The Effects of Aging and Cognitive Tutoring on Planning Abilities

A Cognitive Science Honors Thesis

By Priyam Das

Advisors: Falk Lieder and Thomas Griffiths

Second Reader: William Jagust

Introduction

Many cognitive abilities are known to deteriorate with normal aging. Cognitive abilities can be generally divided into two types — crystallized intelligence and fluid intelligence. Crystallized intelligence refers to using previously acquired knowledge to complete a task, such as taking a vocabulary test or recalling historical events. Fluid intelligence refers to real time cognitive processing and manipulation of information for the task at hand. Previous studies have shown that while abilities relying on crystallized intelligence improve before reaching a plateau at age 60, fluid intelligence steadily declines from the ages of 20 to 80 years old (Murman, 2015). Executive function, memory, and processing speed all fall under the realm of fluid intelligence (Harada, Natelson Love, & Triebel, 2013). In particular, executive functions include abilities such as decision making, problem solving, and planning. These are all necessary in order to live functionally and independently. Thus, it is important to be able to maintain these cognitive abilities throughout life. However, effective methods for combating the decline of fluid intelligence or even improving fluid intelligence in later life have yet to be discovered. Recent work has leveraged resource-rational analysis (Griffiths, Lieder, & Goodman, 2015), which improves cognitive models by accounting for limited cognitive resources, to develop cognitive tutors in the domain of planning (Lieder, 2018). These cognitive tutors have been successful in teaching people how to plan routes more efficiently. In Experiment 1, we examine the efficacy of the cognitive tutors for a range of ages in the hopes that this method may be used for later life maintenance of cognitive abilities. In Experiment 2, we investigated the planning strategies used by different age groups to determine whether this is also something that changes with age. We found that older adults benefit more from using the cognitive tutors than younger adults. We also found that different age groups use certain strategies at different rates, but that both younger and older adults increasingly tend to adopt the optimal strategy over time. Another strategy, depth-first search, is equally employed by older adults in the given planning paradigm, whereas there is no other contender for favorite strategy in the younger group. The changes in planning strategy usage provide evidence of learning during the task and suggest that support throughout the learning process would be beneficial. Our results show the promise of using resource-rational cognitive tutors to improve planning abilities in older adults and shed light on the planning strategies in use by different ages, which can help further improve cognitive tutors.

Background

Planning falls under the domain of executive control, which is known to decline with increasing age (Murman, 2015). Previous studies of the effect of aging on planning have used various paradigms, from planning the steps to solve a puzzle to more open ended tasks such as planning the subtasks necessary for achieving a goal. One popular paradigm in the former category is the Tower of London task, which involves participants manipulating discs on pegs until the setup matches a goal formation. Studies using this task have shown that older adults need more time to formulate plans and use more moves to solve the Tower of London tasks than younger adults, but that increasing the complexity of these tasks does not pose particular problems for older adults. Furthermore, no difference in planning strategies used has been found between older and younger adults (Gilhooly, Phillips, Wynn, Logie, & Sala, 1999). Other studies which use the Tower of Hanoi task and the Zoo Map Test have found that older individuals have trouble with formulating plans and updating them given feedback (Allain et al., 2004; Sorel & Pennequin, 2007). In this study, we plan to use the Mouselab-MDP paradigm (Callaway, Lieder, Das, Gul, Krueger, & Griffiths, 2018) which asks participants to collect information in order to select the path that maximizes the participant's reward. The advantages of using the Mouselab-MDP paradigm over previously used paradigms are two-fold. First, many of the differences in planning performance amongst different age groups is thought to arise from differences in working memory capacity and efficiency (Phillips, MacLeod, & Kliegel, 2004). The Mouselab-MDP paradigm does not require the use of working memory, since all of the information required for planning remains visible to the participant throughout the trial. Second, the Mouselab-MDP paradigm allows us to observe a participant's entire planning process via the actions they take to gather the information needed to formulate plans. All planning must take place before executing any plans and there is little opportunity to make changes to that plan once execution has begun. Lieder (2018) also developed cognitive tutors for use with the Mouselab-MDP paradigm that have been shown to significantly improve planning performance in people. Participants who received feedback from the cognitive tutor during training trials scored over 10 points higher on each test trial compared to those who did not receive any cognitive tutoring.

Questions and Hypotheses

We set out to investigate the effects of aging on learning how to plan. Specifically, we wanted to know whether people in certain age groups learn faster than others and how learning aids, such as feedback, affect this learning process. We were also curious as to whether different age groups rely on different strategies when approaching tasks that require planning, such as planning paths. Finally, we wanted to investigate whether aging has an effect on a person's ability to plan optimally. Given previous findings (e.g. Allain et al., 2004; Gilhooly et al., 1999; Murman, 2015), we hypothesize that the planning performance of older adults is worse than that of younger adults, but that this performance can be improved with the use of directed feedback. Furthermore, we expect to see little difference in the types of planning strategies utilized by both older and younger adults.

Experiment 1

In Experiment 1, we looked at the learning rates of people across a range of ages and whether feedback given by cognitive tutors had any effect on learning during the planning task.

Methods

Procedure

In order to determine whether there was any relationship between a person's age and their ability to plan, we presented participants with a variation of the Mouselab-MDP paradigm (Callaway and Lieder et al., 2018). The Mouselab-MDP paradigm consists of 3-step path planning problems with rewards at each step that are initially hidden (Figure 1). Participants can uncover rewards either by moving along the nodes in the path or clicking on a node and paying $1 to reveal the reward at that step before any moves are made. These clicks serve to elucidate where a participant is planning to move. To encourage planning and disincentivize speeding through the experiment, participants are required to spend a minimum of 7 seconds on each trial before being allowed to move onto the next. A participant collects rewards as they move along a path but once a path is chosen, the participant isn't allowed to backtrack and can only move forward. The sum of the rewards gained minus the cost of uncovering any rewards before moving is added to the participant's score, represented as a dollar amount. The goal of the participant is to maximize the amount of money they have by the end of the experiment.

