

Approximate Inference through Active Sampling of Likelihoods Accounts for Hick's Law and Decision Confidence

Xiang Li (xiangl@nyu.edu)

Department of Psychology, New York University
New York, NY, 10003, USA

Luigi Acerbi (luigi.acerbi@unige.ch)

Department of Neuroscience, University of Geneva
1206 Geneva, Switzerland

Wei Ji Ma (weijima@nyu.edu)

Center for Neural Science and Department of Psychology, New York University
New York, NY, 10003, USA

Abstract

In N -alternative Bayesian categorization, computing exact likelihoods and posteriors might be hard for the brain. We propose an approximate inference framework with active sampling inspired by Bayesian optimization. While it is common in Bayesian models to assume that the agent makes noisy measurements of a state of the world, here we use a more general (and more abstract) starting point. We assume that the true (ideal-observer) likelihoods and posteriors of the categories are unknown to the agent. The agent sequentially makes noisy measurements of those likelihoods, one category at a time, thus refining their beliefs over the true likelihoods and their belief over the true posterior probabilities. To decide whether to make another measurement, the agent simulates the consequences of doing so for the latter belief. This framework accounts for two types of empirical findings. First, we find that the average number of measurements grows approximately logarithmically with N , reminiscent of Hick's law. Second, we account for a puzzling recent finding that decision confidence follows the difference between the two highest posteriors, rather than the highest posterior itself. Our framework provides a novel approach to explain human categorization by combining approximate inference with active sampling.

Keywords: categorization; approximate inference; active sampling; Hick's law; confidence

Introduction

The problem of categorization consists of choosing between two or more discrete alternatives given the available evidence, e.g. telling whether a young woman in an old picture is your mother or your aunt. Facing a similar problem, an optimal agent with access to exact likelihoods would simply apply Bayes' rule and choose according to their loss function (e.g., so as to maximize probability correct). However, the computation of exact likelihoods and posteriors might be infeasible for the human brain, as suggested by both theoretical considerations (e.g. Beck, Ma, Pitkow, Latham, & Pouget, 2012) and empirical results (e.g. Ashby, Waldron, Lee, & Berkman, 2001). How do humans overcome these limitations in categorization? Here we propose a partially normative Approximate Inference framework with Active Sampling (AIAS), inspired by Bayesian optimization (Jones, Schonlau, & Welch, 1998).

In this paper, we first introduce AIAS and discuss its relation with other frameworks. Then we show that AIAS accounts for Hick's law (Hick, 1952) regarding response time, and a recent finding that decision confidence follows the difference between the two highest posteriors of all categories (Li & Ma, 2019).

Framework

Generative Model

We consider an N -alternative categorization task under uncertainty. In this section, we take $N = 3$ for simplicity of exposition. In AIAS (Fig.1A), the true likelihoods L_1, L_2, L_3 and the true posterior vector $\mathbf{P} = (P_1, P_2, P_3)$ are unknown to the agent (Fig.1B,C). We assume that the agent acquires information by sequentially making noisy measurements of the likelihoods (Fig.1C), which we interpret as outputs of limited computations. At each time step, the agent can make a likelihood measurement l_i^k for a chosen category i , where k is the position in the sequence of measurements for category i .

Inference

Given a set of likelihood measurements and a prior over the likelihoods, the agent is able to compute their belief about the true likelihoods, i.e. the posteriors over the likelihoods $p(L_1|\mathbf{l}_1), p(L_2|\mathbf{l}_2), p(L_3|\mathbf{l}_3)$ (Fig.1D), where $\mathbf{l}_i = (l_i^1, l_i^2, \dots)$ is the vector of all likelihood measurements of category i . They then again apply Bayes' rule to compute their belief about the true posterior vector, i.e. the posterior over the posterior vector $p(\mathbf{P}|\mathbf{l})$ (Fig.1E), where $\mathbf{l} = \{\mathbf{l}_1, \mathbf{l}_2, \mathbf{l}_3\}$ is the set of all likelihood measurements of all categories. For computational tractability, in this paper we assume a factorized log-normal distribution for the prior over each likelihood $p(\log L_i) \sim N(\mu_L, \sigma_L^2)$, log-normal noise of the measurements of the likelihoods $p(\log l_i^k | L_i) \sim N(\log L_i, \sigma^2)$, and a uniform prior over categories.

