

Experience-driven discovery of planning strategies

Ruiqi He (ruiqi.he@tuebingen.mpg.de)

Max Planck Institute for Intelligent Systems
Tuebingen, Germany

Falk Lieder (falk.lieder@psych.ucla.edu)

University of California, Los Angeles, Department of Psychology
Los Angeles, United States

Abstract

One explanation for how people can plan efficiently despite limited cognitive resources is that we possess a set of adaptive planning strategies and know when and how to use them. But how are these strategies acquired? While previous research has studied how individuals learn to choose among existing strategies, little is known about the process of forming new planning strategies. In this work, we propose that new planning strategies are discovered through metacognitive reinforcement learning. To test this, we designed an experiment to investigate the discovery of new planning strategies. We then presented metacognitive reinforcement learning models and demonstrated their capability for strategy discovery as well as show that they provided a better explanation of human strategy discovery than alternative learning mechanisms. However, when fitted to human data, these models exhibited a slower discovery rate than humans, leaving room for improvement.

Keywords: metacognitive learning, planning, strategy discovery, cognitive modeling, reinforcement learning

Introduction

From organizing dinner to drafting a summer holiday or mapping out a career, people engage in planning daily. Planning can be visualized as a decision tree that grows rapidly with more actions and steps to consider. While computers can rely on increasing computational power to solve such problems, human cognitive resources are limited. Yet, people often excel in complex domains where even advanced computers struggle, raising the question: how do people plan so efficiently despite cognitive limitations? One explanation is that people use clever strategies to facilitate planning (Callaway et al., 2022), but how do these strategies originate? While previous research showed how people learn to choose between pre-existing strategies from a mental toolbox through reinforcement (Rieskamp & Otto, 2006), little is known about how strategies enter this toolbox in the first place. Developmental psychology research (Siegler & Crowley, 1991) found that children can develop new strategies for solving arithmetic problems suggesting that strategy discovery may be a source of these strategies. While we understand some aspects of strategy discovery in mental arithmetic, the mechanism behind the acquisition of planning strategies is largely unknown. Therefore, in this work we investigated the question of how people discover new planning strategies.

We address this question with two primary contributions: First, we introduced a planning task requiring participants to learn an unfamiliar strategy to measure whether and to what

extent strategy discovery occurs. Second, we presented a computational model for strategy discovery based on the concept of *metacognitive reinforcement learning* (Callaway, Gul, Krueger, Griffiths, & Lieder, 2018), which demonstrated a superior explanation of how people learn to adapt their amount of planning (He, Jain, & Lieder, 2021a) and planning strategies (He, Jain, & Lieder, 2021b). Using the data collected from the task, we then compared metacognitive reinforcement learning (MCRL) against alternative learning mechanisms to evaluate its potential as a model of human strategy discovery.

In the following sections, we will first present our experiment, the empirical findings, then the various learning mechanisms, the resulting models, the procedures for model fitting and selection, and conclude with the modeling results.

Measuring strategy discovery

To measure strategy discovery, we first conducted an experiment using a process-tracing paradigm called Mouselab-MDP (Callaway, Lieder, Krueger, & Griffiths, 2017) and then analyzed the collected data to examine whether and to which extent strategy discovery took place.

Mouselab-MDP

The Mouselab-MDP paradigm (Callaway et al., 2017) extends the widely used Mouselab paradigm (Payne, 1976) to address sequential decision-making problems. Essentially, the Mouselab-MDP presents participants with a planning task designed to render their information-gathering behavior highly indicative of their underlying planning strategies. Figure 1 illustrates several example trials. Participants are asked to plan a spider's route through a maze with the goal to maximize their score. To track the planning process, the values associated with the nodes are initially hidden. Participants can click on a node to reveal its reward or loss. This clicking behavior provides a manifestation of the participant's planning operations through information-gathering (i.e., participants need to sequentially decide which information to gather to inform their planning). The cognitive cost of this planning operation is externalized by imposing a fee for each click, encouraging participants to reveal information only when necessary. Therefore, the final score is the sum of the values of the gray nodes the participant chooses to traverse minus the fee for clicking. This score serves as a measure of resource-

rationality (Lieder & Griffiths, 2020) that balances the expected utility with the cognitive cost of employing a particular strategy. The strategy that yields the highest score given the costs is referred to as the resource-rational (RR) strategy.

