

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

People seek easily interpretable information

Permalink

<https://escholarship.org/uc/item/5sm2b484>

Journal

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

Authors

Cheyette, Samuel J.

Callaway, Frederick

Bramley, Neil R.

et al.

Publication Date

2023

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

People seek easily interpretable information

Samuel Cheyette
cheyette@mit.edu

Frederick Callaway
fredcallaway@princeton.edu

Neil Bramley
neil.bramley@ed.ac.uk

Jonathan Nelson
j.d.nelson@surrey.ac.uk

Joshua Tenenbaum
jbt@mit.edu

Abstract

Research in psychology and artificial intelligence has sought to ground information-seeking behavior in rational terms, typically assuming that people or agents prefer more informative data over less informative data. While this seems reasonable on its surface, it assumes that informativeness is only a property of the data, rather than a joint property of the data and a (potentially bounded) learner. That is, to the extent that it is hard to draw the right inferences from data that are theoretically “high information,” the data will not actually be highly informative to the learner. Here, we investigate active learning in humans using the code-breaking game Mastermind, which requires deductive reasoning from evidence. We find that people make queries that are less informative than random guesses, challenging standard rational or resource-rational accounts of information-seeking. We then show that people make queries are *informative to them* assuming they have a bounded capacity to draw inferences. We also find that participants prefer queries that provide easily-interpretable information over queries that provide more information but are less interpretable. Our results suggest that people are aware of their own cognitive limitations and seek information that they can use.

Introduction

People learn about the world through observation, experimentation, and by asking questions. There is a tradition in Psychology and Artificial Intelligence to understand hypothesis-testing and question-asking in the framework of Optimal Experimental Design (OED) — performing experiments or asking questions that provide maximal information with respect to a set of hypotheses (Horwich, 1982; Oaksford & Chater, 1994). Models based on Bayesian principles of experimental design have had some success in explaining human behavior, especially with respect to other possible modes of inquiry such as falsificationism (Coenen et al., 2019; Popper, 1963).

For instance, Oaksford and Chater (1994)’s account of the Wason Selection Task exemplifies the application of OED principles in explaining a ubiquitous and seemingly irrational aspect of human inquiry: a bias towards *confirmatory* rather than *dis-confirmatory* evidence. In the Wason Selection Task, participants are given four cards, each with a number on one side and a letter on the other, and a rule of the form “if p , then q .” For example, “if there is a vowel on one side (p), then there is an even number on the other side (q).” The four cards might show an A (p card), a K (not- p card), a 2 (q card) and a 7 (not- q card). The logically valid choices are A (p) and 7 (not- q), but people tend to pick A (p) and 2 (q). Oaksford and Chater (1994) demonstrate using OED principles that this type of confirmation bias is actually rational if people assume that p and q are rare — if so, q is a more informative choice than not- q .

However, recent work has shown that people select data and ask questions that are significantly less informative than

an optimal agent (Coenen et al., 2019). For instance, Rothe et al. (2016) ran an experiment involving a game similar to Battleship, where participants saw partially-revealed boards filled with rectangular tiles of varying lengths and colors, and had to ask questions that had single-word answers to reveal the hidden portion of the board. For instance, participants might ask “is the red ship vertical?” They found that participants’ questions were significantly less informative than information-maximizing queries. Other work has demonstrated that human inquiry tends to be myopic, in that people tend to consider only one or a handful of hypotheses (Bramley et al., 2017; Gregg & Simon, 1967; Klayman & Ha, 1989; Markant et al., 2016); and they tend to pick questions in a “greedy” fashion rather than choosing questions that are informative with future questions in mind (Bramley et al., 2015; Meder et al., 2019).

In this paper, we consider the hypothesis that limits on people’s ability to represent and process information determines how people actively learn about the world and what queries they make. Specifically, constraints on learning may bias people toward asking questions that have simpler answers than would seem optimal from a globally information-maximizing perspective — and yet, this may in fact be information-maximizing when accounting for such constraints. We evaluate this hypothesis by testing active learning in humans in a setting where questions return feedback of varying complexity, namely the code-breaking game Mastermind, in which a player attempts to guess a code comprised of colors or digits, and receives feedback about how close each guess was to the true code.