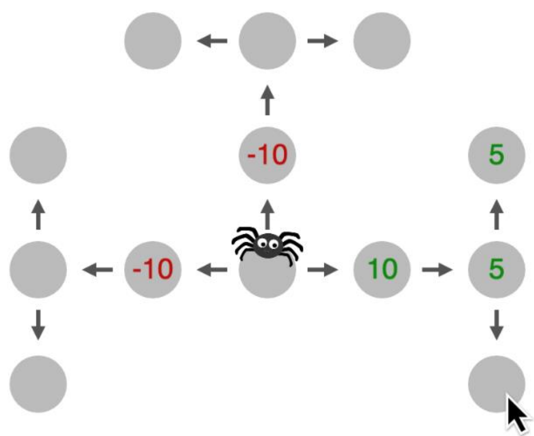

Figure 1: A Mouselab-MDP trial. The visible rewards indicate that those nodes have been clicked. The completely gray nodes have hidden rewards, save for the center, starting node.

In this version of Mouselab-MDP, the range of rewards increased at each step. At the first step, the distribution of rewards was Uniform({-4, -2, 2, 4}), at the second step it was Uniform({-8, -4, 4, 8}), and at the third step, Uniform({-48, -24, 24, 48}). Such an environment favors a backwards planning strategy, where the last node in a path is inspected first since it holds the most information about the expected rewards from that path. Participants were randomly assigned to two groups and counterbalanced: a control group (n=55) and a group that received feedback on their planning performance (n=56). Both groups went through a pretest trial, 10 training trials, and 20 test trials. Feedback was only presented during the 10 training trials to the feedback group.

Feedback was constructed so as to encourage participants to adopt a backwards planning strategy; if participants did not inspect an outermost node first, they were penalized with a delay in the experiment and would see a message asking them to inspect the outermost nodes. These nodes were also highlighted for the duration of the delay. The feedback also encouraged participants to continue inspecting outermost nodes until a high enough reward was uncovered. If a participant started to follow a path prematurely, they were similarly penalized with a delay and uninspected outermost nodes were highlighted with a message informing the participant that inspecting one of those nodes would have been a better action than moving. If a participant performed the correct action, they would simply see a message stating "Good job!" and were able to continue with the experiment without any delays.

These delays were computed by first modeling the problem of deciding how to plan as a meta-level MDP and then solving that MDP using backwards induction. The meta-level MDP is defined as $M_{meta} = (\mathcal{B}, \mathcal{A}, \mathcal{T}, r_{meta})$, where the parameters are belief states, meta-level actions, the transition probability matrix between belief states, and the meta-level rewards, respectively. The belief states represent the reward of each transition using a normal distribution. The metalevel actions are the clicks that a participant can make to uncover the reward at a node, plus the action to terminate planning and begin moving along the path believed to yield the highest expected return. The transition probabilities of a click revealing a certain reward are drawn from the probability density of the normal distribution. The metalevel reward function is the negative click cost for all click actions in every belief state, and the maximum total of the average rewards along each path for the planning termination action (for details, see Lieder, 2018). Upon solving this MDP, the optimal policy is known and the actions that the optimal policy would take in each state can be compared to the actions that participants take in the same state. The length of the feedback delay is calculated as round(2 + score(b_t, c_t)) seconds, where score(b_t, c_t) is defined as $\hat{Q}_{meta}(b, c) - \ max_c \ \hat{Q}_{meta}(b, c)$.

At the end of the experiment there was a comprehension quiz to check that participants understood the experiment directions, and a short survey which asked for the participant's age, among other questions about the participant's planning process.

Participants

We recruited 119 participants from Amazon Mechanical Turk. The data from 8 participants were excluded from analyses because these participants either did not provide their age or we were unable to recover which condition they were in due to a technical error, leaving us with 111 participants. The base pay for the experiment was \$0.75 plus an additional bonus of \$0.01 per every \$5.00 the participant earned during the test block. Participants' ages ranged from 20 to 68 (average age 34.7 ± 9.8 years).

Results

To begin, we compared the average performance on the test trials between the group that received feedback during training and the control group. We found that the feedback group scored 36.16 points on average compared to the control group which scored 24.21 points on average, confirming earlier findings (Lieder, 2018) that people who receive training with the cognitive tutors perform better on the planning task than those who do not.

A multi-way between-subjects ANOVA was conducted to explore the effects of age, trial index, and feedback on the score for each trial. This analysis revealed a significant interaction between the trial index and age ($F(1, 1102) = 6.85$; $p = 0.029$) which supports the hypothesis that the rate of improvement increases with age when given feedback and helps answer the question of whether feedback affects the learning rate of older adults. Another significant effect was that of age ($F(1, 1102) = 14.43$; $p < 0.001$). No other effects were significant. Since we were curious as to whether older adults improved more than younger adults because they started out with lower performance, we added pretest score as another independent variable and ran ANOVA again. As before, there were significant effects for age ($F(1, 1094) =14.44$; $p = 0.001$) and the interaction between trial index and age ($F(1,1094) = 6.85$; $p = 0.014$). Trial index also had a significant effect ($F(1,1094) = 14.27$; $p = 0.041$). However, the results did not show any significant interaction between the pretest score and age ($F(1, 1094) = 0.099$; $p = 0.993$) and no other effects were significant.