Experiment

Materials We utilized the Mouselab-MDP paradigm to create an environment where the RR planning strategy was qualitatively different from any strategies participants might have previously known. In each trial, the environment could assume one of six possible configurations, as illustrated in Figure 1. Each configuration was designed so that exactly one branch began with a positive reward. The outer nodes on the branch starting with an immediate positive reward always contained one node with a high positive reward (+50) and one node with a high negative reward (-50). The cost of revealing a node’s hidden reward (click cost) increased with its distance from the starting point: 1 for the immediate nodes, 3 for the middle nodes, and 30 for the outer nodes. Consequently, the RR strategy involved first examining the immediate nodes until identifying the branch that began with a positive immediate reward. The RR strategy would then examine exactly one of the outer nodes on that branch. If the value of that node was positive, the agent would choose the path leading to it; otherwise, it would choose the path leading to the other outer node on that branch.¹ Given the highly specialized nature of this strategy, we assumed that participants would not plausibly know it before the experiment.

Participants We recruited 420 participants through the platform CloudResearch, of whom 378 completed the experiment (average age 38.77 years; 240 females, 3 identifying as other). Recruitment was restricted to participants located in the United States who had completed more than 100 HITs and had an approval rate of over 80%. Participants earned a bonus of 0.3 cents for each point in their final score. On average, they received a bonus of \$0.56 in addition to the base pay of \$4. The mean duration of the experiment was 25.8 minutes (median 22.8 minutes). Participants who either did not finish the experiment or failed to click anything throughout all trials were excluded from further analysis. This exclusion criterion was designed to filter out inattentive participants. After applying these criteria, 349 participants remained for analysis.

Procedures Each participant was given minimal instructions on how many trials they would face, how to reveal the nodes, and that each click would incur an unknown fee. They were also informed about the maximum and minimum values of all possible rewards and the score that the optimal strategy could achieve in each trial. After passing a quiz, they were asked to complete 120 trials of the planning task.

¹We confirmed that this is indeed the RR strategy for this environment by computing the environment’s optimal metalevel policy using dynamic programming (Callaway et al., 2022). The RR strategy is the only strategy that would consistently result in a positive score in this planning task.

Empirical results

To investigate whether experience-driven strategy discovery was taking place, we first classified the participants’ click sequences into two categories: adaptive strategies and other strategies. Adaptive strategies were defined according to the criteria outlined in the materials section. Using this classification, we plotted the proportion of adaptive strategies across trials and observed an increase from 0.79% (CI: [0; 2.22%]) in the first trial to 28.57% (CI: [27.21%; 29.93%]) in the last trial (see the green line in Figure 2). A logistic regression analysis (adaptive strategy \sim intercept + trial) confirms that this upward trend is significant, with a positive coefficient on the trial regressor, indicating an increasing proportion of adaptive strategies as participants gained more experience (see first two rows in the Table 1). This was further supported by a non-parametric Mann-Kendall test of trend ($S = 6506, p < .001$).

	Coefficient	Std err	p-value
Intercept	-1.29	0.01	<.001
Trial	4e-4	3e-5	<.001
Intercept	-9.39	0.57	<.001
Trial	0.01	0.00	<.001

Table 1: Results of the logistic regression on the proportion of the adaptive strategy (upper two rows) and the mixed linear regression on the score (lower two rows).

Next to the significantly increasing proportion of adaptive strategy, a mixed linear regression on the score (score \sim intercept + trial) also showed a significantly increasing trend (see last two rows in the Table 1 and the blue line in Figure 2).

Both results provide evidence for the occurrence of experience-driven strategy discovery. However, strategy discovery appeared to be a challenging task, as the proportion of participants who discovered the novel adaptive planning strategy increased only very slowly with only about 29% of participants managing to learn it after 120 trials.

Modeling strategy discovery

Having established that experience-driven strategy discovery occurred, we then shifted our focus to examining the underlying mechanism behind *how* new planning strategies are discovered. We will first provide the necessary background, introduce the MCRL models, and demonstrate their ability for strategy discovery. Then to examine its alignment with human strategy discovery, we will introduce alternative learning mechanisms, outline the model fitting and selection procedures, and present the results.

