Title
Adaptivity and optimization under constraints

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 45(45)

Authors
Koenen, Reba
Varma, Sashank

Publication Date
2023

Peer reviewed

Adaptivity and optimization under constraints

Reba Koenen (rkoenen3@gatech.edu)
School of Psychology, 654 Cherry Street NW
Atlanta, GA 30313 USA

Sashank Varma (varma@gatech.edu)
School of Interactive Computing, School of Psychology, 85 5th St NW
Atlanta, GA 30308 USA

Abstract

Human agents confront internal computational limitations and external constraints of the task environment when problem solving. To find optimal solutions under time and material constraints, the agent must adapt their behavior by using various strategies such as offloading (i.e., using external materials to aid their performance). We designed a novel tower building task to investigate adaptive use of strategies under constraints. The task worked as designed: Participants found the optimal solution most often on the least difficult scenario and least often on the most difficult scenario. Surprisingly, offloading led to no significant differences in performance. On the most difficult scenario, some participants found the optimal solution using a *prospective, concurrent,* or *retrospective* strategy based on experience with the constraints of the task environment. This study shows how optimality can be understood as a trend over time and investigated in tasks that allow multiple attempts.

Keywords: adaptivity, optimization, strategies, problem solving, constraints

Introduction

Optimal problem solving is studied across various disciplines with foundations in economics (Pingle, 1992; Pingle & Day, 1996; Simon, 1955; Smith et al., 1982) and with more recent work in cognitive science (Hawkins et al., 2012; Howes et al., 2014; Howes et al., 2016; Jain et al., 2022; Lewis et al., 2014; Lieder & Griffiths, 2020). An important subtopic is problem solving under constraints. The constraints of a task environment (e.g., time, materials) impact the strategies by which a person solves a problem, and in turn affects their ability to find an optimal solution. The current study investigates optimal problem solving under time and material constraints using a novel tower building task. It also evaluates the role of offloading in potentially supporting optimal problem solving. Lastly, it documents the different strategy progressions people use over multiple attempts on the same difficult problem.

Prior Research

Foundational research on optimality begins with classic economic theories of maximization (Smith et al., 1982). The goal of maximization is to find the highest-value solution given the defined costs and benefits of a task environment. Maximization assumes unlimited resources in terms of both time and computational capabilities to arrive at a solution. Simon (1955) pointed out that a limitation of maximization is that the human problem solver is bounded in their computational capabilities. He proposed *bounded rationality*, claiming the computational capabilities of the human agent are constrained both internally by their computational limitations (i.e., their ability to conduct mental calculations) and externally by their experience with a task environment. Thus, it is often not possible to exhaustively search for the best solution. Instead, humans must *satisfice,* determining a preset threshold (i.e., aspiration level) and stopping once they arrive at such an outcome (Simon, 1955).

Satisficing, too, is limited in its assumption that people terminate problem solving after arriving at a predefined threshold (Payne et al., 1988). If the constraints of the task environment allow it, people may seek to adjust their original threshold, and aim instead for an *optimal* outcome as opposed to settling for a satisfactory one (Bhui et al., 2021). The current study investigates whether and how people adapt their behavior under varying time and material constraints to find the optimal solution.

Conlisk (1988, 1996) considered the question of optimization under constraints from an economics perspective. He argued that searching for relevant information and finding the optimal solution are both costly activities, describing these activities as *optimization costs*. He claimed that traditional economic theories are insufficient because they do not account for these costs, which are central to the decision outcome (Conlisk, 1988). From an AI perspective, Boddy and Dean (1994) considered the design of intelligent agents that perform in time-constrained environments. They argued that management of an agent's limited computational resources is crucial for optimal decision making. They proposed algorithms for the proper allocation of these resources during planning and problem solving.

A very different approach to problem solving under constraints comes from distributed cognition, which proposes sharing the task's cognitive load between the agent and the environment (Hutchins, 1995). Similarly, Tversky (2011) reviews the benefit of visualizing meaning in the external world. Utilizing the world as a source of searchable information can aid people in their ability to solve problems. For example, Zahner and Corter (2010) found evidence for the importance of external visual representations in problem solving. Across a variety of tasks, lower performers tended to

In M. Goldwater, F. K. Anggoro, B. K. Hayes, & D. C. Ong (Eds.), *Proceedings of the 45th Annual Conference of the Cognitive Science Society*.

rely on the resources of the external environment, compared to higher performers.

Use of the external world as a problem solving aid is sometimes referred to as cognitive *offloading.* Offloading is when people utilize external resources (e.g., paper, pen) to distribute their cognitive load during problem solving (Risko & Gilbert, 2016). It is an effective strategy in constrained problem solving environments, used most often in challenging or novel task environments where the agent has little experience with the specific constraints. The current study contributes to this literature, investigating the use of offloading on optimal performance during time-constrained problem solving.

Recently, Lieder and Griffiths (2020) expanded on prior work on distributing resources during problem solving in their model of *resource-rational analysis*. This model proposes that allocating resources optimally enables the agent to find a solution. They propose that allocating resources effectively is based on planning. The specific type, and timing, of planning varies based on the agent and the constraints of the task environment (Lieder & Griffiths, 2020).

Lieder and Griffith's (2020) proposal is consistent with Simon's (1955) original formulation of bounded rationality which called for *adaptive* mechanisms for handling the constraints of a task environment. Adaptation occurs when a person changes their behavior in response to obstacles during problem solving. For example, they may adjust their strategy based on their understanding of the task instructions or their experience with the task environment and its constraints. Crucially, time constraints may result in a person modifying their plan to find an optimal solution within the allotted time.

