Deciphering the meta-cognitive experiences in creative problem-solving

Yang-Fan Lu (vanesslu@pku.edu.cn)
Academy for Advanced Interdisciplinary Studies and Center for Life Sciences, Peking University
Beijing 100871, China

Qianli Yang (qlyang@ion.ac.cn)
Center for Excellence in Brain Science and Intelligence Technology, Chinese Academy of Sciences
Shanghai 200031, China

Hang Zhang (hang.zhang@pku.edu.cn)
School of Psychological and Cognitive Sciences, Peking University
Beijing 100871, China

Abstract

Previous studies on human meta-cognition, represented by confidence in perceptual decisions, often focus on over-simplified environments that yield experiences with limited semantic dimensions. However, in real-life situations such as solving a new problem, people need to make sequential decisions in a complex environment, exploring vast combinations of actions that unfold over time. How do people make meta-cognitive evaluations out of the rich, high-dimensional cognitive experiences in such situations? Here we develop a computational method that models each individual's meta-cognitive ratings (e.g., difficulty) of problem-solving experience in a visual puzzle game based on information-theoretic metrics derived from their own action sequences. Individuals are assumed to Bayesian update their "thought-space distributions" with their own behavioral distributions on different semantic categories. Our preliminary results show that information discrepancies between beliefs at different moments can predict the individual differences in self-reported difficulty.

Keywords: problem-solving; Bayesian inference; meta-cognition; cognitive effort; information theory

Introduction

Meta-cognitive evaluations are crucial for humans to flexibly refine behaviors in inference and decision-making. Though confidence judgments have been widely studied in static and dynamic environments (Meyniel et al., 2015; Sanders et al., 2016), with perceptual and value-based tasks (da Silva Castanheira et al., 2021; Rahnev et al., 2015), demonstrating accuracies and biases (Lebreton et al., 2019; Ting et al., 2023), much less is known about meta-cognition in complex sequential decision-making tasks such as solving new problems.

To predict meta-cognitive evaluations in these more realistic situations, we model humans as Bayesian observers by extending previous work on confidence judgment (Fleming & Daw, 2017) and the sub-goaling and policy updating in repetitively solving similar problems (Binz & Schulz, 2023; Donnarumma et al., 2016). We assume that human problem-solvers use their own action sequences to update the probability distribution of potentially useful actions in their thinking. We then derive information-theoretic metrics from these "thought-space distributions", similar to quantifying cognitive effort (Zenon et al., 2019) and capacity (Binz & Schulz, 2023; Prat-Carrabin et al., 2021) using information theory. We tested the predictive power of these metrics in human experimental data using human experimental data of retrospective self-reports on the difficulty and creativity of problem-solving in visual puzzle games.

Model

We model the problem-solver as a Bayesian ideal observer of one's own behaviors $[a_0, a_1, \cdots, a_{T-1}]$ and changing states of the environment $[o_0, o_1, \cdots, o_T]$, at arbitrary time points $[t_0, t_1, \cdots, t_T]$. The granularity of the time series depends on different modeling contexts and needs: it can be as detailed as each step of action selection or as coarse as the beginning and end of each problem.

Online updating of thought-space distribution

The problem-solver represents the observations of behaviors $[a_0, a_1, \cdots, a_{T-1}]$ and $[o_0, o_1, \cdots, o_T]$ in a general "thought space" $\boldsymbol{\mathcal{X}}$. The primitive unit of this "thought" could be action category or state abstraction, as long as the representation of such thoughts is universal across multiple trials or problems.

The "thought space" serves as an effective abstraction of the problem space. As the problem-solver interacts with the environment, the more frequently used thoughts become more activated in their internal model. Mathematically, the likelihood of activating each thought during period $[t_j, t_{j+1})$ follows a multinomial distribution over K units:

$$\mathrm{P}(\boldsymbol{x}|\boldsymbol{p}, N) = \frac{N!}{x_1!\, x_2! \ldots x_K!} \prod_{i=1}^{K} p_i^{x_i},$$

where $\boldsymbol{x} = (x_1, x_2, \ldots, x_K)$ is the observed counts of each thought, $N = \sum_{i=1}^{K} x_i$ is the total number of observations, and $\boldsymbol{p} = (p_1, \ldots, p_k)$ is the probability of activating each thought.