Reinforcement learning

Reinforcement learning (RL) has proven to be a promising framework for modeling how people learn from interactions with the environment (Niv, 2009). Similar to human trial-and-error learning, a common type of RL involves predicting




Figure 1: Each trial is randomly sampled from one of those six configurations.

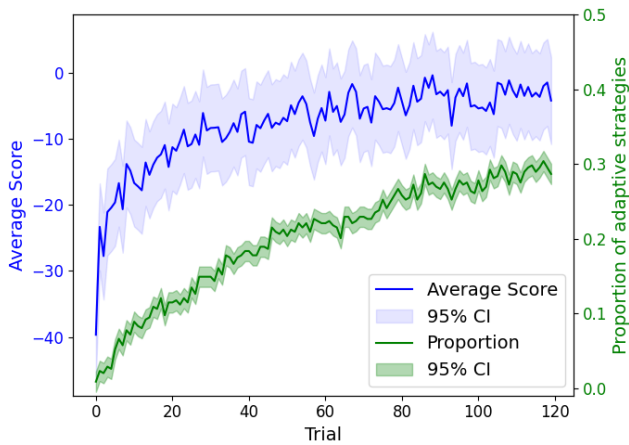


Figure 2: The average score and the proportion of the adaptive strategy plotted across 120 trials of all 349 participants.

the expected reward for a specific action a in a given state s and updating this prediction based on the observed outcome r .

Metacognitive reinforcement learning (MCRL)

While standard reinforcement learning focuses on learning policies for external actions, metacognitive learning operates on internal mental operations. To explore metacognitive learning, previous work (Krueger, Lieder, & Griffiths, 2017) utilized the MCRL framework to study the problem of deciding how to decide (meta-decision-making (Boureau, Sokol-Hessner, & Daw, 2015)) as a meta-level MDP (Hay, Russell, Tolpin, & Shimony, 2014): $M_{meta} = (\mathcal{B}, \mathcal{C} \cup \{\perp\}, T_{meta}, r_{meta})$. In this framework, $b_t \in \mathcal{B}$ represents the mental belief state, $c_t \in \mathcal{C}$ denotes the mental computations including deliberation termination \perp . T_{meta} de-

scribes the meta-level transitions between belief states, and r_{meta} is the cost of computations as well as the expected return from terminating thinking and acting based on the current belief state. Solving a meta-level MDP optimally is often computationally challenging. However, the optimal solution can be approximated using MCRL by assuming that the brain estimates optimal meta-decision-making through reinforcement learning mechanisms (Russell & Wefald, 1991; Callaway et al., 2018) that learn to approximate the optimal strategy directly (He et al., 2021b).

Before diving into the MCRL mechanisms, we will first introduce how the planning strategies are represented.

Strategy representation In the MCRL framework, planning strategies are represented by a set of features and corresponding feature weights. To represent the unique planning strategy of our planning task, we extended the feature set outlined by Jain et al. (2022), who developed these features based on the Mouselab-MDP paradigm, by adding seven new features that are essential to capture preferences and avoidance related to specific node levels, as well as one additional stopping criterion. This extension results in a total of 68 features, which can be grouped into six broader categories believed to influence planning (Keramati, Smittenaar, Dolan, & Dayan, 2016; Daw, 2018): Pavlovian features, model-free and heuristic features, mental-effort avoidance features, satisficing and stopping criteria features, model-based metareasoning features, and habitual features that reflect mental habits related to how frequently a node, branch, or level has been examined in the past and how many planning operations have been performed before. Full details on these features can be found in the supplementary materials: <https://osf.io/yqw9a>. This strategy representation allows us to characterize an individual’s learning trajectory as a series of combinations of features where the feature weights adjust over trials.

