C^2Tutor: Helping People Learn to Avoid Present Bias During Decision Making

Calarina Muslimani[1,2(✉)], Saba Gul[1], Matthew E. Taylor[1,2], Carrie Demmans Epp[1], and Christabel Wayllace[1,2]

[1] University of Alberta, Edmonton, AB T6G 2R3, Canada
{musliman,matthew.e.taylor,demmanse,wayllace}@ualberta.ca
[2] Alberta Machine Intelligence Institute (Amii), Edmonton, AB T5J 3B1, Canada

Abstract. Procrastination can harm many aspects of life, including physical, mental, or financial well-being. It is often a consequence of people's tendency to prefer immediate benefits over long-term rewards (i.e., present bias). Due to its prevalence, we created C^2Tutor, an intelligent tutoring system (ITS) that can potentially reduce procrastination habits by teaching planning strategies. C^2Tutor teaches people how to make decisions aligned with long-term benefits. It will discourage present bias behavior while allowing for differences in user cognitive abilities. Our study found that C^2Tutor encourages far-sighted behavior while reducing maladaptive planning-strategy use.

Keywords: Procrastination · Behavior Change · Lifelong Learning

1 Introduction

Procrastination, the action of voluntarily and unnecessarily delaying something despite adverse consequences, is a failure in self-regulation, where a person is unable to regulate their thoughts, emotions, impulses, or behavior [1]. Several theories from economics and psychology suggest that procrastination can be a consequence of people's tendency to prefer immediate benefits over more critical long-term rewards, a phenomenon known as present bias [12].

Despite its prevalence, there is a substantial shortage of mental healthcare professionals [11]. This shortage creates an opportunity and the need for more virtual interventions. Recent work has developed intelligent tutoring systems (ITSs) to teach people strategies for reducing present bias behavior [2,8].

Therefore, the present work aims to develop an ITS, C^2Tutor, that teaches planning strategies to reduce present bias. Unlike prior work, we developed a flexible backwards planning strategy that discourages present bias behavior without assuming each individual has the same cognitive ability. We also developed a teaching methodology for C^2Tutor based on theories of formative feedback, meta-cognitive reinforcement learning, and behavior change. To that end, we performed a study to investigate the use of C^2Tutor in a simulated environment in which present bias is likely to occur. We address the following questions:

N. Wang et al. (Eds.): AIED 2023, LNAI 13916, pp. 733–738, 2023.
https://doi.org/10.1007/978-3-031-36272-9_64

1. Do people learn our flexible backwards planning strategy using C^{2}Tutor?
2. Can C^{2}Tutor improve people's decision-making in a simulated environment?
3. Do people find C^{2}Tutor useful for improving decision-making?

2 Literature Review and Research Gaps

Present bias or near-sightedness is a fundamental factor influencing how much people procrastinate [10,12]. Prior work [2,8] attempted to reduce procrastination using an ITS to discourage present bias in the Mouselab-MDP paradigm [3] (Fig. 1 Left). These ITSs teach an optimal planning strategy that obtains the best possible tradeoff between the expected cost of making a decision and the expected utility of that decision [7]. To teach this strategy, these systems use meta-cognitive reinforcement learning, where positive or negative reinforcement is given based on how a decision was made rather than what decision was made.

Although this prior work showed the potential of using ITSs to discourage present bias, the optimal planning strategy and teaching methodology were limited. A primary weakness is their reliance on a single optimal planning strategy for all individuals. This choice has two underlying assumptions: (1) all learners have the same cognitive abilities, and (2) the environment is consistent. For example, following this optimal planning strategy, users will start reviewing the rewards of different options in the future and finalize the plan when they find the first maximum possible reward. Any further exploration is discouraged. Therefore, this strategy is not flexible to how much exploration an individual might want to perform when planning. Moreover, relying on meta-cognitive reinforcement learning does not fully account for what is known about feedback delivery [5,9]. For example, these prior ITSs [2,8] used praise as the primary source of positive reinforcement, even though praise can discourage learning [5].

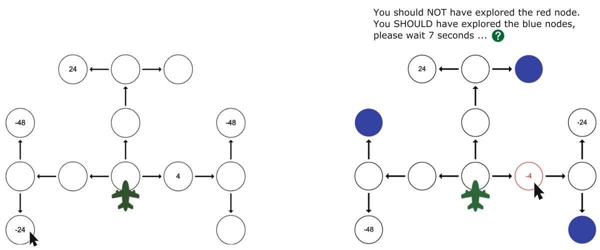

Fig. 1. Mouselab-MDP - Participants click to reveal the value at future states (Left). C^{2}Tutor incorporated in Mouselab-MDP (Right).