Figure 1: An example Mastermind game. The guesses are shown on the left; feedback is given in the green and yellow squares. The green square tells how many digits of the guess are in the code at the correct position. The yellow square tells how many digits of the guess are in the code but at the wrong position.

We first present a model of Mastermind to capture how an agent with limited working memory might rationally play and to help formalize a notion of processing complexity. Specifically, the model compresses the feedback received during a game into semantic statements about which codes are still possible given the queries and feedback in the game, such as “codes containing the digit 1” or “codes with the digit 1 at the second or third position”. The model makes different predictions about rational behavior depending on how long these expressions are allowed to be: to the extent that these statements are required to be short, information-maximizing guesses (relative to the agent) will be those that will likely provide highly compressible feedback; to the extent that these expressions are allowed to be long, guesses will be *truly* information-maximizing.

We then perform three experiments to test how humans play, using this model to guide experimental design and analysis. In Experiment 1, we first simply have participants play Mastermind and examine how informative their guesses are overall. We find that participants’ guesses are remarkably uninformative — less informative even than guesses sampled at random. In Experiment 2, we use the model to partially solve games optimized at different levels of processing complexity. We find that participants have a much easier time solving games — requiring both fewer guesses and less time — when games are optimized to lower levels of processing complexity. In Experiment 3, we use a forced-choice paradigm to test whether participants actually *prefer* lower-complexity (and lower-information) queries to higher-complexity (and higher-information) queries. We find that participants initially show little preference, but after playing several games strongly prefer lower-complexity queries over information-maximizing ones.

Mastermind

The goal of Mastermind is to guess a code consisting of four colors or digits (here we use digits)¹ by repeatedly making queries and receiving feedback. Each time a player makes a guess they receive two pieces of feedback: how many digits of their guess are in the code at the correct position and how many are in the code but in the wrong position. So, if a player guesses 1233 but the true code is 1344, they will learn that one digit of their code is in the code at the right position (in this case the 1) and one digit is in the code but at the wrong place (in this case the 3). They are not told which digits were correct or partially correct — only how many of each.

While Mastermind is often called a “deductive” game and is cognitively demanding to play, it can be solved perfectly by a simple computer program without using any complex reasoning. This is done by maintaining a list of all possible codes and ruling out codes that are incompatible with previous guesses. Given the limitations of human memory, however, it is unlikely that this is how any person solves the

game. In order to account for people’s limited memory capacity, we modeled game-play in a “Language of Thought” (LoT) that generates short semantic expressions to determine which codes are ruled out (or still valid) after each guess.

Bounded LoT model

The LoT model generates λ expressions consisting of first-order logic statements that take as input a potential code and a potential guess and return a truth value, i.e. $\lambda C, G \rightarrow Bool$. The truth value indicates whether a code is still valid given the guess. For instance, if the model guesses “2222” and learns that one digit is correct, the only remaining codes are those that have exactly one “2”, which can be expressed in the model as: $\lambda C, G \rightarrow equals(count(C, 2), 1)$. The expansion rules to generate valid expressions are essentially those of a context free grammar, though with the addition of bound-variables for quantifiers. The primitive operations we used for generating expressions under the model are shown below. While there is always some flexibility in choosing primitive operations, we note that we selected these primitives without reference to the data, using a fairly minimal set similar to those used to model other tasks (e.g. Goodman et al., 2008; Piantadosi et al., 2016).

Codes

$Code \rightarrow C$
 $Code \rightarrow G$

Booleans

$Bool \rightarrow equals(Code, Code)$
 $Bool \rightarrow equals(Val, Val)$
 $Bool \rightarrow equals(Num, Num)$
 $Bool \rightarrow gt(Num, Num)$
 $Bool \rightarrow and(Bool, Bool)$
 $Bool \rightarrow or(Bool, Bool)$
 $Bool \rightarrow xor(Bool, Bool)$
 $Bool \rightarrow not(Bool)$
 $Bool \rightarrow Quant(Bool)$