Gradient ascent through the strategy space (Reinforce)

The metacognitive Reinforce model (Jain, Callaway, & Lieder, 2019), based on the Reinforce algorithm by Williams (1992), assumes that individuals adapt their planning strategies within the space of possible strategies. The strategies, akin to the term policy in the context of policy-gradient reinforcement learning, map features of belief states to values of cognitive operations, which are then transformed into probabilities of specific planning operations via a softmax function (Williams, 1992). Specifically, the strategy is parameterized by a weight vector \mathbf{w} , which is updated after each trial to reflect the effectiveness of different cognitive operations:

$$\mathbf{w} \leftarrow \mathbf{w} + \alpha \cdot \sum_{t=1}^O \gamma^{t-1} \cdot r_{meta}(b_t, c_t) \cdot \nabla_{\mathbf{w}} \ln \pi_{\mathbf{w}}(c_t | b_t), \quad (1)$$

where O represents the number of planning operations carried out during the trial, b the belief state, c the click under consideration, C_b the set of cognitive operations available in belief state b , α the learning rate, and γ the discount factor. The learned weights, combined with the features described earlier are then used to select the planning operations using a softmax function:

$$\pi_{\mathbf{w}}(c|b) = \frac{\exp\left(\frac{1}{\tau} \cdot \sum_{k=1}^{68} w_k \cdot f_k(b, c)\right)}{\sum_{c \in C_b} \exp\left(\frac{1}{\tau} \cdot \sum_{k=1}^{68} w_k \cdot f_k(b, c)\right)} \quad (2)$$

where τ is the inverse temperature that controls the balance between exploration and exploitation. The free hyperparameters are α , γ , and τ as well as the 68 initial feature weights.

We derived two variants of the Reinforce model: a hybrid variant that employs the full set of features including the model-based metareasoning features that capture the uncertainty of a node, as defined by the standard deviation and a purely model-free variant that utilizes a subset of the features by excluding the model-based metareasoning features, therefore positing that adaptation is driven purely by interactions with the environment (see first two rows in Table 2).

Capability to represent and to discover the novel planning strategy

To evaluate our model’s capacity to represent and discover the adaptive strategy, we conducted model simulations using hyperparameters that were found by optimizing for the optimal click sequence. Figure 3 shows that both the hybrid and model-free Reinforce models were capable of discovering the adaptive strategy from scratch (Mann Kendall test of trend for all simulations: $S > 3442$, $p < .001$). While both the full and reduced set of features were sufficient to represent the adaptive strategy, it remains uncertain whether one model variant consistently outperforms the other which could be subject to feature importance or ablation studies.

Alternative learning mechanisms

Having demonstrated that our MCRL models are capable of strategy discovery, we now turn to examining how well our models can explain the human strategy discovery process. For this, we introduce two alternative learning mechanisms to compare against as well as a non-learning model.

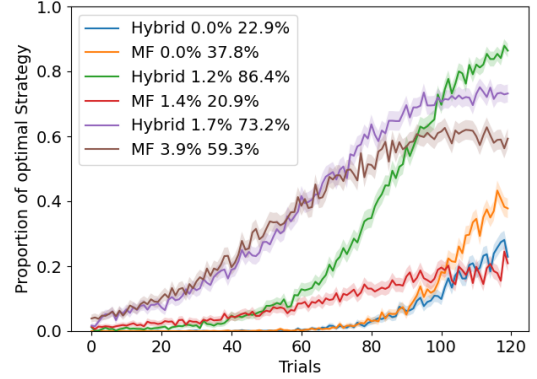


Figure 3: Simulations of the hybrid and model-free Reinforce models, each with 3 different sets of hyperparameters. Each curve shows the averaged performance across 1000 runs and the corresponding 95% confidence intervals. Details on the hyperparameters of one of the simulations can be found in the supplementary materials: <https://osf.io/yqw9a>

Rational strategy selection learning Strategy selection learning (Rieskamp & Otto, 2006), posits that individuals possess a repertoire of cognitive strategies for problem-solving and that these strategies are reinforced through feedback. Similar to this idea, Lieder and Griffiths (2017) developed the model of rational strategy selection learning (RSSL), which frames the problem of deciding how to plan as an n -armed bandit problem, with each arm representing a different strategy. In this model, each planning strategy is parameterized by the mean and variance of its expected return. Bayesian inference is performed on the expected return of each strategy, and strategies are selected using Thompson sampling from the posterior distribution of the expected return of the strategies. We will adopt the implementation described by Lieder and Griffiths (2017) and use the set of 79 strategies pre-defined by Jain et al. (2022), which were extensively tested on the Mouselab-MDP. Although the optimal strategy is not included in this set to reflect the constraints of working with a predefined set of strategies, the set contains a strategy that closely approximates the optimal one—specifically, a strategy that involves examining all immediate nodes and then one outer node at random.