3 C^{2}Tutor - Reducing Present Bias

To teach people how to plan, we must be able to observe how people make decisions. To that end, we employed an implementation of Mouselab-MDP since it externalizes people's planning behavior [3] - See Fig. 1 Left. Mouselab-MDP represents planning problems using simple directed acyclic graphs (DAGs). Along each path, there are nodes with hidden rewards. Users click on nodes to reveal their values. After revealing as many nodes as desired, the individual chooses the path they think is best (i.e., the one with the highest total reward). A node-reveal corresponds to a user evaluating the quality of a future state, which is a fundamental cognitive operation in planning. As this process requires mental effort, there is a cost of 1 per click operation. The participant's goal is to maximize their cumulative reward for a decision-making task.

To represent real-life settings where present bias is commonly a factor, we used a three-level DAG in which the rewards are randomly drawn from the following distributions: In level 1, closest to the root, $\{-4, -2, +2, +4\}$; in level 2 $\{-8, -4, +4, +8\}$; and in level 3 $\{-48, -24, +24, +48\}$. The order in which nodes are explored reveals the planning strategy used. Common near-sighted strategies include only inspecting the immediate nodes (i.e., those in level 1) or inspecting an immediate node first. Other poor planning strategies include zero-planning and random exhaustive search. In zero-planning, no nodes are inspected and a random path is taken. For random exhaustive search, all nodes are inspected before taking the path with the maximum cumulative reward [6]. For a full description of all planning strategies, we refer readers to Jain et al. [6].

C^2Tutor aims to teach a general planning strategy that does not assume individuals' cognitive limits are fixed. To that end, we develop our flexible backwards planning strategy that explicitly considers future rewards first and immediate rewards last. As in typical backwards planning, this strategy explores node rewards by level, going from the leaf nodes to the root. Thus, it first explores the reward of all leaf nodes. If there is not enough information to choose a path, it explores nodes of the upper levels. As the range of rewards is higher for the leaf nodes and decreases towards the root, one can prune paths with low rewards.

A limitation of exhaustive backwards planning (i.e., reviewing all nodes from the leaves to the root) is that it can be time-consuming or infeasible for large graphs. For example, if a person has limited time or has more cognitive constraints, they can make a decision (i.e., choose a path) based solely on reviewing the rewards of the leaf nodes. Otherwise, they can plan backwards for more than one step. Critically, C^2Tutor does not assume there is one optimal planning strategy, making our planning strategy flexible to individual differences.

C^2Tutor also provides theoretically-grounded formative feedback to support behavior change. It uses twice as many practice rounds as prior systems [2,8] since an individual's performance on a task is directly related to the amount of deliberate practice performed [4]. To guide design decisions about the timing and types of feedback provided, we used meta-cognitive reinforcement learning and a computer-based feedback framework [9]. To determine the type of feedback to provide, we first assumed that learners have low prior knowledge of effective planning strategies. Therefore, we provided correct response and response-contingent feedback after each planning operation [9]. To alleviate the potential for unnecessarily elaborate feedback, we gave learners control over response-contingent feedback. Learners could choose whether to view this type of feedback (see the green question mark in Fig. 1 Right). This allows learners to obtain clarification if needed or wanted. We also included negative reinforcement as a time penalty because it has been effective when teaching planning strategies [2,8].

4 Experimental Design

Following an institutional ethics review, we conducted a between-subjects study with two conditions: a Tutor experimental condition and a No Tutor control.

We had 15 participants (10 male and 5 female): 7 in the Tutor condition and 8 in the No Tutor condition. Age ranges were collected: 7 were aged 18–24, 4 were 25–30, 2 were 31–40, and 2 were aged 41–50. Participants self-identified their membership in racial groups: 6 were Asian, 5 were White, 3 were Middle Eastern or North African, and 1 was multi-racial.

Following consent procedures, participants were given instructions on how to complete tasks in Mouselab-MDP. In the Tutor condition, participants also read an instructional page describing C^{2}Tutor. Participants were then given the instructions quiz, for which they had to answer all questions correctly to continue. Qualifying participants interacted with the Mouselab-MDP for 1 pre-test round, 20 practice rounds, and 10 test rounds. In the Tutor condition, participants had access to C^{2}Tutor during the practice rounds only. After the test rounds, participants completed an ITS preference and demographics questionnaire.

For all rounds, we collected data on which nodes users revealed, the order of node reveals, and which path they chose. We used this information and the planning strategies defined by Jain et al. [6] to identify the strategies used. To compare performance across groups, we used a one-tailed Mann-Whitney U Test ($\alpha = .05$). We report r as the effect size.

5 Findings

Do People Learn the Backwards Planning Strategy Using C^{2}Tutor? Figure 2 Right shows the number of times participants used various planning strategies during the test rounds. Those in the Tutor condition used our backwards planning strategy close to twice as much ($U = 43.5$, $p = .038$, $r = 0.47$). Meanwhile, participants in the No Tutor condition had significantly higher use of maladaptive planning strategies during both the training ($U = 8.5$, $p = .012$, $r = 0.59$) and the test rounds ($U = 14.0$, $p = .021$, $r = 0.545$).