An independent-samples t-test was also conducted to compare average performance over test trials in younger adults and older adults. We used the median age of our participants, 33, to determine these groups (any participant older than 33 was included in the older group). There was a significant difference in the average scores for the younger adults ($M = 33.11$, $SD = 12.70$) and the older adults ($M = 27.11$, $SD = 17.27$); $t(109) = 2.08$, $p = 0.04$. Looking further at age and average score over all test trials, there was a significant negative correlation in the control condition ($\rho(53) = -0.335$; $p = 0.012$) but not in the feedback condition ($\rho(54) = -0.016$; $p > 0.9$). However, across both groups, there was a significant negative correlation between age and the pretest score ($\rho(109) = -0.217$; $p = 0.022$). This contradiction with our ANOVA result is likely due to the fact that each participant was only given one pretest trial and thus the data may be very noisy.

Next, we looked at improvement by subtracting a participant's pretest score from their average score over their last five test trials. The average of the last five test trials was used as a proxy for the participant's final absolute performance since participants still showed signs of learning throughout the test trials (see Figure 2). We found a significant positive correlation for the feedback group (Figure 3; $\rho(54) = 0.264$; $p = 0.0497$) while the correlation for the control group was insignificant (Figure 3; $\rho(53) = -0.046$; $p > 0.7$), revealing that feedback was beneficial for learning.

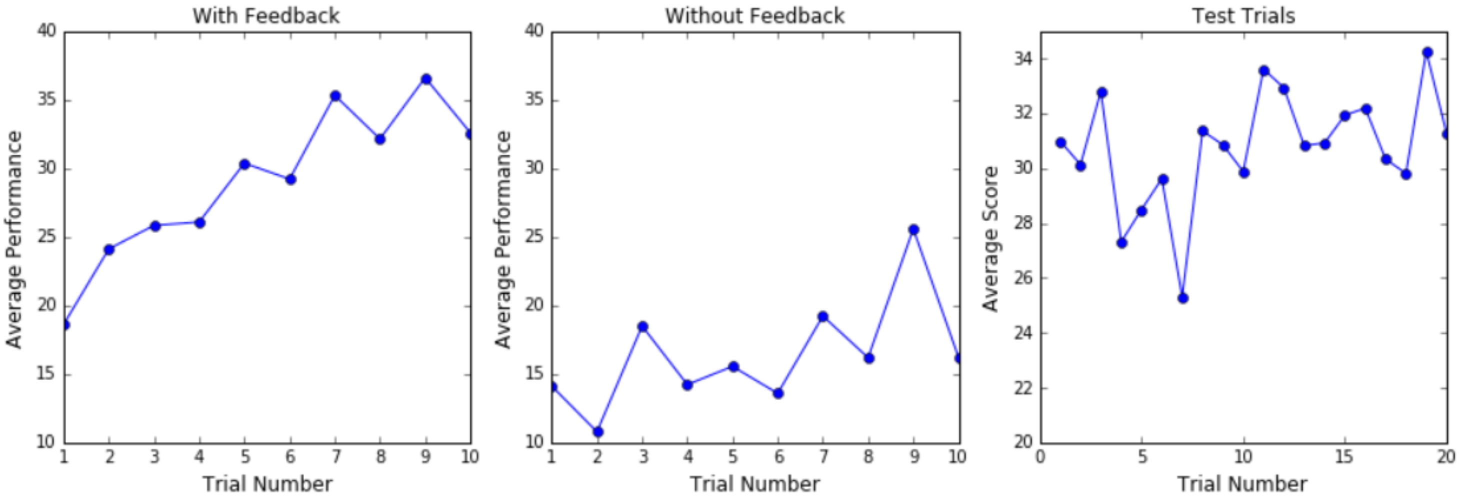

Figure 2: The learning curves for the training trials in both conditions and test trials. The average score for each trial of the experiment is shown. Note that the slope is steeper for the learning curve in feedback condition. Although no more feedback was presented during test trials, the average learning curve indicates that additional learning took place.

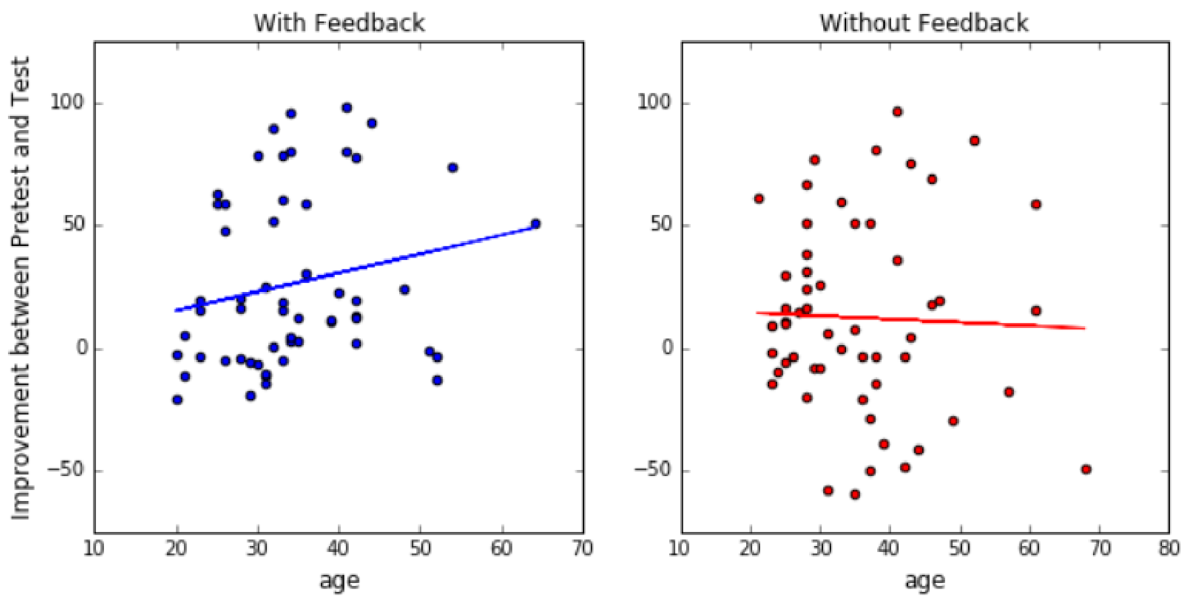

Figure 3: The difference between the average score of the last 5 test trials and the pretest trial plotted against age for both conditions. The solid line depicts the fit of a linear regression model.