Active Sampling, Termination, and Decision

Under the assumptions that measurements have a cost (time or computational resources), picking the category to measure randomly is unreasonable. For example, if a category seems improbable to be the correct answer, the agent should not waste more measurements on it. Analogously, because of limited resources, the agent does not make infinite measurements. At some point, they stop making measurements, se-



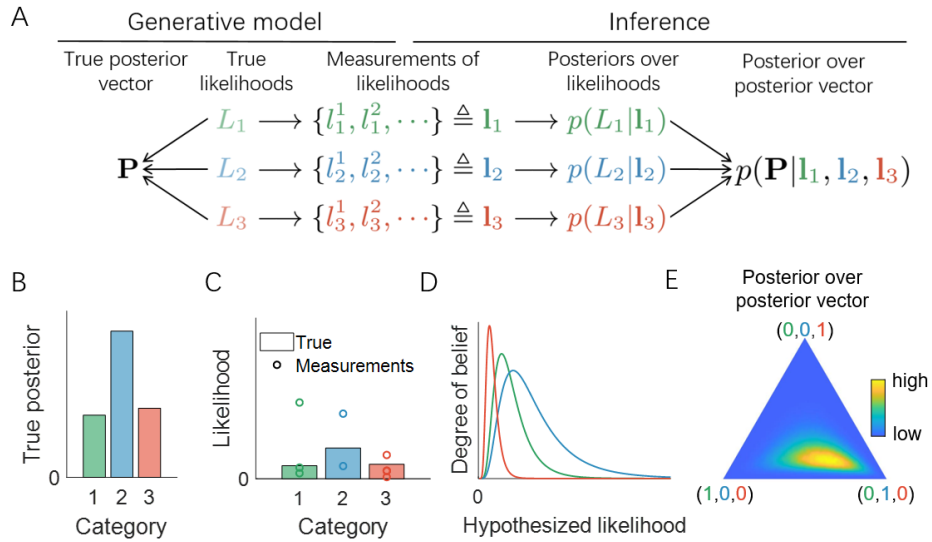


Figure 1: Generative model and inference process in AIAS. We consider a 3-alternative categorization problem for illustrative purposes. Colors correspond to the three categories. A) Flow chart of the generative model and inference process. The true likelihoods and posterior are unknown to the agent. B-C) Example true likelihoods, the corresponding true posterior vector, and the agent’s noisy measurements of the likelihoods. D) Posterior distributions over the true likelihoods of the categories. E) The posterior distribution over the posterior vector. Each point corresponds to a hypothesized posterior vector. The color of the point represents the degree of belief.

lect a category, and get rewarded for being correct in the categorization. The two key questions are, thus, the problems of *active sampling* (which category to make a measurement from?) and of *termination* (when to stop?). In reinforcement learning, the Bellman equation optimally solves this problem by taking into account all future steps, but the general case is computationally intractable. Instead, we consider here a myopic policy, which contemplates these questions only one step ahead.

Given the posterior over the posterior vector, $p(\mathbf{P}|\mathbf{I})$, the agent can calculate their belief over the maximal value in the posterior vector $p(\max\{P_1, P_2, P_3\}|\mathbf{I})$ and the variance of such distribution, Var . In the next step, the agent considers each category as a candidate for making another measurement. Separately for each category i , they simulate a new likelihood measurement based on current knowledge. The simulated measurement leads to a change of the corresponding posterior over the likelihood, a change of the posterior over the posterior vector, and thus an absolute change of the aforementioned Var , i.e. $|\Delta\text{Var}_i|$. If the maximum of all $|\Delta\text{Var}_i|$ ($i = 1, 2, 3$) is larger than a termination threshold ϵ , the agent chooses the argmax category to make an actual measurement of the likelihood. The agent then updates their posterior over the posterior vector and repeats the whole simulation process. Otherwise, the agent stops measuring and chooses the category i that maximizes $E[P_i|\mathbf{I}]$, i.e. the category that they believe maximizes probability correct.