Mental habit formation According to this model, changes in planning strategy are not driven by value but are influenced by mental habits. It suggests that an individual’s tendency to execute a particular planning operation increases with the number of times they have previously performed it regardless of the outcome, consistent with findings by Miller, Shenhav, and Ludvig (2019). This approach involves using a weighted sum of frequency-based data, such as the number of prior clicks on the same node, branch, and level, within a softmax decision process. Technically, we implemented this model as a Reinforce model that uses the same set of features as the model-free variant but does not update the weights of the features. This model therefore adjusts its planning based on the

Model name	Updating parameters	Features
Hybrid Reinforce	Yes (updating weights)	Pavlovian, model-free, satisficing, habitual, model-based
Model-free Reinforce	Yes (updating weights)	Pavlovian, model-free, satisficing, habitual
RSSL	Yes (updating mean and variance of each of the 79 strategies)	Not applicable
Mental habit	No	Same as model-free Reinforce
Non learning	No	Pavlovian, model-free, satisficing

Table 2: Overview of the models, their features, and updating mechanism.

frequency of previous planning operations from past experiences through the habitual features.

Non-learning model This model is implemented as a Reinforce model that does not perform any weight updates and does not use habitual features.

In summary, we considered five different models: two MCRL models (hybrid and model-free Reinforce), the mental-habit, the RSSL, and the non-learning model. They differ in the set of features used as well as whether and how the parameters are being updated (see overview in Table 2).

Model fitting and model selection

After having introduced all the models, we now turn to comparing the MCRL mechanism against the alternatives. For this, all the models were fitted to each participant individually by maximizing the likelihood of their click sequences using Bayesian optimization (Bergstra, Yamins, & Cox, 2013) for 60,000 iterations. The likelihood of a click sequence is calculated as the product of the likelihood of the individual clicks. After fitting the models, we applied family-level Bayesian model selection (BMS) at three different aggregation levels. Specifically, we estimated the expected proportion of participants best described by each model family (r) and the exceedance probability (ϕ), which indicates the likelihood that this proportion is significantly higher than that for any other model family using random effect Bayesian model selection (Rigoux, Stephan, Friston, & Daunizeau, 2014). The first level distinguished between non-learning models and learning models. The second level compared the basic learning mechanisms: Reinforce, RSSL, and the mental habit model. The third level compared the two Reinforce variants: hybrid and model-free. In addition to BMS, we used the Bayesian Information Criterion (BIC) (Schwarz, 1978) to assess the relative fit of each model to individual participants' data.

How well can the models explain human strategy discovery?

BMS indicated that the majority of participants are better explained by a learning model (90%) rather than a non-learning model (see Table 3). Among the learning models, the Re-

inforce models accounted for approximately 61% of participants, outperforming the mental habit model (39%) and RSSL (< 1%) with 100% exceedance probability. Within the Reinforce variants, the hybrid Reinforce model explained 62% of participants better than the model-free variant (38%). These observations indicated that while the MCRL models, especially the hybrid Reinforce, explains more participants better than the alternative learning mechanisms, there were individual differences in the learning mechanisms. Therefore, we next looked into the performance of the different participant groups. When examining participants grouped

Level	Model	r	ϕ
1	Learning models	89.65%	100.00%
	Non-learning model	10.35%	0.00%
2	Reinforce models	60.89%	99.99%
	Mental habit model	38.83%	0.01%
	RSSL	0.29%	0.00%
3	Hybrid Reinforce	62.31%	100.00%
	Model-free Reinforce	37.69%	0.00%

Table 3: BMS results where r described the proportion of participants best explained by the model and ϕ the probability that this proportion is higher than the alternative across different aggregation levels.

