participant consistently sampled the High-level feature followed by the relevant feature and then responded, they would have bright values in the blue dot column (one choice), the blue-green column (two-choices), and the blue-green-black column (three choices).

.

formation state was represented by three one-hot features, denoting whether each feature of the stimulus was covered or uncovered. During a trial, the model engaged in information search as described in the sections above, and stored a memory trace every time a feature was sampled, with traces added to K and V between trials. When feedback was experienced, category labels were stored as part of participants' memory traces. Because many traces did not contain category labels, we used a weighted one-hot code when storing labeled traces. For example, the code for category one can be given as $[l, 0, 0, 0]$, where $l \geqslant 0$ is a free "label-determinism" parameter representing the relative salience of the label. Additionally, to allow for variable exploration, we used separate d_{samp} and τ parameters for the two halves of Training and the Testing phase of the experiment. Finally, the subjective value gained by uncovering a dimension was defined as the reduction in normalized entropy of the four vector-places denoting category labels (y_{t+1} above) after being normalized to a probability distribution. Thus, when a sampled feature increases the probability that the stimulus is in one category over the others, this action will be assigned a large value.

With this implementation, we simulated 100 participants using two sets of parameters. The Two-Level simulation (Middle of Figure 2) used the parameters $\beta = 2$, $\gamma = .1$, $\tau = .95, .6, .6$, $l = 10$, and $d_{samp} = 5, 20, 20$ during the two halves of training and testing respectively. In contrast, the One-Level simulation (Bottom of Figure 2) used the parameters $\beta = 2$, $\gamma = .4$, $\tau = .85, .85, .85$, $l = 10$, and $d_{samp} = 5, 5, 10$. Furthermore, the Two-Level simulation used the premise monotonicity assumption, while the One-Level simulation did not.

Results Our results can be seen in Figure 2. Two predominant patterns are seen in the Two-Level solution. The first is that many simulated participants show an optimal search pattern by the Testing Phase, choosing the feature which distinguishes higher-level categories first, followed by the correct Relevant feature, and then giving a response. This group appears to have learned that the feature distinguishing higher-level categories is useful because it determines which of the remaining features to sample. In other words, this group has mastered both levels of the hierarchical task, qualitatively matching the performance of participants who performed best. The existence of this group shows that memory recurrence and rapid sequential storage is sufficient to accomplish information sampling in this task.

The second pattern seen in this simulation is a collection of non-learners who sample either the Relevant or Irrelevant feature first. The size of this group diminishes through training, however some participants remain even during testing. This can occur because the model stores memory traces in a confirmatory manner, selecting the dimensions that will ultimately be stored according to its current assessment of the value of that dimension. Thus, a misleading stimulus order at the beginning of training, or unlucky early sampling choices can lock the model into a sub-optimal search pattern which some never escape from. This shows that the model sometimes falls into *information traps* which have empirically been observed in human participants, including in the data presented here (Rich & Gureckis, 2018; Blanco, Turner, & Sloutsky, 2023).

In the One-Level simulation, a similar pattern of non-learners can be observed. However, there are also a selection of qualitatively different learners. These individuals learn that the F1 dimension is useful for distinguishing between higher-level categories; however they do not develop a preference to sample the Relevant or Irrelevant features. This pattern can occur because non-monotonic information gain can fluctuate widely through training, leading to unstable value estimates which the model does not escape from.

Comparing these patterns to participant data in Figure 2, we observe that, with the exception of a small number of off-task participants who respond without uncovering any features during testing, all qualitative search patterns can be seen in either the Two-Level or the One-Level simulation. This shows that ATHENA-RSS reproduces the qualitative patterns seen in participants without attention weights or a computationally-intensive search algorithm.

General Discussion

The goal of the current study was to provide evidence that an instance-based memory model, ATHENA-RSS, is able to reproduce humans' information sampling patterns in a hierarchical category learning task. In our experiments, we found that human learners displayed different search patterns during learning. Some were able to follow an optimal search pattern by sampling features based on their dynamic information gain, while some were able to prioritize the Higher-level feature but got confused between two Lower-level features. Finally, non-learners tended to sample all the features without following a specific order.

ATHENA-RSS successfully reproduced the search patterns exhibited by human learners via recurrent memory and rapid sequential storage using two sets of parameters. It reproduced both the optimal path of human learners and some learning traps that learners could fall into. This suggests that, while attention optimization plays a large role in the literature, learners may use other mechanisms in addition to attention when sampling information. However, the success of both attention optimization models (i.e., AARM) and our model indicates that this paradigm cannot distinguish between these two learning mechanisms, and future investigations can explore when the two mechanisms may perform differently.

Finally, the model considered here has several limitations. In our investigations, we focused on the memory storage and retrieval component of ATHENA-RSS to distinguish it from alternatives. As a result, we omitted attentional learning and used a simplified response mechanism. To make the model more realistic, future research can implement it with more psychologically-plausible mechanisms to understand the predictions that ATHENA makes more clearly and how changes in model parameters affect these predictions.

Acknowledgments

This study was supported by National Institutes of Health Grant R01HD078545 to Vladimir Sloutsky.

