Learning task rule updating strategies requires extensive practice

Shengjie Xu (Shengjie.Xu@ugent.be)
Department of Experimental Psychology, Ghent, Belgium

Tanya Wen (2tanyawen@gmail.com)
Meta Reality Labs, Redmond, U.S.A
Center for Cognitive Neuroscience, Duke University, U.S.A

Tobias Egner (tobias.egner@duke.edu)
Center for Cognitive Neuroscience, Duke University, U.S.A
Department of Psychology and Neuroscience, Duke University, U.S.A

Tom Verguts (Tom.Verguts@ugent.be)
Department of Experimental Psychology, Ghent, Belgium

Senne Braem (Senne.Braem@ugent.be)
Department of Experimental Psychology, Ghent, Belgium

Abstract

People can adjust how fast they update task rules, depending on the volatility of their environment. We investigated whether this adaptivity is primarily driven by recently experienced volatility in task demands, or can also be shaped by learned, environment-specific associations with expected levels of volatility. We trained participants on a Wisconsin Card Sorting Task where different environments required different speeds of task rule updating. Initially, participants updated strategies depending on the most recently experienced levels of volatility (Experiment 1). However, after extensive (four days) training (Experiment 2), participants also developed environment-specific associations. Our findings provide important insights in how people learn to regulate cognitive flexibility.

Keywords: cognitive control; cognitive flexibility; reinforcement learning; multi-day training

Introduction

Contemporary theories on cognitive control emphasize that the regulation of cognitive control requires learning (Abrahamse et al., 2016; Braem et al., 2024; Braem & Egner, 2018; Egner, 2014, 2023; Verguts & Notebaert, 2009), specifically meta-learning (Griffiths et al., 2019; Wang, 2021). Humans can learn to set up control parameters, sometimes referred to as meta-control (Eppinger et al., 2021), by associating environmental features to different control processes (Chiu & Egner, 2017; Xu et al., 2024). This way, learned control settings can be evoked when humans revisit these environments, allowing for faster changes in task updating.

Evidence from a recent task switching study suggests that such learning of environment-specific strategies might require extensive task experience (Xu et al., 2024). However, traditional task switching paradigms leave little uncertainty about the current need for task rule updating, potentially demotivating humans to learn environment-specific statistics. Contrasting a high versus low volatility environment in a Wisconsin Card Sorting Task, Wen and colleagues (2023) recently found that participants used different learning rates for task updating in high versus low task volatility, and generalized these to a test phase with neutral volatility. However, this finding could reflect a learning and memorizing of environment-specific control settings, or a carry-over effect of recently experienced volatility. Here, we manipulated volatility in a Wisconsin Card Sorting Task within-subjects to investigate if humans can learn to strategically regulate cognitive flexibility in response to different environments. We hypothesized that participants first rely on locally experienced differences in demands for task rule changes, while learned, environment-specific task updating strategies can only be observed after extensive training.

Methods

Participants. We recruited 65 and 62 participants in Experiment 1 and 2. After excluding participants with accuracy below 65%, we had 55 participants (43 females) and 55 participants (49 females) in Experiment 1 and 2.

Task and procedure. Every trial, participants chose between two cards to match the reference card on top (Fig. 1A). Each card varied across four features: color, filling, item number, and shape, and shared only one feature value with the reference card. Two feature dimensions were randomly selected as task-relevant, counterbalanced across participants. Feedback was 80% accurate, with matching rules changing every 10 trials in the high volatility environment, and every 30 trials in the low volatility environment. Each environment was associated to different backgrounds: a wood or stone table (Fig. 1B). In probe blocks, rules changed every 20 trials on either table to test if subjects learned and used the environment-volatility associations. Participants were trained for either one or four days before testing in Experiment 1 and 2, respectively (Fig. 1C-D).

Reinforcement learning model. We fitted choices to a dual-rates (DR) model, with individual learning rates for positive and negative feedback. We used a hierarchical framework, to account for both fixed and random effects within the data structure for parameter estimates.