To further assess the effect of feedback on learning, we calculated the slope of the learning curve (refer to Figure 2) over all training trials using linear regression with performance as the dependent variable (Figure 4). For the feedback condition, there was a significant positive correlation between age and slope ($\rho(54) = 0.285$; $p = 0.033$) suggesting that older adults benefit more from feedback. The same analysis for the control condition also revealed a positive correlation with age ($\rho(53) = 0.269$; $p = 0.047$), although less significant.

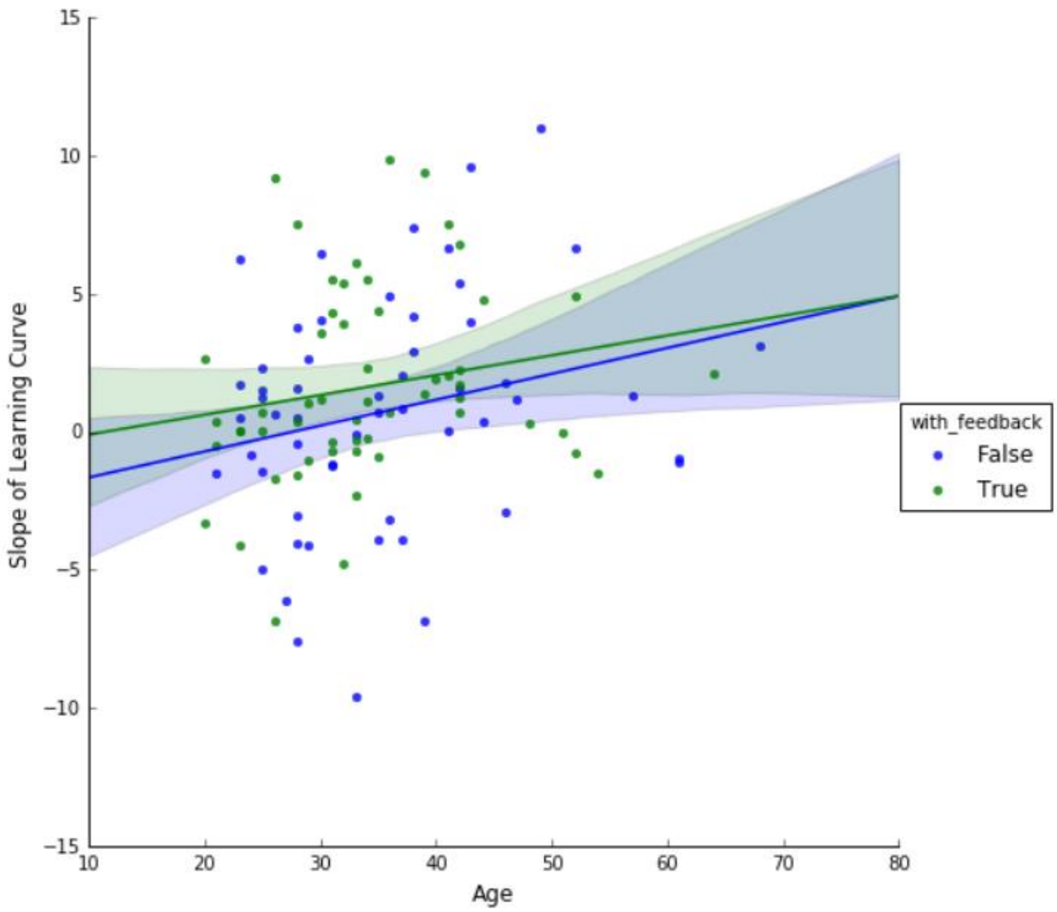

Figure 4: The slope of the participants' learning curves plotted against age. The solid lines show the fits of a simple linear regression model.

Discussion

In this experiment, we found a significant positive correlation between age and improvement on the task for those who received feedback from the tutors compared to those who didn't, suggesting that older people benefit much more from such feedback than younger people. However, it is also possible that older people simply had much more room for improvement given that they tended to start with lower pretest scores than the younger adults. While an ANOVA did not reveal any significant interactions between pretest score and age despite there being a significant negative correlation, our experiment only gave participants one pretest trial which is likely to add lots of noise to the data. To get a more robust and reliable sense of pretest performance, more pretest trials are necessary in future versions of this experiment, but the balance of having enough pretest trials without introducing learning effects is delicate and should be carefully considered.

Another limitation of Experiment 1 is that we didn't have an equal number of participants for each age range. We particularly had a small number ($n = 10$) of participants over 50 years of age, so it is difficult to conclusively say that older adults will benefit greatly from the feedback of the cognitive tutors, but given previous results (Lieder, 2018) and our own, we can say that the method shows promise and simply requires further testing with better sample sizes for different age groups.

Experiment 2

In Experiment 2, we sought to characterize which planning strategies were utilized by people of different age groups in order to answer the question of whether age affects strategy usage. We were also careful to recruit a bigger sample size of underrepresented ages in order to answer the question from Experiment 1 about whether older adults start the experiment with worse performance than younger adults.