References

Aha, D. W., & Goldstone, R. L. (1992). Concept learning and flexible weighting. In *Proceedings of the fourteenth annual conference of the cognitive science society* (pp. 534–539). Bloomington, IL: Lawrence Erlbaum Associates.

Bellman, R. (Ed.). (1957). *Dynamic programming*. Princeton, NJ, USA: Princeton University Press.

Blair, M. R., Watson, M. R., Walshe, R. C., & Maj, F. (2009). Extremely selective attention: Eye-tracking studies of the dynamic allocation of attention to stimulus features in categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *35*(5), 1196–1206.

Blanco, N. J., Turner, B. M., & Sloutsky, V. M. (2023). The benefits of immative cognitive control: How distributed attention guards against learning traps. *Journal of Experimental Child Psychology*, *226*.

Braunlich, K., & Love, B. C. (2022). Bidirectional influences of information sampling and concept learning. *Psychological Review*, *129*(2), 213–234.

Davis, T., Love, B. C., & Preston, A. R. (2011). Learning the exception to the rule: Model-based fmri reveals specialized representations for surprising category members. *Cerebral Cortex*, *22*(2), 260–273.

Deng, W., & Sloutsky, V. M. (2015). The development of categorization: Effects of classification and inference training on category representation. *Developmental Psychology*, *51*(3), 392–405.

Deng, W., & Sloutsky, V. M. (2016). Selective attention, diffused attention, and the development of categorization. *Cognitive Psychology*, *91*, 24–62.

Doll, B. B., Simon, D. A., & Daw, N. D. (2012). Waiting and weighting: Information sampling is a balance between efficiency and error-reduction. *Current Opinion in Neurobiology*, *22*(6), 1075–1081.

Gershman, S. J., & Uchida, N. (2019). Believing in dopamine. *Nature Reviews Neuroscience*, *20*(11), 703–714.

Ghirlanda, S., Lind, J., & Enquist, M. (2017). Memory for stimulus sequences: A divide between humans and other animals? *Royal Society Open Science*, *4*(161011), 1-12.

Jacoby, L. L., Wahlheim, C. N., & Coane, J. H. (2010). Test-enhanced learning of natural concepts: Effects on recognition memory, classification, and metacognition. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *36*(6), 1441–1451.

Kruschke, J. K. (1992). Alcove: An exemplar-based connectionist model of category learning. *Psychological Review*, *99*(1), 22–44.

Kruschke, J. K. (2001). Toward a unified model of attention in associative learning. *Journal of Mathematical Psychology*, *45*(6), 812–863.

Love, B. C., Medin, D. L., & Gureckis, T. M. (2004). Sustain: A network model of category learning. *Psychological Review*, *111*(2), 309–332.

Meier, K. M., & Blair, M. R. (2013). Waiting and weighting: Information sampling is a balance between efficiency and error-reduction. *Cognition*, *126*(2), 319–325.

Nosofsky, R. M. (1986). Attention, similarity, and the identification–categorization relationship. *Journal of Experimental Psychology: General*, *115*(1), 39–57.

Nosofsky, R. M., & Hu, M. (2022). Category structure and region-specific selective attention. *Memory & Cognition*.

Osherson, D. N., Smith, E. E., Wilkie, O., López, A., & Shafir, E. (1990). Category-based induction. *Psychological Review*, *97*(2), 185-200.

Paivio, A., & Csapo, K. (1971). Short-term sequential memory for pictures and words. *Psychonomic Science*, *24*(2), 50–51.

Ralston, R., & Sloutsky, V. M. (2022). Query-based memory approximates rational induction: Applications to infant statistical learning. In *Proceedings of the 44th annual conference of the cognitive science society.*

Ransom, K. J., Perfors, A., & Navarro, D. J. (2016). Leaping to conclusions: Why premise relevance affects argument strength. *Cognitive Science*, *40*, 1775-1796.

Rich, A. S., & Gureckis, T. M. (2018). The limits of learning: Exploration, generalization, and the development of learning traps. *Journal of Experimental Psychology: General*, *147*, 1553-1570.

Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, *275*, 1593–1599.

Sutton, R. S., & Barto, A. G. (Eds.). (1998). *Introduction to reinforcement learning*. Cambridge, MA: MIT Press.

Unger, L., & Sloutsky, V. M. (2023). Category learning is shaped by the multifaceted development of selective attention. *Journal of Experimental Child Psychology*, *226*, 105549.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A., … Polosukhin, I. (2017). Attention is all you need. *ArXiv*, *abs/1706.03762*.

Voorspoels, W., Navarro, D. J., Perfors, A., Ransom, K., & Storms, G. (2015). How do people learn from negative evidence? non-monotonioc generalizations and sampling assumptions in inductive reasoning. *Cognitive Psychology*, *81*, 1-25.

Weichart, E. R., Galdo, M., Sloutsky, V. M., & Turner, B. M. (2022). As within, so without, as above, so below: Common mechanisms can support between- and within-trial category learning dynamics. *Psychological Review*, *129*(5), 1104–1143.