by the model with the best goodness-of-fit (lowest BIC) an interesting pattern emerged. Participants best explained by the mental habit model (referred to as habitual participants), which does not learn but merely tends to repeat previous planning operations, performed better than those explained by the hybrid Reinforce mechanism (Mann-Whitney U test: mental habit vs. hybrid Reinforce $S > 1e8, p < .001$; no significant difference noted when comparing habitual with the model-free Reinforce participants $S > 9e7, p = .108$). The corresponding mean and standard deviations can be found in Table 4. This counterintuitive observation can, however, be explained by some participants classified as habitual learners showing rapid improvement in the first 20 trials, followed by a flattened learning curve (see Figure 4). Looking into the reason behind this explosive improvement reveals that 21 out of 121 habitual participants examined all the nodes early on (hence the extremely low score in the first trials), thereby quickly learned the environment's structure, and then adopted the optimal strategy. This type of insight-like learning was much less pronounced in other participant groups (15 out of 228 non-habitual participants). Therefore, these participants may be misclassified, as the mental habit model better captured their relatively flat learning curve following their abrupt change in behavior. This type of abrupt learning is currently not represented by any of the current models as they exhibit rather gradual improvements (see Figure 3). Excluding the abrupt-learners resulted in similar empirical results (see supplementary materials).

This limitation of the current MCRL models is reflected when simulating the discovery process using the models with

	Mean	Std
Mental Habit	-7.68	38.23
Model-free Reinforce	-5.97	33.78
Hybrid Reinforce	-9.04	37.99
Kruskal Wallis test	$S = 37.76, p < .001$	

Table 4: Mean and SD of the score of the participant groups best fitted by a model type and Kruskal-Wallis test testing the score for differences between all three samples.




Figure 4: Average scores of the participants grouped by the model that has the lowest BIC. Participants who have been best explained by the mental habit model have been further subdivided into those who examined all the nodes in the first trial and those who did not.

hyperparameters fitted to the participants' clicking behavior. While the two Reinforce models, the mental habit model and RSSL, exhibit significant increasing trends in the proportion of adaptive strategies according to four Mann-Kendall tests (all $S > 3024$ and all $p < .001$) as opposed to the non-learning model ($S = -115, p = .792$), participants discovered the new planning strategies at a different pace (see the increase in proportions in the legend of Figure 5). For each of the models, the slope was significantly lower (all $p < .001$).

Discussion and further work

This work investigated the underlying mechanisms of strategy discovery by designing an experiment that required learning a highly task-specific planning strategy that was unlikely part of participants' pre-existing mental toolbox. Empirical results showed that while experience-driven strategy discovery was taking place, it proved to be a challenging task, with only about 29% of participants discovering the novel strategy after 120 trials. We then introduced the metacognitive reinforcement learning (MCRL) models and demonstrated their capability to discover new planning strategies. Fitting our models and alternative learning mechanisms to participant data revealed that 61% of individuals were better explained by the metacognitive Reinforce model, while 39% were bet-




Figure 5: Model simulations using fitted hyperparameters of participants who did not examine all nodes at the first trial ($n = 311$) and their proportion of the adaptive strategy at the first and last trial.

ter accounted for by a value-free mental habit model. These findings highlight two key insights: (i) metacognitive reinforcement learning can explain a larger proportion of participants, suggesting that people likely adjust their planning strategies directly based on observed outcomes; and (ii) there are notable individual differences in metacognitive learning. Moreover, while our models are in principle able to discover new planning strategies (see Figure 3 for model simulations fitted with optimised hyperparameters), fitting the models' hyperparameters on the participants' click sequences caused the models to discover the adaptive strategy at a slower rate than participants (see Figure 5 for model simulation with hyperparameters fitted to participants' data). We attribute this discrepancy to several factors. First, some participants seemed to experience sudden, significant improvements in performance as well as engaged in active learning - an aspect not captured by our current set of MCRL models. To address this, future work could explore additional mechanisms incorporating active learning or insight-based models, potentially through model-based approaches that infer environmental structure to derive optimal strategies. Second, while our current feature set effectively represents the optimal strategy and captures a wide range of strategies, we cannot ensure that it encompasses all potential intermediate strategies participants might have adopted.

Overall, our work advances our understanding of metacognitive learning and how people discover adaptive planning strategies as well as contributes to the advances in artificial intelligence (AI) to reverse-engineer or even surpass human strategy discovery capacity. While MCRL may not explain every aspect of this process, it is a key mechanism for uncovering strategies in planning tasks. We hope these insights enhance our understanding of human strategy discovery and inspire adaptive AI models across diverse domains.