Before observing any behaviors, the problem-solver holds a prior about the probability of each thought, modeled as a Dirichlet distribution:

$$\mathrm{P}(\boldsymbol{p}|\boldsymbol{\alpha}) = \mathrm{Dir}(K, \boldsymbol{\alpha}) = \frac{1}{\mathrm{B}(\boldsymbol{\alpha})} \prod_{i=1}^{K} p_i^{\alpha_i - 1}$$

where $\boldsymbol{\alpha} = (\alpha_1, \alpha_2, \ldots, \alpha_K)$ is the parameters of the Dirichlet distribution and $\mathrm{B}(\cdot)$ is the beta function.

As Dirichlet distribution is the conjugate prior distribution of the multinomial distribution, the problem-solver's belief posterior after observing the behavior $\boldsymbol{x} = (x_1, x_2, \ldots, x_K)$ is also a Dirichlet:

$$\mathrm{P}(\boldsymbol{p}|\boldsymbol{x}, \boldsymbol{\alpha}) = \mathrm{Dir}(K, \boldsymbol{\alpha} + \boldsymbol{x}).$$

In other words, $\boldsymbol{\alpha}^{j+1} = \boldsymbol{\alpha}^j + \boldsymbol{x}^j$. The final belief is $\boldsymbol{\alpha}^T = \boldsymbol{\alpha}^0 + \sum_j \boldsymbol{x}^j$.

Offline meta-cognitive computation

After completing the problem-solving process and the "thought"-updating process concurrently, the problem-solver will perform offline computations when asked for retrospective evaluations on the problem-solving process during period $[t_j, t_{j+1})$. We now propose six metrics (Figure 2).

Observation entropy measures the uncertainty in the observed distribution of thought units: $\boldsymbol{H}(\frac{\boldsymbol{x}}{\sum \boldsymbol{x_i}})$.

Update likelihood measures how probable the observation is prospectively from time point $\boldsymbol{t_j}$:
$\boldsymbol{L}(\boldsymbol{x}|\boldsymbol{p}, \boldsymbol{\alpha^j}) = \int \mathbf{Multi}(\boldsymbol{x}|\boldsymbol{p})\mathbf{Dir}(\boldsymbol{p}|\boldsymbol{\alpha^j})\mathbf{d}\boldsymbol{p}$.