Can C^{2}Tutor Improve People's Decision-Making in a Simulated Environment? We compared the participants' cumulative training and test scores across conditions. We found a significant difference ($U = 52.0$, $p = .003$, $r = 0.72$) in scores during training between the Tutor ($M = 735.57$, $SD = 39.24$) and No Tutor conditions ($M = 560.25$, $SD = 278.61$). None was not found during testing ($U = 33.0$, $p = .301$, $r = 0.15$) across the Tutor ($M = 466.43$, $SD = 15.56$), and No Tutor ($M = 403.88$, No Tutor $SD = 124.08$) groups.

Across all rounds, participants in the Tutor condition achieved significantly higher scores compared to those in the No Tutor condition ($U = 658.0$, $p = .006$, $r = 0.32$). This highlights that using C²Tutor resulted in consistently higher performance during both the training and test phases (Fig. 2 Left).

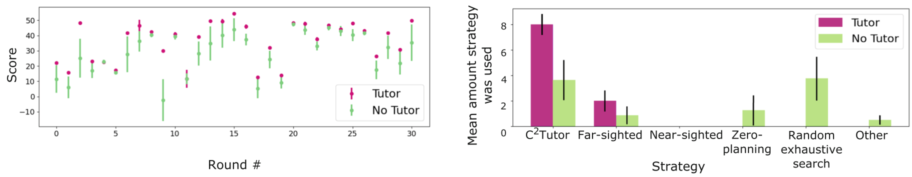

Fig. 2. Left: Average score per round. Right: Average number of times participants used a planning strategy (test rounds). Error bars show standard error for both plots.

Do People Find C²Tutor Useful for Improving Decision-Making? Overall, participants in the Tutor condition found C²Tutor helpful. They also found the feedback easy to understand and use (Fig. 3). Moreover, approximately a third of participants believed that the concepts they learned during the study would be helpful for future decision-making tasks, suggesting the potential for the learned planning strategies to transfer to future decision-making tasks.

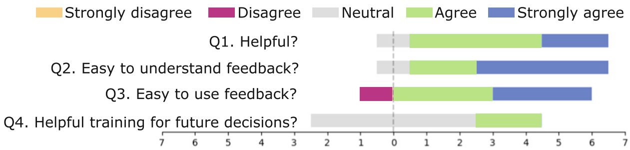

Fig. 3. Participant perceptions of C²Tutor.

6 Conclusion and Future Work

In this work, we sought to create an intervention that has the potential to reduce procrastination habits. To that end, we developed an ITS that teaches people a planning strategy that can reduce present bias tendencies. Unlike prior work, C²Tutor teaches a planning strategy that discourages present bias behavior in a manner that enables individualization. In addition, C²Tutor uses a principled teaching methodology that is based on theories of formative feedback, metacognitive reinforcement learning, and behavior change.

We found that participants who practiced decision-making with C^2Tutor were significantly more likely to use our flexible backwards planning strategy. Meanwhile, participants who did not use our ITS were significantly more likely to use maladaptive strategies. While these results are promising, we note that C^2Tutor only focused on reducing present bias in a simple simulated environment. Future work should consider how educational technologies can be used to tempter other contributors to procrastination in more complex environments.

Acknowledgements. This work was supported in part by grants from Amii; a Canada CIFAR AI Chair, Amii; Compute Canada; Huawei; Mitacs; and NSERC.

References

1. Baumeister, R.F., Heatherton, T.F., Tice, D.M.: Losing control: how and why people fail at self-regulation (1994)
2. Callaway, F., et al.: Leveraging artificial intelligence to improve peoples planning strategies. In: Proceedings of the National Academy of Sciences, vol. 119, no. 12 (2022)
3. Mouselab-MDP: A new paradigm for tracing how people plan (2017)
4. Ericsson, K.A., Krampe, R.T., Tesch-Römer, C.: The role of deliberate practice in the acquisition of expert performance. Psychol. Rev. **100**(3), 363–406 (1993)
5. Hattie, J., Timperley, H.: The power of feedback. Rev. Educ. Res. **77**(1), 81–112 (2007)
6. Jain, Y.R., et al.: A computational process-tracing method for measuring people's planning strategies and how they change over time. Behav. Res. **55**, 2037–2079 (2023)
7. Lieder, F., Griffiths, T.L.: Resource-rational analysis: understanding human cognition as the optimal use of limited computational resources. Behav. Brain Sci. **43**, e1 (2020)
8. A cognitive tutor for helping people overcome present bias (2019)
9. Mason, B.J., Bruning, R.: Providing feedback in computer based instruction: what the research tells us. Tech. rep. 9 (2001)
10. O'Donoghue, T., Rabin, M.: Doing it now or later. Am. Econ. Rev. **89**(1), 103–124 (1999)
11. Psychologists struggle to meet demand amid mental health crisis. Tech. rep. DOE-SLC-6903-1. American Psychological Association (2022)
12. Steel, P., König, C.J.: Integrating theories of motivation. Acad. Manag. Rev. **31**(4), 889–913 (2006)