Thompson sampling is a heuristic decision rule that approximates the intractable marginalization over beliefs by choos-

ing the action that maximizes the expected utility for a single sampled belief (Thompson, 1933). If we define the utility in AIAS as the absolute change of the variance after simulation $|\Delta\text{Var}|$, the above sampling process is indeed Thompson sampling. The intuition behind our approach is that categories with high true posteriors – which are the important ones for the decision – are likely to contribute more to the agent’s belief of the maximal value in the posterior vector, and thus simulations from these categories are more likely to lead to large $|\Delta\text{Var}|$.

Related Frameworks

Race model (Ratcliff, 1978): Race models assume that humans accumulate evidence for all alternatives simultaneously, until the evidence for one alternative is larger than a threshold. Such models have succeeded in explaining many behavioral results (e.g., response time) in decision-making under multiple alternatives. However, their process of evidence accumulation is mostly descriptive¹. In AIAS, the agent is also accumulating evidence (the measurements of the likelihoods), but the evidence is always incorporated in a Bayes-optimal way. Thus, AIAS is more normative than race models. Also, most race models do not have active sampling (but see Krajbich & Rangel, 2011).

Resource rationality (Griffiths, Lieder, & Goodman, 2015): If we assume the same generative model and inference, the resource rationality framework would propose to solve the problem by using Bellman equation. It performs optimally

¹Under some conditions, race models can be normative (Bitzer, Park, Blankenburg, & Kiebel, 2014), but this is not general.

(under imprecise likelihoods), but is still computationally intractable. Our approach provides a more computationally realistic inference method as a solution by a combination of the myopic step, heuristic termination rule, and Thompson sampling. This is also the reason why we describe AIAS as “partially normative”.

Information sampling (Oaksford & Chater, 1994): Information sampling is a general idea that humans selectively sample from the environment to maximize information gain. AIAS is consistent with and extends information sampling, by considering the brain’s computations themselves as a means of information gathering.

Bayesian optimization (Jones et al., 1998): Bayesian optimization is a machine learning technique to optimize costly and possibly noisy functions by building a posterior over functions and actively choosing which point to sample next according to an *acquisition function* that balances exploration and exploitation. Our approach is similar to Bayesian optimization on a discrete space of (log) likelihoods, with a specific acquisition function and termination rule, and a non-trivial correlation between choices induced by the normalization step going from likelihoods to posterior.

Results

Hick’s Law

We now investigate the behavioral predictions of AIAS. Response time is perhaps the most commonly probed behavioral metric. In a multiple-alternative choice task, Hick’s law states that response time increases logarithmically with the number of alternatives (Hick, 1952). In AIAS, we define as response time the total number of measurements for all categories. Thus, to reproduce Hick’s law, the number of measurements should increase logarithmically with the number of categories.

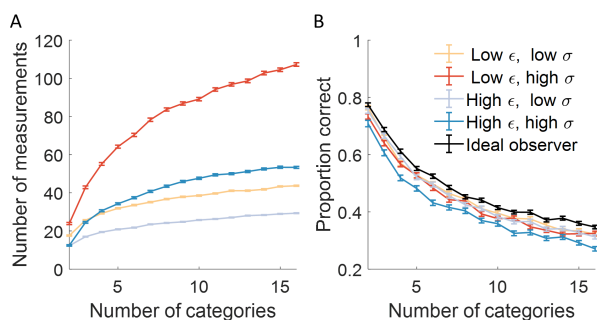


Figure 2: AIAS replicates Hick’s law. Error bars represent ± 1 s.e.m. A) The number of measurements increases logarithmically with the number of categories, for different combinations of termination threshold ϵ and measurement noise σ . B) Performance remains close to the ideal observer in categorization.