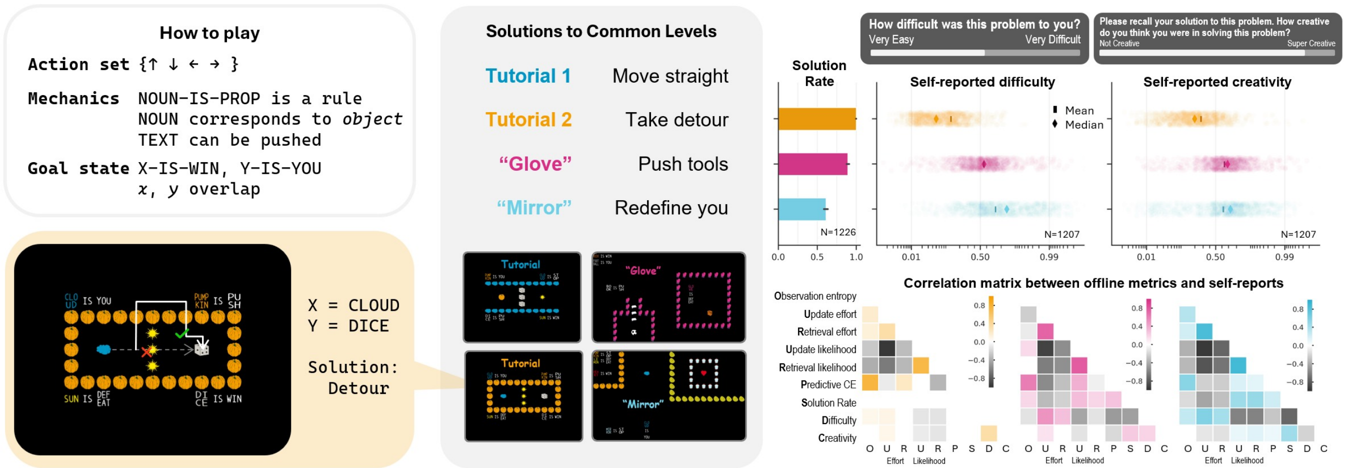

Figure 1. **(A)** Game mechanics and levels. *Bottom-left*: configuration of Tutorial 2 which requires a detour to avoid defeat. **(B)** *Left*: percentage of winning. *Middle & Right*: self-reported difficulty and creativity; overlapping dots indicate data density. **(C)** Spearman's correlation between offline metrics, solution rate, and self-reports in tutorial (*left*), "Glove" (*middle*), "Mirror" (*right*). Cells show significant correlations only (p < 0.05 after fdr correction).

Retrieval likelihood measures how probable the observation is retrospectively from the final time point $\boldsymbol{t_T}$: $\boldsymbol{L(x|p,\alpha^T)} = \int \mathbf{Multi}(\boldsymbol{x}|\boldsymbol{p})\mathbf{Dir}(\boldsymbol{p}|\boldsymbol{\alpha^T})\mathbf{d}\boldsymbol{p}$.

Update effort measures the information cost to transform the belief prospectively from $\boldsymbol{\alpha^j}$ into $\boldsymbol{\alpha^{j+1}}$: $\boldsymbol{D_{KL}}(\mathbf{Dir}(\boldsymbol{\alpha^{j+1}}), \mathbf{Dir}(\boldsymbol{\alpha^j}))$, where $\boldsymbol{D_{KL}}$ denotes the KL divergence metric.

Retrieval effort measures the information cost to transform the belief retrospectively from $\boldsymbol{\alpha^T}$ into $\boldsymbol{\alpha^j}$: $\boldsymbol{D_{KL}}(\mathbf{Dir}(\boldsymbol{\alpha^j}), \mathbf{Dir}(\boldsymbol{\alpha^T}))$.

Predictive cross-entropy (CE) measures the discrepancy between the past observation and the predictions made from the final belief $\boldsymbol{\alpha^T}$: $\int \boldsymbol{H}\left(\frac{\boldsymbol{x}}{\sum \boldsymbol{x_i}}, \boldsymbol{p}\right)\mathbf{d}\boldsymbol{p}$, given $\boldsymbol{p} \sim \mathbf{Dir}(\boldsymbol{K}, \boldsymbol{\alpha})$.

Note that not all computations are strictly offline. Metrics that do not involve the final belief can be computed during online updating.

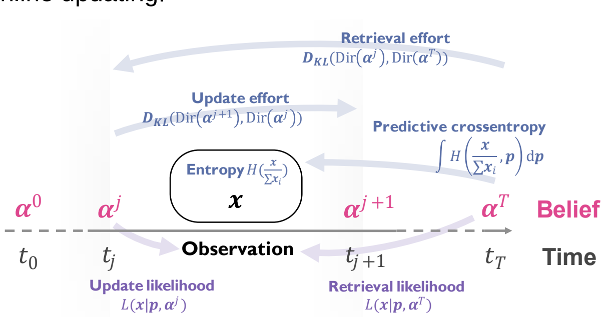

Figure 2. Offline metrics. Blue: information theoretic, Purple: related to the inference process.

Predicting self-reported difficulty

In a computer visual puzzle game adapted from "Baba Is You" (Fig 1A; Teikari, 2019), participants were given eight minutes to solve each of the five problems: two tutorials, one helper (5 variants, providing potential skills), two targets ("Glove" and "Mirror"). The order of two targets is counter-balanced between participants. After gameplay, participants made subjective evaluations on the 2nd tutorial, and two targets (Fig 1B).

A total of 1226 participants were recruited from Prolific platform (Age: 26.34 ± 5.14; Sex: 604 Females, 622 Males). 19 were excluded from analysis for failing ≥1 tutorial or having an action reaction time ≥7 minutes.

Table 1: Model comparison results (Bold: better fit).