Methods

Procedure

Upon starting the experiment, participants were presented with a short demographics survey that asked for their gender, their age, the country they lived in, their employment status, and the number of people in their household, including themselves. We used the answer given for age to determine whether the participant proceeded with the experiment; the other questions were distractor questions to avoid situations where participants respond dishonestly. If a participant was younger than 25 or older than 47, they proceeded to the second part of the experiment, where they were presented with 30 test trials of the Mouselab-MDP paradigm without any feedback (see the Methods section of Experiment 1 for details). If a participant was between 25 and 47 years of age, they saw a screen that thanked them for their participation in the survey and the experiment ended. The lower bound of 48 years old for our older adults group was chosen due to the dearth of data from this subset of the population in Experiment 1; their data made up 10.8% of the overall data. We subsequently found that setting the upper bound of 25 years old allowed us to have a comparable sample size of younger adults; they also made up 10.8% of the participants in Experiment 1. Those participants that completed the 30 trials of the Mouselab-MDP paradigm were given a comprehension check afterward as well as a survey about their planning strategies.

Participants

We recruited 348 participants from Amazon Mechanical Turk, only 90 of which passed the age screening and completed the full experiment. A further 12 participants were excluded from analysis due to failing the comprehension check (defined as getting at least 2 out of the 4 questions incorrect). This left us with 49 participants in the younger group and 29 participants in the older group. The base pay of the experiment was $0.05; those that passed the age screening received an additional $0.45 for the additional time spent on the experiment plus a performance dependent bonus of 1.5¢ for every $10 made in the experiment.

Analysis

In order to characterize and determine which planning strategies participants were using to solve the task, we built models of planning strategies to compare with participants' click sequences. The six strategies selected for comparison were the optimal planning strategy, depth-first search, breadth-first search, best-first search, progressive deepening, and random. The optimal strategy was derived by solving the MDP using the method described in Experiment 1. It first inspects the end nodes of each path, terminating once it finds a value of 48, the highest possible reward, and selects that path. If it inspects all end nodes and there are multiple nodes with equivalent value, the optimal strategy then inspects the second-to-last node in each path, effectively planning backwards, before terminating and selecting the path that awards the maximum reward. Depth-first search is the well-known strategy of first exploring a path, ignoring all branches, until its end, then returning up to the last fork in the path and subsequently exploring each branch (and recursively exploring any additional branch) to the end until all paths are explored. Breadth-first search first explores the first node of each possible path, then all of the second node, and so on, until it has explored all possible paths. Best-first search explores paths in the order of how promising they are — the path that is expected to return the highest rewards is explored before all others. Progressive-deepening search is a strategy proposed by Newell and Simon (1972) and is similar to depth-first search in the way that once chosen, a path is explored until the end. The difference is that once the strategy has finished with one path, rather than exploring the other branches of that path and completing the subtree, it may begin to explore another path from the beginning (i.e. another path that branches off from the start node). The random strategy is one that explores the nodes across all the paths at random.

We model participants' planning strategies as a combination of following one of the aforementioned strategies and moves that are not captured by any of these strategies and assumed to be made at random. Formally, this is defined as

$$\Pr(c|b, M, \theta_M) = (1 - \varepsilon) \cdot \sigma\big(c;\ V_{b,M}, \tau\big) + \ \varepsilon \cdot \mathcal{U}(c; C_b) \qquad \text{(Equation 1)}$$

The first term represents the probability that a particular node will be inspected under a particular strategy by taking a softmax over the possible clicks c in belief state b when following strategy M. The full equation for the softmax function is

$$\sigma\left(c;\ V_{b,M}, \tau\right) = \frac{exp\left(\frac{1}{\tau} \cdot V_{b,M}(c)\right)}{\sum_{c' \in C_b} exp\left(\frac{1}{\tau} \cdot V_{b,M}(c')\right)} \qquad \text{(Equation 2)}$$

where τ is temperature, a free parameter that determines sensitivity to actions with low probability. The second term in the model, $\mathcal{U}(c; C_b)$, represents the actions not explained by any particular strategy as a uniform distribution over all possible clicks and the terminal action to stop planning. ε is another parameter that determines how much randomness there was in a participant's clicks.

Apart from the random strategy, the definition of $V_{b,M}$ differs according to strategy M, but the rest of the model stays the same. The random model only has the second term, $\mathcal{U}(c; C_b)$. For the optimal strategy model, V_{b,M_O} is simply Q^*_{meta}, the solution to the MDP. For the rest of the strategies, $V_{b,M}$ reproduces their behavior in the way that values are assigned to each action. For example, in our depth-first search model, the depth of a node on a partially explored path, c, is the value assigned to $V_{b,M_{DFS}}$ so that deeper nodes are prioritized and paths that have been partially explored will be explored to completion before considering others. Similarly, our breadth-first search model prioritized shallower nodes on partially observed paths by setting the value of $V_{b,M_{BFS}}$ to -1 * the depth of c. For best-first search, $V_{b,M_{BestFS}}$ was defined as the expected sum of rewards along the same path as c. so that the most promising paths were considered first. $V_{b,M_{PD}}$ for progressive deepening was defined similarly to depth-first search but once a path is fully explored, the start of any sub-paths that branch off from it receive the same value as the start of all other completely unexplored paths. For all strategies, paths are explored in the order that they would be traversed. The depth-first search, breadth-first search, best-first search, and progressive deepening models had additional satisficing (Simon, 1956) and pruning (Huys et al., 2012) threshold parameters. When the expected reward for terminating in belief state b exceeded the satisficing parameter, then $V_{b,M}$would assign the value of 10^{10} to the terminating action (⊥) so that all strategies would strongly favor stopping. If the expected sum of rewards for any path fell below the pruning threshold, then $V_{b,M}$ would assign the value of 10^{20} to all of the remaining unobserved nodes on that path so that the strategy is discouraged from continuing to explore that unprofitable path.

We fit all models to every trial for each participant and calculated the log likelihood for values of τ, ε, and the satisficing and pruning thresholds that maximized it. We then used the Bayesian information criterion (Schwarz, 1978) to select the most likely used strategy for each trial.