In simulation, we vary the noise of likelihood measurements σ and the termination threshold ϵ . Higher noise or lower termi-

nation threshold increase the number of total measurements (Fig.2A). Within a reasonable region of parameter space, multiple parameter pairs replicate the log form of Hick’s law and produce performance close to the ideal observer, who knows the true posteriors (Fig.2B).

Decision Confidence

In categorization, decision confidence provides another behavioral metric to investigate the underlying process. A leading hypothesis is that confidence reflects the subjective belief of the probability correct of a decision (Kepecs & Mainen, 2012). However, a recent study has found evidence contradicting this (Li & Ma, 2019). In Li & Ma’s experiment (Fig.3A), spatial configurations of three colored dot clouds are displayed on the screen. A black target dot is randomly generated from one of the three clouds. Subjects need to judge which color cloud the target dot comes from and report their confidence. There are four different color cloud configurations (Fig.3B top row). Li & Ma considered three models: the leading hypothesis that confidence represents the highest posterior of all categories (Max model), that confidence represents the entropy of the posterior distribution (Entropy model), and that confidence represents the difference between the two highest posteriors (Difference Model). Surprisingly, they found that the Difference model fitted their data best (Fig.3C, top row) and well (Fig.3B, middle row).

The win of the Difference model is curious because it conflicts with intuitive notions of confidence. AIAS offers a different perspective on the same findings. Since in AIAS, the posterior itself is unknown, a natural construct of confidence is the probability that the chosen category *has the highest true posterior probability* – which is different from the probability that the choice is correct. Because of sampling noise, the agent might accidentally pick the category with the second-highest posterior, but this would be associated with lower confidence. Intuitively, this form of confidence will depend on the difference between the two highest posteriors. To test this more thoroughly, we generated data from AIAS with parameters that reproduce features of the human data. We then fitted the Li & Ma models to these model-generated data. The generated data resemble human data (Fig.3B bottom row) and the model ranking is the same on generated data as on human data, with the Difference model being the best-fitting model (Fig.3C bottom row). This result suggests that the Difference model is a high-level description of an underlying process that is described by the AIAS, and it offers a new view of confidence.

Discussion

In this work, we develop a partially normative approximate inference framework with active sampling to address the problem of categorization when the agent does not have access to exact likelihoods. In AIAS, the agent selectively makes measurements of the likelihoods to refine beliefs about the posterior vector. We found that AIAS can account for two rather different types of regularities in human data, namely Hick’s law