Model		Tutorial	"Glove"	"Mirror"
Full	Adj. R^2	**0.0640**	**0.4331**	**0.6334**
	AIC	-822.6	**-318.7**	**-705.6**
Baseline	Adj. R^2	0.0630	0.3313	0.5778
	AIC	**-825.1**	-123.3	-539.3

For each participant, we assumed a weak prior $\boldsymbol{\alpha}^0 = (1, 1, \cdots, 1)$, simulated the online updating process, and computed offline metrics (Fig 1C).

To validate the hypothesis that participants' self-evaluations on past problem-solving experience could be effectively modeled by offline metrics, we compared two linear regression models. The baseline model included between-participant conditions and basic performance measures:

$$\text{Difficulty} \sim \text{order of targets} * \text{problem completion time} + \text{type of helper} + \text{problem solution}.$$

The full model extended the baseline model by including the offline metrics. To address concerns of multicollinearity, only four were added:

$$\text{Difficulty} \sim \cdots + \text{observation entropy} + \text{update effort} + \text{retrieval likelihood} + \text{predictive CE}.$$

Model comparison suggested that inclusion of offline metrics substantially enhances the model's ability to account for individual differences in difficulty reports.

References

Binz, M., & Schulz, E. (2023). Reconstructing the Einstellung Effect. *Computational Brain & Behavior*, *6*(3), 526–542. https://doi.org/10.1007/s42113-022-00161-2

da Silva Castanheira, K., Fleming, S. M., & Otto, A. R. (2021). Confidence in risky value-based choice. *Psychonomic Bulletin & Review*, *28*(3), Article 3. https://doi.org/10.3758/s13423-020-01848-y

Donnarumma, F., Maisto, D., & Pezzulo, G. (2016). Problem Solving as Probabilistic Inference with Subgoaling: Explaining Human Successes and Pitfalls in the Tower of Hanoi. *PLOS Computational Biology*, *12*(4), e1004864. https://doi.org/10.1371/journal.pcbi.1004864

Fleming, S. M., & Daw, N. D. (2017). Self-evaluation of decision-making: A general Bayesian framework for metacognitive computation. *Psychological Review*, *124*(1), 91–114. https://doi.org/10.1037/rev0000045

Lebreton, M., Bacily, K., Palminteri, S., & Engelmann, J. B. (2019). Contextual influence on confidence judgments in human reinforcement learning. *PLOS Computational Biology*, *15*(4), e1006973. https://doi.org/10.1371/journal.pcbi.1006973

Meyniel, F., Schlunegger, D., & Dehaene, S. (2015). The Sense of Confidence during Probabilistic Learning: A Normative Account. *PLoS Computational Biology*, *11*(6). https://doi.org/10.1371/journal.pcbi.1004305

Prat-Carrabin, A., Meyniel, F., Tsodyks, M., & Silveira, R. D. (2021). Biases and Variability from Costly Bayesian Inference. *Entropy*, *23*. https://doi.org/10.3390/e23050603

Rahnev, D., Koizumi, A., McCurdy, L. Y., D'Esposito, M., & Lau, H. (2015). Confidence Leak in Perceptual Decision Making. *Psychological Science*, *26*(11), 1664–1680. https://doi.org/10.1177/0956797615595037

Sanders, J. I., Hangya, B., & Kepecs, A. (2016). Signatures of a Statistical Computation in the Human Sense of Confidence. *Neuron*, *90*(3), 499–506. https://doi.org/10.1016/j.neuron.2016.03.025

Teikari, A. (2019). *Baba Is You* [Computer software]. https://www.hempuli.com/baba/

Ting, C.-C., Salem-Garcia, N., Palminteri, S., Engelmann, J. B., & Lebreton, M. (2023). Neural and computational underpinnings of biased confidence in human reinforcement learning. *Nature Communications*, *14*(1), Article 1. https://doi.org/10.1038/s41467-023-42589-5

Zenon, A., Solopchuk, O., & Pezzulo, G. (2019). An information-theoretic perspective on the costs of cognition. *Neuropsychologia*, *123*, 5–18. https://doi.org/10.1016/j.neuropsychologia.2018.09.013