Values

$Val \rightarrow Code[Index]$

Indexes

$Index \rightarrow 1 \dots L$

Numbers

$Num \rightarrow 1 \dots N$
 $Num \rightarrow count(Code, Val)$

Quantifiers

$Quant \rightarrow \exists x \in Code$
 $Quant \rightarrow \forall x \in Code$

We can use the model to measure the complexity of representing a certain piece of information — i.e., the description length of an expression in the model². In many cases, the shortest expression to exactly represent the remaining valid codes are quite long. Crucially, however, there are other expressions that are shorter and less precise that capture *partial* information about which codes are still valid. Assuming that there is a cost to representing long semantic expressions — or that it is even impossible to do so beyond some limit — the game is still solvable but will take longer. Figure 2 shows how the model’s performance changes assuming it is only allowed to represent information about which codes are valid up to some bounded length. As the model is allowed longer expressions, it rules out more codes more quickly (blue lines)

¹In the original game of Mastermind, there are six allowable digits (or colors) in each of four slots. We only allowed four digits (1-4) in the four slots here.

²What we call the “description length” in this paper is actually the negative log probability of an expression, i.e. the number of bits it takes to represent under the model.




Figure 2: The number of codes remaining (y-axis) as a function of the number of guesses made (x-axis) in the model, grouped by the allowable description length. As the model is allowed greater expressivity (higher description-length bound), it makes more precise deductions and rules out more codes.

and as it is allowed shorter expressions it rules them out much less quickly (yellow lines).

As an illustrative example, consider the very first guess. The information-maximizing first guess happens to be “1123,” (or any equivalent code with two elements of one kind and two singles) from which 88% of codes will be ruled out afterward on average. The *least* informative first guess is “1111,” from which only 67% of codes can be ruled out after receiving feedback on average. However, the feedback received from guessing “1123” is often complicated to process, since multiple digits are involved. On the other hand, “1111” will only ever give simple feedback, returning how many 1’s are in the code. It turns out that the relative informativity of the two guesses is swapped assuming a low processing complexity bound. For a complexity bound of 6, for instance, “1123” only rules out 14% of codes on average, whereas “1111” rules out 44%.

Experiment 1

We first set out to simply test how informative participants’ queries are in the game Mastermind without any experimental manipulation and relate them to a model that picks maximally-informative queries. Based on previous studies of human performance in Mastermind (Schulz et al., 2019), as well as in other similar settings (Rothe et al., 2017), we expect that participants will pick queries that are not maximally informative. To the extent that a standard bounded-rational account of active learning might explain the data, however, participants’ queries should be more informative than random but less informative than optimal. If participants’ queries are *less* informative than random, that would suggest that participants may be using a strategy to *avoid* high-information queries, perhaps because they also tend to be complex.

Method

We recruited 40 adult participants from the online platform Prolific to play 10 rounds of the game Mastermind. Players

had up to 15 guesses total, otherwise they lost. They were given a base pay of \$2.50 with the possibility of earning a maximum of \$0.50 extra per game depending on their performance: each additional guess they took lowered their total bonus by 5 cents. If they ever guessed the correct code (even after 10 guesses) they received an additional bonus of 5 cents. All experiments were created using the PsiTurk framework (Gureckis et al., 2016).

Results

We removed participants who solved fewer than 5 games ($n = 9$), leaving 31 participants. The remaining participant pool eventually guessed the correct code in 83% of the games overall, so 17% of the time the game ended after 15 guesses without them guessing the correct code. In the games where participants eventually guessed the correct code, they took 7.6 guesses on average ($SD = 3.4$). This performance is remarkably poor compared with an optimal model, which takes on average 3.8 guesses and at most 5 for any given game.

An interesting question is whether participants were making uninformative queries, and thus taking a long time to narrow down the space of possible codes; or whether they were simply ignoring or not fully processing the information available to them. These are, of course, not mutually exclusive possibilities. If participants were fully processing the available information but asking uninformative questions, we should expect that once the code was fully determined by the queries and feedback — meaning, the game *could* be solved by a perfect-reasoning agent at that point — they should get the correct answer on the next turn. Of the 7.6 guesses it took participants to solve a game on average (provided that they ever solved it), it took them 4.5 guesses before the code was perfectly determined and 3.1 guesses after the code was perfectly determined to guess the correct answer. So, participants took on average 2.1 extra guesses on average once they had determined the true code. This implies that participants were likely both making uninformative queries *and* not fully processing all of the available information. The inefficiency remained even after removing redundant guesses, which only slightly lowered the mean guesses required to solve a game to 7.4.