References

- Bergstra, J., Yamins, D., & Cox, D. (2013). Making a science of model search: Hyperparameter optimization in hundreds of dimensions for vision architectures. In *International conference on machine learning* (pp. 115–123).
- Boureau, Y.-L., Sokol-Hessner, P., & Daw, N. D. (2015). Deciding how to decide: Self-control and meta-decision making. *Trends in cognitive sciences*, 19(11), 700–710.
- Callaway, F., Gul, S., Krueger, P., Griffiths, T. L., & Lieder, F. (2018). Learning to select computations. In *Uncertainty in artificial intelligence: Proceedings of the thirty-fourth conference*.
- Callaway, F., Lieder, F., Krueger, P. M., & Griffiths, T. L. (2017). Mouselab-mdp: A new paradigm for tracing how people plan. In *The 3rd multidisciplinary conference on reinforcement learning and decision making*.
- Callaway, F., van Opheusden, B., Gul, S., Das, P., Krueger, P. M., Lieder, F., & Griffiths, T. L. (2022). Rational use of cognitive resources in human planning. *Nature Human Behaviour*, 1–14.
- Daw, N. D. (2018). Are we of two minds? *Nature Neuroscience*, 21(11), 1497.
- Hay, N., Russell, S., Tolpin, D., & Shimony, S. E. (2014). Selecting computations: Theory and applications. *arXiv preprint arXiv:1408.2048*.
- He, R., Jain, Y. R., & Lieder, F. (2021a, December). Have i done enough planning or should i plan more? In *Workshop on metacognition in the age of ai. thirty-fifth conference on neural information processing systems*. Long Paper.
- He, R., Jain, Y. R., & Lieder, F. (2021b). Measuring and modelling how people learn how to plan and how people adapt their planning strategies the to structure of the environment. In *International conference on cognitive modeling*. Retrieved from <https://mathpsych.org/presentation/604/#/document>
- Jain, Y. R., Callaway, F., Griffiths, T. L., Dayan, P., He, R., Krueger, P. M., & Lieder, F. (2022). A computational process-tracing method for measuring people's planning strategies and how they change over time. *Behavior Research Methods*. doi: <https://doi.org/10.3758/s13428-022-01789-5>
- Jain, Y. R., Callaway, F., & Lieder, F. (2019). Measuring how people learn how to plan. In *Cogsci* (pp. 1956–1962).
- Keramati, M., Smittenaar, P., Dolan, R. J., & Dayan, P. (2016). Adaptive integration of habits into depth-limited planning defines a habitual-goal-directed spectrum. *Proceedings of the National Academy of Sciences*, 113(45), 12868–12873.
- Krueger, P. M., Lieder, F., & Griffiths, T. (2017). Enhancing metacognitive reinforcement learning using reward structures and feedback. In *Cogsci*.
- Lieder, F., & Griffiths, T. L. (2017). Strategy selection as rational metareasoning. *Psychological Review*, 124(6), 762–794.
- Lieder, F., & Griffiths, T. L. (2020). Resource-rational analysis: Understanding human cognition as the optimal use of limited computational resources. *Behavioral and brain sciences*, 43.
- Miller, K. J., Shenhav, A., & Ludvig, E. A. (2019). Habits without values. *Psychological review*, 126(2), 292.
- Niv, Y. (2009). Reinforcement learning in the brain. *Journal of Mathematical Psychology*, 53(3), 139–154.
- Payne, J. W. (1976). Task complexity and contingent processing in decision making: An information search and protocol analysis. *Organizational behavior and human performance*, 16(2), 366–387.
- Rieskamp, J., & Otto, P. E. (2006). Ssl: a theory of how people learn to select strategies. *Journal of Experimental Psychology: General*, 135(2), 207.
- Rigoux, L., Stephan, K. E., Friston, K. J., & Daunizeau, J. (2014). Bayesian model selection for group studies—revisited. *Neuroimage*, 84, 971–985.
- Russell, S., & Wefald, E. (1991). Principles of metareasoning. *Artificial intelligence*, 49(1-3), 361–395.
- Schwarz, G. (1978). Estimating the dimension of a model. *The annals of statistics*, 461–464.
- Siegler, R. S., & Crowley, K. (1991). The microgenetic method: A direct means for studying cognitive development. *American psychologist*, 46(6), 606.
- Williams, R. J. (1992). Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 8(3), 229–256.