Results

Since we were careful to recruit a larger sample size (n = 29) for the older group in this experiment, we revisited the question of whether older adults perform worse at the beginning of the experiment compared to younger adults in the absence of any learning aids such as feedback. We looked at only the first 5 trials for each participant in order to produce a more reliable estimate of starting performance while hoping to keep any learning effects from practice to a minimum. Taking the average score over the first 5 trials revealed that the average baseline performance for younger adults is 19.05 points while the baseline for older adults was just 6.78 points. Thus, we can conclude that older adults start the task with worse performance compared to younger adults.

Examining the usage frequency of strategies in both age groups revealed that while the optimal strategy was the most preferred by younger adults, having been used for nearly 48% of the total trials completed by their age group, older adults employed the optimal strategy and depth-first search strategy equally, both being used for 40% of the overall trials for their age group.

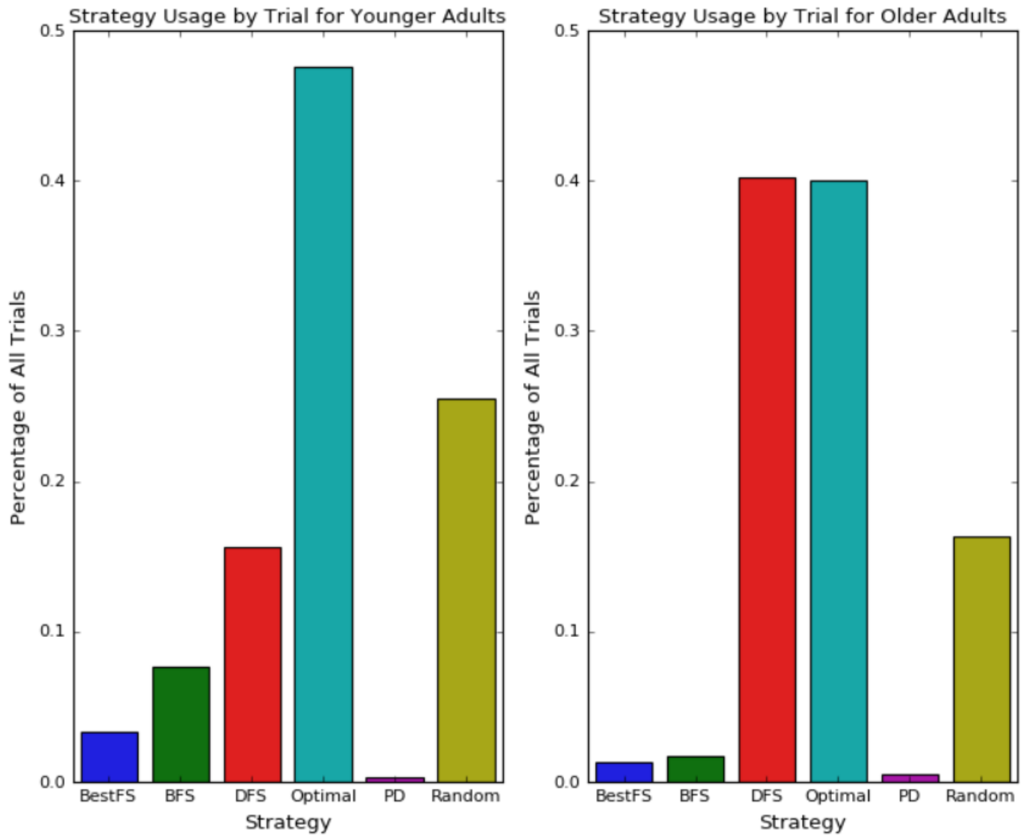

Figure 5: Strategy usage frequencies for younger adults and older adults over all trials. The strategies that were used are (from left to right): Best-first search, breadth-first search, depth-first search, optimal, progressive deepening, and random.

A chi-squared test revealed that the strategy usage frequencies shown in Figure 5 are significantly different between older versus younger adults ($X^2(5) = 205.43$, $p < .001$). Younger adults used best-first search, breadth-first search, the optimal strategy, and the random strategy more often than older adults, while older adults used depth-first search more often than younger adults. We ran additional chi-squared tests of independence for the frequency of a particular strategy in both age groups. Interactions for all strategies save for progressive deepening were significant and are summarized in Table 1. It must be noted that the number of occurrences for progressive deepening did not exceed 5 for the elderly group and thus the chi-squared test may be invalid for this case.

Strategy	**$X^2(1)$**	**p**
Best First	8.55	.003
Breadth First	35.86	< .001
Depth First	175.85	< .001
Optimal	12.30	< .001
Progressive Deepening	0.01	> 0.05
Random	26.28	< .001

Table 1: Results of the chi-squared test of independence for strategy usage frequencies in younger and older adults. Interactions for all strategies except for progressive deepening were significant.

We also looked at the frequency of strategy usage over time in both younger and older adults. For the younger adults, the random strategy was most favored for the first few trials until the optimal strategy gained popularity around trial 7 and its usage rose over time while the usage of other strategies eventually dropped. For the older adults, the depth-first search strategy is the most popular at the beginning of the experiment and isn't overtaken by the optimal strategy until around trial 11 (see Figure 6). After these two strategies, the next popular strategy for older adults was the random strategy. For both groups, the most frequently used strategy at the end of the experiment was the optimal strategy with 67% of the younger adults using it on the last trial versus 62% of the older adults.