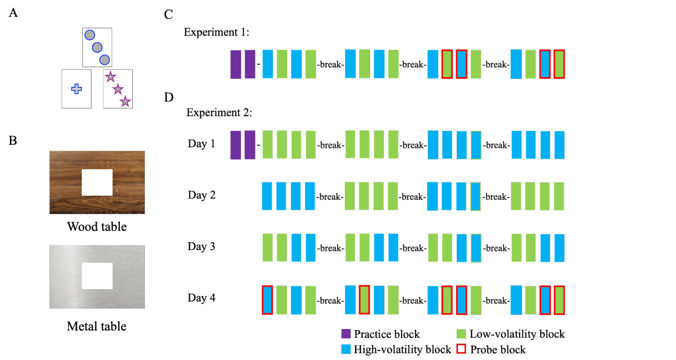

Figure 1: Wisconsin cart sorting task paradigm. (A) Example cards presented on each trial. (B) Table pictures. (C) Overview of Experiment 1. (D) Overview of Experiment 2.

Results

Learning blocks. Positive learning rates were significantly higher in the high-volatility environment in training blocks (Figure 2A), posterior probability (*p*_post) = 0.019, but not the negative learning rates, *p*_post = 0.908, in Experiment 1. Similarly, after four days training, participants robustly applied higher positive learning rates in response to high volatility (Figure 2B), *p*_post < 0.001. However, there was no difference between negative learning rates in two environments, *p*_post = 0.261, in Experiment 2.

Probe blocks. Critical to our hypothesis, we expected that the observed parameter patterns in the learning blocks would extend to the probe blocks. There was no difference between positive learning rates in two environments, *p*_post = 0.176, nor the negative learning rates, *p*_post = 0.691 in Experiment 1. However, in Experiment 2, negative learning rates were higher in the high volatility environment, *p*_post = 0.011, but no difference in positive learning rates, *p*_post = 0.199.

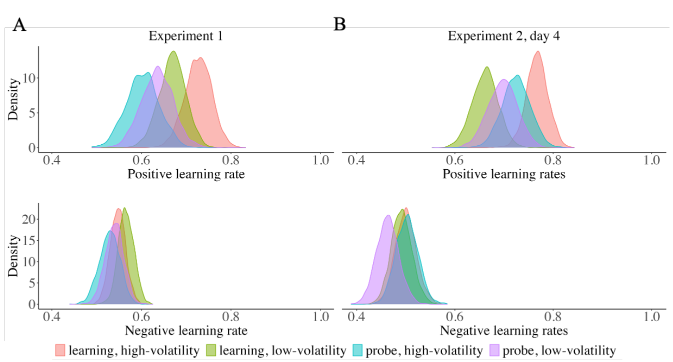

Figure 2: Model parameter estimates in Experiment 1 (A) and Experiment 2, day 4 (B).

Learning rates evolution over learning. We further observed that the size of positive learning rates was consistent over learning (Fig 3), all *p*_posts > 0.05, while the negative learning rates decreased on Day 2, *p*_post < 0.001 and became stable as of Day 3 *p*_post = 0.343. Notably, the difference between positive learning rates in two volatility environments reduced over learning (95% highest density interval (HDI), Day 1: [0.06, 0.15]; Day 2: [0.02, 0.08]; Day 3: [0.00, 0.07]), while the difference between negative learning rates increased (95% HDI, Day 1: [-0.02, 0.05]; Day 2: [0.01, 0.07]; Day 3: [0.02, 0.07]).