To better address this question, we can compare the performance of two Mastermind-playing models against humans: an agent that picks a query to rule out as many codes as possible; and an agent that picks a random query until the code is perfectly determined, and then picks the correct code. Figure 3a shows the cumulative probability of having won the game (y-axis) as a function of the number of guesses made (x-axis) for people and both models. People are worse not only than the optimal model, but also significantly worse than the *random* model, which takes 4.9 guesses on average, or 2.5 fewer than participants did. More striking, however, is Figure 3b, which shows how many codes are remaining (y-axis) as a function of the number of guesses made (x-axis). This illustrates that participants’ queries are less informative even than random queries, since they rule out fewer codes af-




Figure 3: Results from Experiment 1. (a) The cumulative probability of guessing the correct code (y-axis) as a function of the number of guesses attempted (x-axis) for an information-maximizing agent (red), randomly-guessing agent (blue) and humans (black). (b) The number of codes remaining (y-axis) as a function of the number of guesses attempted (x-axis) for an information-maximizing agent (red), randomly-guessing agent (blue) and humans (black). (c) A histogram of the difference between the expected information gain of humans’ queries versus random queries across all guesses where there was at least some information to gain.

ter each guess than a purely random query would. Figure 3c shows the expected information gain of participants’ queries against purely random queries when there are at least some codes remaining. The EIG of participants’ queries and random queries differed significantly from 0 ($\mu = -0.25$ bits, $t_{1,1397} = -14.7$, $p < 0.001$).

Participants’ inefficiency was not merely about being forgetful or repeating themselves, as their queries were significantly less informative than random queries on each of their first four choices; separate t-tests revealed significant differences at each point ($ts < -3$, $ps < 0.01$). There was, however, an effect of time such that when participants spent more time thinking about their guess, the difference between the expected information gain of queries and a random query narrowed. As illustrated in Figure 4, however, even when participants spent a significant amount of time thinking (over 30 seconds), their queries never became more informative than random on average.

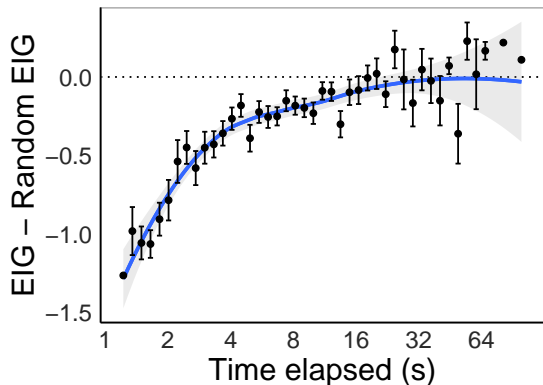


Figure 4: Data from experiment 1. The expected information gain of participants’ queries relative to a random query as a function of the amount of time they spent before guessing.

Experiment 2

Experiment 1 demonstrates that people make uninformative queries in Mastermind. However, this does not mean that

they would necessarily benefit from more informative data. If people have a bounded capacity for reasoning — and are unable to make all the available deductions when they are complex — then the most informative guesses may lead to feedback that is difficult to reason about. We can test this by partially solving games for participants where queries have been chosen either to be maximally informative or to be less informative but more interpretable. Our prediction is that participants will struggle to finish games that have been solved with the maximum-EIG guesses relative to games with less informative but lower-complexity guesses.

Method

We used the bounded-LoT model to play games of 4-digit, 4-slot Mastermind in one of two ways: either picking the max-EIG guess subject to a relatively low description length bound (length 9) or picking the max-EIG guess overall. Participants were then presented with partially-solved games from either one of these two conditions, with anywhere between 1 and 8 guesses already made by the model (along with feedback). Participants’ job was to finish the game from where the model left off. We recruited 35 participants from the online platform Prolific, each of whom played 10 games, five of which came from the bounded-DL model and five of which came from the max-EIG model. Participants were paid \$2.50.