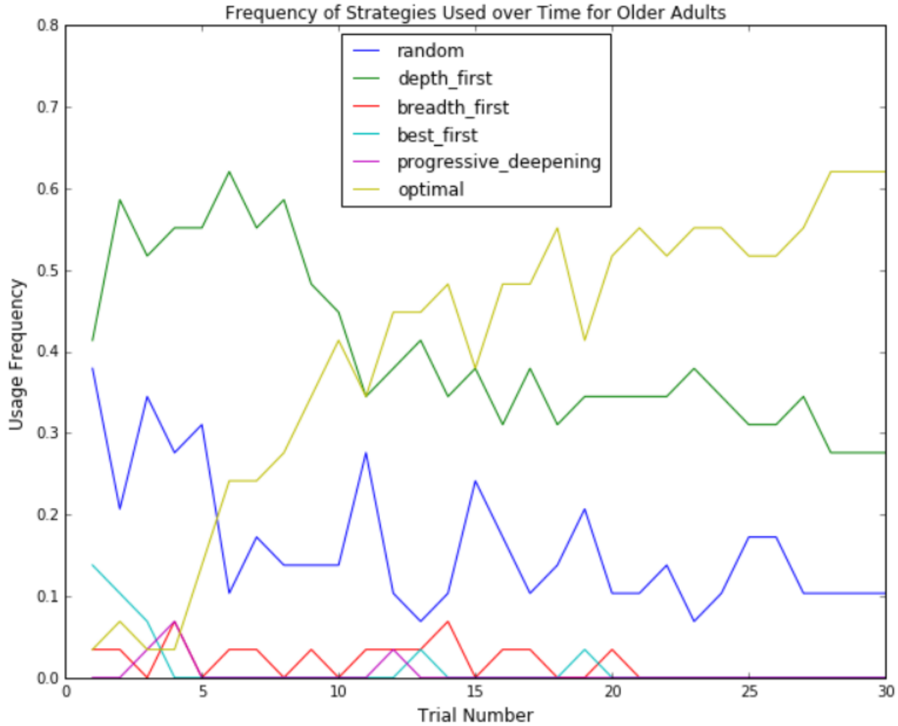

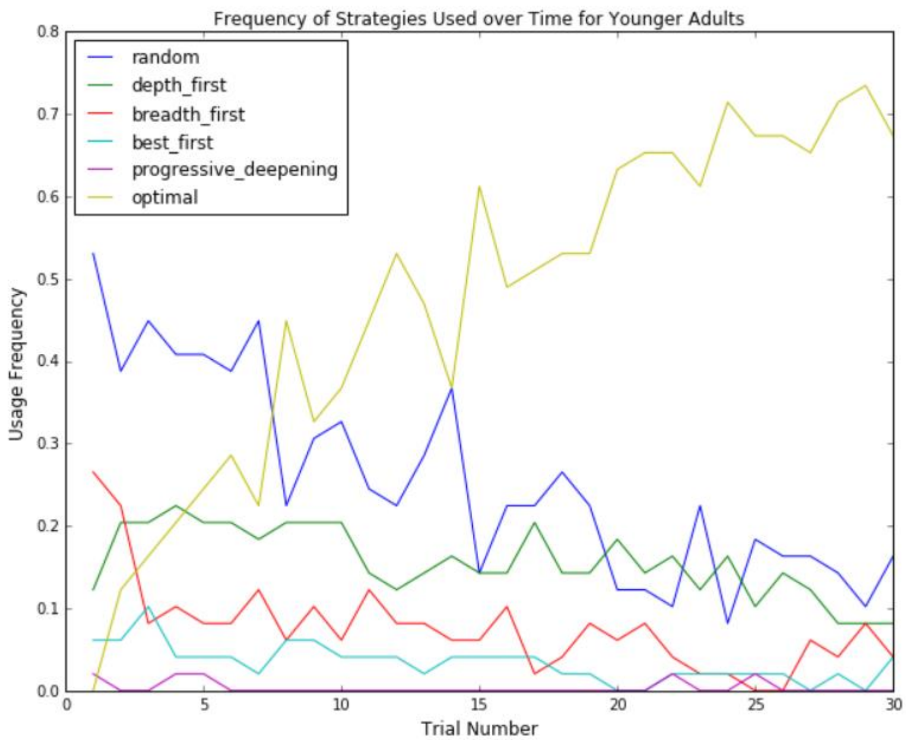

Figure 6: The frequencies of strategy usage for every trial in the experiment for younger adults (top) and older adults (bottom).

Discussion

In this experiment, we found evidence that older adults use certain strategies at differing rates compared to younger adults. While both groups heavily favored the optimal strategy, collectively using it at least 40% of the time, when we looked at the frequency of strategy usage over time, the data revealed that the use of the optimal strategy only became more favorable as the experiment went on, suggesting that both groups discovered and learned the optimal strategy for the task on their own (since cognitive tutors were not used in this experiment). However, the usage of the optimal strategy in the older group was still not as high as usage in the younger group, which indicates that not as many participants in the older group discovered or favored the optimal strategy. On average, participants in the older group scored 20.19 points per trial compared to the 30.65 points per trial scored by the younger group, a difference greater than 10 points. Since following the optimal strategy will maximize the overall score (Lieder, 2018), this difference in points must be due in part to the differing frequencies of optimal strategy usage in both groups. If this difference in usage is due to older participants being unaware of the existence of the optimal strategy, then this is a case where it would be useful to introduce the cognitive tutors and examine the effect on strategy usage frequencies. If this is a case of older participants favoring another strategy, this opens up further lines of research as to why other strategies are perceived as more effective and how interventions can be tailored to convince these participants to use more effective alternatives. It is also particularly interesting to look at strategy usage frequencies at the start of the experiment for both groups, when the task was relatively novel and there has been little opportunity to test several different strategies. For the younger group, the random strategy is dominant for the first few trials whereas the older group preferred the depth-first search strategy more than any other. This unexpected result contradicts the earlier finding reported by Gilhooly et al. (1999) that strategies used for the Tower of London planning task did not differ by age. This earlier study also considered the depth-first search, breadth-first search, and progressive deepening strategies and while they found that younger adults made "deeper" plan searches than older adults, this difference is largely explained by the limited working memory capacity of older adults, since all plans were formed mentally and described verbally. Since our Mouselab-MDP task does not require the use of working memory, our findings allow us to update the belief that different age groups do not use different planning strategies, at least, for situations that do not require working memory. Of course, the difference in usage of strategies between age groups is most likely dependent on the task as well, since in Gilhooly et al.'s (1999) study, they found that 10.14% of younger adults and 6.37% of older adults used progressive deepening, which does not align with our findings that both groups used progressive deepening less than 1% of the time. However, our results do agree that neither group has a tendency to use progressive deepening.