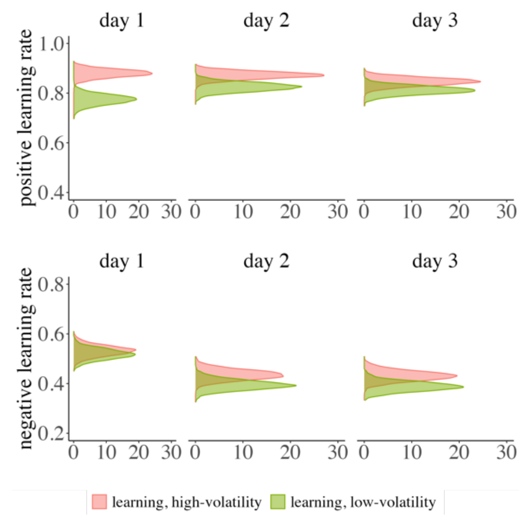

Figure 3: Model parameter estimates in Experiment 2, day 1 to 3.

Discussions

The current study aimed to examine the environment-specific regulation of task rule updating strategies using a probabilistic Wisconsin Card Sorting Task. Across two experiments, we observed that participants were able to learn associations between different strategies and co-occurring environmental features, i.e., table pictures, which enabled participants to use environmental features to regulate task rule updating strategies. In sum, our study investigated behavioral mechanisms underlying the environment-specific regulation of task updating strategies, providing empirical evidence to support a critical contribution of multi-day training schemas to form environment-control associations.

References

Abrahamse, E., Braem, S., Notebaert, W., & Verguts, T. (2016). Grounding cognitive control in associative learning. Psychological Bulletin, 142(7), 693–728. https://doi.org/10.1037/bul0000047

Braem, S., Chai, M., Held, L. K., & Xu, S. (2024). One cannot simply 'be flexible': Regulating control parameters requires learning. Current Opinion in Behavioral Sciences, 55, 101347. https://doi.org/10.1016/j.cobeha.2023.101347

Braem, S., & Egner, T. (2018). Getting a Grip on Cognitive Flexibility. Current Directions in Psychological Science, 27(6), 470–476. https://doi.org/10.1177/0963721418787475

Chiu, Y.-C., & Egner, T. (2017). Cueing cognitive flexibility: Item-specific learning of switch readiness. Journal of Experimental Psychology: Human Perception and Performance, 43(12), 1950–1960. https://doi.org/10.1037/xhp0000420

Egner, T. (2014). Creatures of habit (and control): A multi-level learning perspective on the modulation of congruency effects. Frontiers in Psychology, 5, 1247. https://doi.org/10.3389/fpsyg.2014.01247

Egner, T. (2023). Principles of cognitive control over task focus and task switching. Nature Reviews Psychology, 1–13. https://doi.org/10.1038/s44159-023-00234-4

Eppinger, B., Goschke, T., & Musslick, S. (2021). Meta-control: From psychology to computational neuroscience. Cognitive, Affective, & Behavioral Neuroscience, 21(3), 447–452. https://doi.org/10.3758/s13415-021-00919-4

Griffiths, T. L., Callaway, F., Chang, M. B., Grant, E., Krueger, P. M., & Lieder, F. (2019). Doing more with less: Meta-reasoning and meta-learning in humans and machines. Current Opinion in Behavioral Sciences, 29, 24–30. https://doi.org/10.1016/j.cobeha.2019.01.005

Verguts, T., & Notebaert, W. (2009). Adaptation by binding: A learning account of cognitive control. Trends in Cognitive Sciences, 13(6), 252–257. https://doi.org/10.1016/j.tics.2009.02.007

Wang, J. X. (2021). Meta-learning in natural and artificial intelligence. Current Opinion in Behavioral Sciences, 38, 90–95. https://doi.org/10.1016/j.cobeha.2021.01.002

Wen, T., Geddert, R. M., Madlon-Kay, S., & Egner, T. (2023). Transfer of Learned Cognitive Flexibility to Novel Stimuli and Task Sets. Psychological Science, 34(4), 435–454. https://doi.org/10.1177/09567976221141854

Xu, S., Simoens, J., Verguts, T., & Braem, S. (2024). Learning where to be flexible: Using environmental cues to regulate cognitive control. Journal of Experimental Psychology: General, 153(2), 328–338. https://doi.org/10.1037/xge0001488