Results

We removed participants who solved fewer than half of the games ($n = 5$), leaving 30 participants. In games that were partially solved using the max-EIG policy, there were fewer codes remaining than in games that were partially solved using the bounded-DL model (13.8 codes versus 26.5 codes on average). If processing complexity did not matter, then we would expect participants to finish the game faster in the games that were partially solved using the max-EIG policy. However, as Figure 5a shows, the opposite is true: the number of guesses required to solve the game in the lower-complexity, lower-information games was less ($\mu = 4.3$, $\sigma = 3.9$) than in the max-EIG games ($\mu = 6.7$, $\sigma = 4.7$). We ran

a regression predicting the number of guesses it took to solve the game from the number of codes remaining at the beginning and the policy type, with random subject intercepts and slopes. We found a significant effect of the number of initial codes remaining ($B_{\text{start}} = 0.04$; $t_{2,298} = 3.35$; $p < 0.001$); however, this effect was dominated by the policy type: games partially solved with the max-EIG policy required several more guesses to finish than games solved with the low-complexity policy ($B_{\text{max}} = 3.35$; $t_{2,298} = 4.88$; $p < 0.001$). Not only did participants require fewer guesses before reaching the correct solution in the low-complexity condition, but they were quicker as well, finishing 35 seconds faster on average (85s vs 120s).

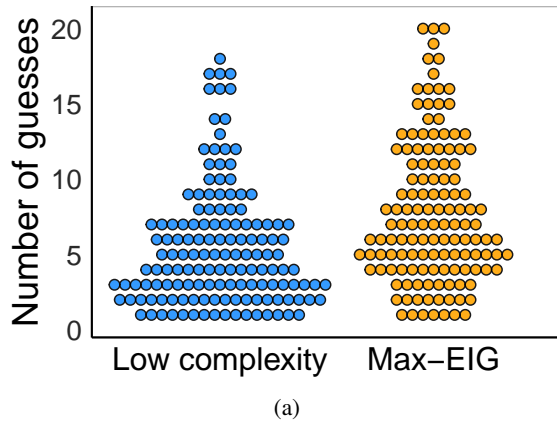


Figure 5: Results from Experiment 2. The number of guesses taken to find a solution in the low-complexity and max-EIG conditions. Participants took more queries on average to finish games that were partially solved using the max-EIG queries.

Experiment 3

Experiment 2 shows that people struggle to process information-maximizing guesses in Mastermind, and are able to make much more effective use of guesses that provide feedback that is easy to interpret. However, it is still not clear that people are deliberately *choosing* less informative but lower-complexity guesses when they play Mastermind. That is, we have not demonstrated that people are using a “bounded-rational” strategy in any sense, just that they struggle to process certain information and do not pick high-information queries. In Experiment 3, we use a forced-choice paradigm to evaluate whether people will choose low-information, low-complexity guesses over high-information, high-complexity guesses — and also, if so, whether this is to their benefit.

Method

In this experiment, participants played a modified version of Mastermind with two phases. In phase one, participants were repeatedly given two codes to select between as a guess. This portion of the game continued until the true code was perfectly determined, after which phase two began. In phase two, participants tried to guess the code (no longer as a forced-choice) but without any feedback. The codes in the forced-choice portion (phase one) were designed such that one

was maximally-informative and the other was maximally-informative subject to a bounded description length (length 9). Additionally, the codes presented as options to the participant in phase one were never the true code. We recruited 38 participants from Prolific, each of whom played 10 games. Participants were paid \$2.50.

Results

Our main question of interest here is how participants behaved in the forced-choice portion of the game — specifically, whether participants chose the guess that maximized expected information gain or whether they chose the lower-complexity guess that maximized information gain at a bounded description length. Figure 6a shows the probability that participants chose the low-complexity guess as a function of trial (how many games they had played). Participants were initially near chance in picking between the two options, but came to have a strong preference for the lower-complexity guess as they had played more games. We ran a logistic regression with random subject effects to predict the probability that participants chose the low-complexity guess as a function of the trial³, and found that the intercept was positive but not significantly different than 0 ($B_0 = 0.19$; $z = 1.05$; $p = 0.29$); however, there was a significant effect of trial number ($B_t = 0.09$; $z = 3.86$; $p < 0.001$). This regression thus predicts that participants start out picking the low-complexity option about 55% of the time but are picking the low-complexity option about 74% of the time by the final game. This is notable because it implies that participants are *learning* a strategy for picking lower-complexity codes.