General Discussion and Conclusion

In order to determine whether planning performance could be maintained or improved as a person ages, we used cognitive tutors, developed from the ideas of resource-rational analysis, to teach people how to use more effective planning strategies. We also examined the types of strategies younger and older adults use in route planning tasks without the assistance of such cognitive tutors. We found that older adults performed worse than younger adults on average, but that older adults benefited more from the cognitive tutors and improved their planning at a faster rate. We also found that age affects the choice of strategies used for planning, contrary to the findings of an earlier study (Gilhooly et al., 1999), and that older adults tend to use the optimal strategy less frequently than younger adults. Taken together, these findings suggest that cognitive tutors can help reduce the gap in performance between older and younger adults by teaching older adults the optimal strategy, which they might not otherwise discover on their own. Learning with cognitive tutors would be beneficial outside the lab as well, since recent work (Lieder, 2018; Lieder et al., 2018) has shown that the skills gained from training carry over to different tasks.

Of course, as with any cross-sectional study of aging, this study suffers from limitations such as selection bias and cohort bias. Participants were not screened for any cognitive impairments beforehand either. These are all important considerations for future replications of this study, though by using a service such as Amazon Mechanical Turk, we can be fairly confident in claiming that any differences found between age groups was not due to unfamiliarity with the technology used to run this study.

Now that we have a better idea of how age affects learning to plan and the planning process, future applications involve developing tailored cognitive tutoring to best meet the needs of the elderly population. For example, a person who favors a random strategy may just need a gentle push towards a strategy with structure, whereas a person who approaches problems from the start state with depth-first search may need to be shown that sometimes planning with the end goal in mind yields better results. If cognitive tutoring can reliably produce lasting improvements in cognition, as preliminary results indicate (Lieder et al., 2018), then the quality of life for elderly people will drastically increase along with the capacity for independent living. Our findings bring us one step closer to this future.

Acknowledgements

The author would like to thank her supervisors, Dr. Falk Lieder, for being extremely patient and understanding while she learned how to conduct a research project from the ground up, and Professor Thomas Griffiths, for this document would not exist without his encouragement and his belief in the author's abilities. Sincere thanks are also in order to Professor William Jagust for his time spent reading and evaluating this work. The author also expresses deep gratitude to Fred Callaway, for his assistance with the modeling sections of this work, and to all the friends who never doubted that she would finish, no matter how long it took.

References

Allain, P., Nicoleau, S., Pinon, K., Etcharry-Bouyx, F., Barré, J., Berrut, G., … Le Gall, D. (2004). Executive functioning in normal aging: A study of action planning using the Zoo Map Test. https://doi.org/10.1016/j.bandc.2004.08.011

Callaway, F., Lieder, F., Das, P., Gul, S., Krueger, P. M., & Griffiths, T. L. (2018). A resource-rational analysis of human planning. In *Proceedings of the 40th Annual Conference of the Cognitive Science Society*. Madison, WI. https://doi.org/10.13140/RG.2.2.15636.40326

Gilhooly, K. J., Phillips, L. H., Wynn, V. E., Logie, R. H., & Sala, S. Della. (1999). Planning processes and age in the five-disc Tower of London task. *Thinking and Reasoning*, *5*, 339–361. https://doi.org/10.1080/135467899393977

Griffiths, T. L., Lieder, F., & Goodman, N. D. (2015). Rational use of cognitive resources: Levels of analysis between the computational and the algorithmic. *Topics in Cognitive Science*, *7*(2), 217–229.

Harada, C. N., Natelson Love, M. C., & Triebel, K. L. (2013). Normal cognitive aging. *Clinics in Geriatric Medicine*, *29*(4), 737–52. https://doi.org/10.1016/j.cger.2013.07.002

Huys, Q. J. M., Eshel, N., O'Nions, E., Sheridan, L., Dayan, P., & Roiser, J. P. (2012). Bonsai Trees in Your Head: How the Pavlovian System Sculpts Goal-Directed Choices by Pruning Decision Trees. *PLoS Computational Biology*, *8*(3), e1002410. https://doi.org/10.1371/journal.pcbi.1002410

Lieder, F. (2018). *Beyond bounded rationality: reverse-engineering and enhancing human intelligence*. University of California, Berkeley.

Lieder, F., Callaway, F., Krueger, P. M., Das, P., Griffiths, T. L., & Gul, S. (2018). Discovering and Teaching Optimal Planning Strategies. In *Proceedings of the 14th Biannual Conference of the German Society for Cognitive Science, GK*. Darmstadt.

Murman, D. L. (2015). The Impact of Age on Cognition. *Seminars in Hearing*, *36*(3), 111–21. https://doi.org/10.1055/s-0035-1555115

Newell, A., & Simon, H. A. (1972). *Human problem solving*. Englewood Cliffs, NJ: Prentice-Hall.

Phillips, L. H., MacLeod, M. S., & Kliegel, M. (2004). Adult Ageing and Cognitive Planning. In R. Morris & G. Ward (Eds.), *The Cognitive Psychology of Planning* (pp. 111–135). Taylor & Francis.

Schwarz, G. (1978). Estimating the Dimension of a Model. *The Annals of Statistics*, *6*(2), 461–464. https://doi.org/10.1214/aos/1176344136

Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review*, *63*(2), 129–138. https://doi.org/10.1037/h0042769

Sorel, O., & Pennequin, V. (2007). Aging of the Planning process: The role of executive functioning. https://doi.org/10.1016/j.bandc.2007.07.006