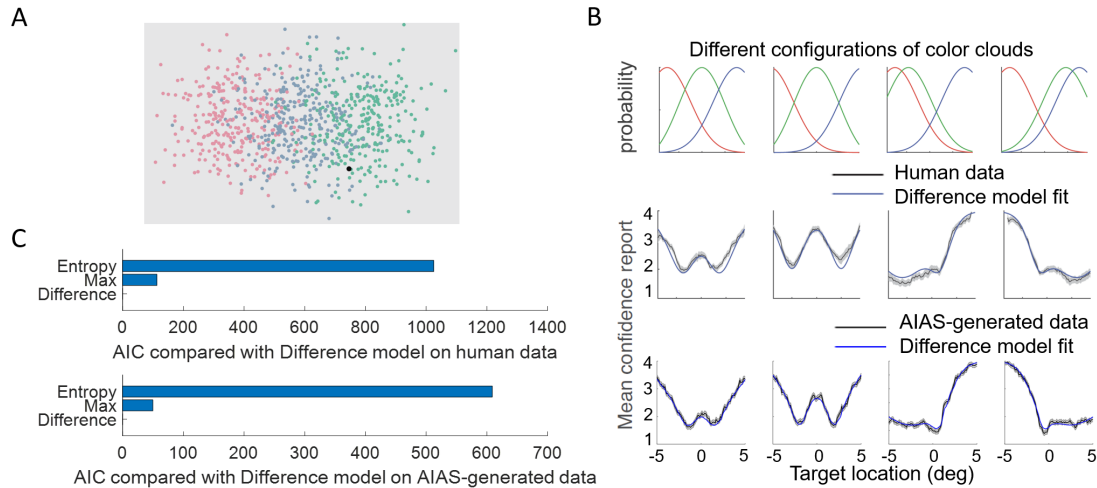


Figure 3: AIAS accounts for confidence reports in three-alternative categorization. A) Li and Ma (2019)'s experiment paradigm. Subjects see the configurations of three colored dot clouds. A black target dot is generated randomly from one of three color clouds. The subjects judge which color cloud the target dot comes from and report their confidence. B) Top row: four different configurations of color clouds. Middle row: human mean confidence reports (black) as a function of the target dot's position, with the fitting results of the Difference model (blue). Shaded areas represent ± 1 s.e.m. Bottom row: as middle row, but for AIAS-generated data. C) Model comparison results (AIC): the Difference model outperforms the Max model and the Entropy model for both human confidence data (top row) and AIAS-generated confidence data (bottom row).

for response times and confidence ratings in multi-alternative categorization.

These results form a basis for further testing AIAS. In particular, would it be possible to more directly interrogate the active sampling process? Eye movements might provide a window into the agent's thought processes, potentially signaling which category the agent is actively sampling from. For example, fixation times are reported to correlate with the final choice in a multiple-alternative choice task (Krajbich & Rangel, 2011), thus providing a potential empirical correlate of AIAS's measurement process.

Acknowledgments

Funding for this project was provided by NIH grants R01 EY020958-09 and R01 EY026927-02.

References

- Ashby, F. G., Waldron, E. M., Lee, W. W., & Berkman, A. (2001). Suboptimality in human categorization and identification. *Journal of Experimental Psychology: General*, *130*(1), 77–96.
- Beck, J. M., Ma, W. J., Pitkow, X., Latham, P. E., & Pouget, A. (2012). Not noisy, just wrong: The role of suboptimal inference in behavioral variability. *Neuron*, *74*, 30–39.
- Bitzer, S., Park, H., Blankenburg, F., & Kiebel, S. J. (2014). Perceptual decision making: drift-diffusion model is equivalent to a bayesian model. *Frontiers in human neuroscience*, *8*, 102–118.
- Griffiths, T. L., Lieder, F., & Goodman, N. D. (2015). Rational use of cognitive resources: Levels of analysis between the computational and the algorithmic. *Topics in cognitive science*, *7*(2), 217–229.
- Hick, W. E. (1952). On the rate of gain of information. *Quarterly Journal of experimental psychology*, *4*(1), 11–26.
- Jones, D. R., Schonlau, M., & Welch, W. J. (1998). Efficient global optimization of expensive black-box functions. *Journal of Global Optimization*, *13*(4), 455–492.
- Kepecs, A., & Mainen, Z. F. (2012). A computational framework for the study of confidence in humans and animals. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *367*(1594), 1322–1337.
- Krajbich, I., & Rangel, A. (2011). Multialternative drift-diffusion model predicts the relationship between visual fixations and choice in value-based decisions. *Proceedings of the National Academy of Sciences*, *108*(33), 13852–13857.
- Li, H.-H., & Ma, W. J. (2019). Confidence reports in decision-making with multiple alternatives violate the bayesian confidence hypothesis. *bioRxiv*, 583963–583400.
- Oaksford, M., & Chater, N. (1994). A rational analysis of the selection task as optimal data selection. *Psychological Review*, *101*(4), 608–631.
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological review*, *85*(2), 59–108.
- Thompson, W. R. (1933). On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, *25*(3/4), 285–294.