We next looked at how participants’ behavior in the forced-choice portion of the task predicted their ability to determine the true code in the second phase of the task, which was free response and without feedback. Figure 6b shows the probability that participants eventually solved the game (within an allotted 10 guesses) given how often they had picked a low complexity guess in the forced-choice portion of the task. There is a striking relationship, with picking low-complexity guesses in the forced-choice phase strongly predicting eventually finding a solution. In a logistic regression with random subject effects predicting solving the game from the proportion of low-complexity guesses, the intercept significantly differed from 0 ($B_0 = -0.78$; $z = -2.34$; $p = 0.02$), as did the proportion of low-complexity guesses ($B_{\text{low}} = 1.74$; $z = 4.01$; $p < 0.001$). This indicates that on trials where only the max-EIG option was picked, participants solved the game only about 31% of the time; but when only the low-complexity option was chosen, participants solved the game nearly 72% of the time. Similarly, as Figure 6c shows, on trials where participants picked the low-complexity more frequently, they also required *fewer* guesses to get the right answer in phase two. A linear regression with subject parameters found a significant intercept ($B_0 = 8.03$; $t = 15.0$; $p <$

³For ease of interpretation, trial 1 was set to trial 0. So the intercept represents how often participants picked the low complexity option on the first game.




Figure 6: Results from Experiment 3. (a) The probability that participants picked the low-complexity code (y-axis) in the forced-choice portion of each game across trials (x-axis). (b) The probability that participants eventually solved a game (within 10 guesses) in phase two (y-axis) given the proportion of low-complexity guesses they made in phase one (x-axis). (c) The number of guesses participants required to solve the game in phase two (y-axis) given the proportion of low-complexity guesses they made in phase one (x-axis), non-solvers excluded.

0.001) and a significant effect of picking the low-complexity choice ($B_{\text{low}} = -2.65$; $t = -3.07$; $p < 0.001$), indicating that participants required more guesses to get the correct answer when they did not pick the low-complexity option frequently in the forced choice phase. This effect remained even removing trials where participants eventually solved the task.

Discussion

In this paper, we asked how human active learning is shaped by processing demands, which we tested using the game Mastermind. Across our three experiments, we demonstrated: 1) that people make queries that are far from maximally-informative; 2) that they in fact cannot *use* informative queries effectively because they are too complex to reason about; and 3) that when given the choice, they actually prefer lower-complexity, lower-information queries. These findings challenge simple Optimal Experimental Design theories of active learning that do not account for costs and limitations associated with learning, and instead support a more nuanced story (cf. Gong et al., 2023): that people are cognitively limited in what data they can process effectively; they know this (or can learn their limits quickly); and they make queries with their limited capacity for processing information in mind.

Previous research on human active learning has used a modified version of Mastermind, called “Entropy Mastermind,” in which the distribution over digits/colors in a code is allowed to be non-uniform (Schulz et al., 2019). In that work, they found that people required fewer queries, were faster to generate queries, and had faster learning rates if the entropy of the generating distribution was lower. They also found that participants adapted to the generating distribution across trials. However, it is evident from their results that participants’ queries were inefficient in absolute terms, but the authors’ primary interest was in directional effects relating to the entropy over codes — and the authors’ analyses did not address this issue. The work we present here complements these results by explaining how people might efficiently seek maximal information under a limited processing capacity, and

hence might seem inefficient in absolute terms.

The LoT-based approach we took shares similarities with the “mental models” theory of reasoning developed by Philip Johnson-Laird and others, which has been used to capture the incorrect inferences people draw from sets of logical postulates (e.g. Craik, 1967; Johnson-Laird et al., 1998, 2004). The mental models approach assumes that participants reduce propositions into sets of independent possibilities about what might be true. This theory has been used to explain a number of fallacies and “illusions” people make when evaluating sets of logical assertions, e.g. with the form “if A then B or else C.” In one study, they found that as the number of possibilities compatible with the assertions increases, the difficulty of the task increases, and that reasoners represent what is true according to assertions, but do not keep track of what is false (Johnson-Laird et al., 2000) — consistent with their theory (and with ours). The LoT framework we employed here extends the mental models approach to a potentially unbounded space of representations, such that it allows for *perfect* reasoning in the limit of expressivity (i.e. when unboundedly long descriptions are allowed) but also for constrained reasoning when expressivity is restricted.

There are two limitations worth noting. First, as with all such “Language of Thought” based modeling approaches, the particular choice of primitive operations in the model allows for a significant degree of flexibility. We attempted to mitigate this issue by choosing a fairly agnostic set of primitives that have been used to model similar tasks (Goodman et al., 2008; Piantadosi et al., 2016), and by not fitting production weights for the primitives to data. Second, our model does not address at an algorithmic level how people actually choose codes — we have only shown that people find certain information difficult to process, and that they prefer information that has low expected complexity under our model. How people compute the expected complexity of a query, or whether they actually make this computation at all, remains unclear.

Acknowledgements

We would like to thank Lara Bertram, Florian Elsaesser, and Alberto Feducci for helpful discussions. We would also like to acknowledge Anselm Rothe, who co-developed a separate bounded-optimal Mastermind model called “Braveheart.”

References

- Bramley, N. R., Dayan, P., Griffiths, T. L., & Lagnado, D. A. (2017). Formalizing neurath’s ship: Approximate algorithms for online causal learning. *Psychological review*, *124*(3), 301.
- Bramley, N. R., Lagnado, D. A., & Speekenbrink, M. (2015). Conservative forgetful scholars: How people learn causal structure through sequences of interventions. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *41*(3), 708.
- 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*(5), 1548–1587.
- Craik, K. J. W. (1967). *The nature of explanation* (Vol. 445). CUP Archive.
- Gong, T., Gerstenberg, T., Mayrhofer, R., & Bramley, N. R. (2023). Active causal structure learning in continuous time. *Cognitive Psychology*, *140*, 101542.
- Goodman, N. D., Tenenbaum, J. B., Feldman, J., & Griffiths, T. L. (2008). A rational analysis of rule-based concept learning. *Cognitive science*, *32*(1), 108–154.
- Gregg, L. W., & Simon, H. A. (1967). Process models and stochastic theories of simple concept formation. *Journal of Mathematical Psychology*, *4*(2), 246–276.
- Gureckis, T. M., Martin, J., McDonnell, J., Rich, A. S., Markant, D., Coenen, A., Halpern, D., Hamrick, J. B., & Chan, P. (2016). Psiturk: An open-source framework for conducting replicable behavioral experiments online. *Behavior research methods*, *48*(3), 829–842.
- Horwich, P. (1982). How to choose between empirically indistinguishable theories. *The Journal of Philosophy*, *79*(2), 61–77.
- Johnson-Laird, P. N., Girotto, V., & Legrenzi, P. (1998). Mental models: A gentle guide for outsiders. *Sistemi Intelligenti*, *9*(68), 33.
- Johnson-Laird, P. N., Girotto, V., & Legrenzi, P. (2004). Reasoning from inconsistency to consistency. *Psychological Review*, *111*(3), 640.
- Johnson-Laird, P. N., Legrenzi, P., Girotto, V., & Legrenzi, M. S. (2000). Illusions in reasoning about consistency. *Science*, *288*(5465), 531–532.
- Klayman, J., & Ha, Y.-w. (1989). Hypothesis testing in rule discovery: Strategy, structure, and content. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *15*(4), 596.
- Markant, D. B., Settles, B., & Gureckis, T. M. (2016). Self-directed learning favors local, rather than global, uncertainty. *Cognitive science*, *40*(1), 100–120.
- Meder, B., Nelson, J. D., Jones, M., & Ruggeri, A. (2019). Stepwise versus globally optimal search in children and adults. *Cognition*, *191*, 103965.
- Oaksford, M., & Chater, N. (1994). A rational analysis of the selection task as optimal data selection. *Psychological Review*, *101*(4), 608.
- Piantadosi, S. T., Tenenbaum, J. B., & Goodman, N. D. (2016). The logical primitives of thought: Empirical foundations for compositional cognitive models. *Psychological review*, *123*(4), 392.
- Popper, K. R. (1963). Science as falsification. *Conjectures and refutations*, *1*(1963), 33–39.
- Rothe, A., Lake, B. M., & Gureckis, T. (2017). Question asking as program generation. *Advances in neural information processing systems*, *30*.
- Rothe, A., Lake, B. M., & Gureckis, T. M. (2016). Asking and evaluating natural language questions. *CogSci*.
- Schulz, E., Bertram, L., Hofer, M., & Nelson, J. D. (2019). Exploring the space of human exploration using entropy mastermind. *bioRxiv*, 540666.