The current study examines three forms of strategy adaptation. First, people may engage in *prospective planning*, expending time upfront to ultimately get closer to an optimal solution. In our study, participants considered the potential value associated with each outcome using pen and paper to try to find the optimal solution. Second, people may use *concurrent planning*, adapting their strategy as they gain experience with the affordances of a novel task environment. In our study, participants laid out blocks in front of them to help them visualize an optimal solution. Third, people may engage in *retrospective planning*, adjusting their plan only after spending their initial time finding an optimal solution. In our study, participants tended to calculate the value of the solution they previously found before proceeding. Each of these adaptive strategy progressions is used to aid the agent in their plan to find the optimal solution.

Research Questions

The current study examined optimal problem solving under time and material constraints. College undergraduates completed a novel tower-building task under three different scenarios. Each scenario had different time and material constraints, resulting in different optimal solutions. This study addressed five research questions:

1. Does the optimality of problem solving decrease as the number of constraints to manage increases?
2. Do people learn to perform more optimally across successive attempts on a scenario?
3. Do people who choose to offload during problem solving (i.e., use paper and pen resources) find more optimal solutions than those who do not?
4. Are people's metacognitive ratings of scenario difficulty associated with the optimality of their performance?
5. Across successive attempts on the most difficult scenario, C, do people exhibit different strategy progressions as they move toward more and more optimal solutions?

Method

Participants

The sample consisted of 45 undergraduates from a public technical university in the American Southeast. Their mean age was 19.9 years (*SD* = 1.7). Participants were compensated with course credit for one hour of their time. The protocol was approved by the local IRB.

Design

The study had a within-subjects design. The first factor was scenario (A, B, C). The second factor was attempt (1, 2, 3): participants completed each scenario three times in succession. Thus, there were a total of nine attempts. The time (i.e., number of seconds provided for each attempt) and materials (i.e., number of blocks provided for each attempt) differed across the scenarios. (See below for details on these constraints.)

The primary dependent variable was the number of optimal solutions. This was aggregated for each scenario (i.e., the total number of the three attempts for which a participant found the optimal solution) and aggregated across all scenarios and attempts (i.e., the total number of the nine attempts for which a participant found the optimal solution). Scenario C had an additional dependent variable: overall benefit (out of 540, which was the highest score participants could achieve on this scenario).

We collected three additional measures via a strategy questionnaire. This instrument asked, for each scenario, what strategy a participant used and if their strategy changed over the duration of the scenario. It also asked participants to rate their perceived difficulty of the scenario on the scale of 1 ("not difficult") to 10 ("very difficult").

Materials

We developed a novel task to investigate optimality of problem solving under time and material constraints. It consists of three scenarios. For each, participants were asked to build towers. Each tower was composed of three colored blocks; see Figure 1. The *benefit* of each tower was consistent across the entire task; left = 50, middle = 60, and right = 100.

Scenarios A and B also had *costs* associated with using each block; pink = 5, green = 10, and black = 20. Participants could build any number of each of these three block towers. The goal was to optimize the *net benefit* for each attempt of each scenario. The allotted time and number of blocks available for each scenario varied. Therefore, the optimal solution for each scenario varied, as did the difficulty of finding it.

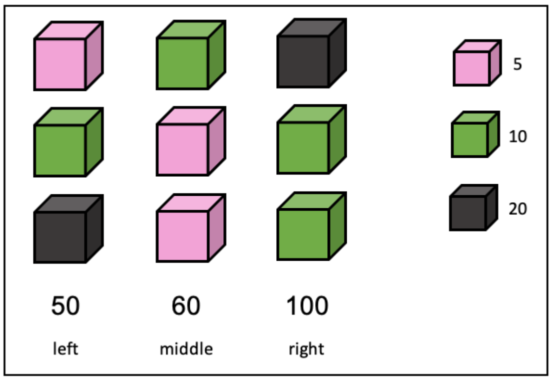

Figure 1. The three tower options and their associated point benefit, and also, the block costs.

Scenario A. Scenario A was always completed first as it was intended to be the easiest one. Participants were given 30 seconds for each attempt. The goal was to build the combination of the towers shown in Figure 1 resulting in the highest net benefit. The costs of blocks were denominated in terms of points. For example, the left tower of Figure 1 has a benefit of 50 points and a cost of 5 + 10 + 20 = 35 points. Thus, the left tower has a net benefit of 50 – 35 = 15 points. The middle tower has a net benefit of 40 points and the right tower a net benefit of 60 points. Thus, the optimal strategy is to build as many right towers as possible in the allotted time. Participants were *unconstrained* in this scenario.

Scenario B. Scenario B was intended to be of moderate difficulty, and so participants were allotted 90 seconds for each attempt. Again, the goal was to build the combination of towers resulting in the highest net benefit. Here, participants were subjected to a *temporal constraint*: the costs of the blocks were denominated in terms of time: participants had to request blocks one at a time and had to wait the specified amount of time to receive each one (e.g., 10 seconds for a green block). Recall that the left tower of Figure 1 has a benefit of 50 points. It has a cost of 35 seconds, and thus choosing to build this tower yields a benefit of 50 / 35 = 1.43 points/sec. The middle tower yields 3 points/sec and the right tower 2.5 points/sec. Thus, the optimal strategy is to build as many middle towers as possible during the allotted time.

Scenario C. Scenario C was intended to be the most difficult. Participants were given 180 seconds for each attempt. This scenario had a *material constraint*: Participants were provided with exactly 10 blocks of each color with no costs associated with using the blocks. There were two tower combinations that achieved the highest possible net benefit (540) using these 30 blocks: (1) four middle towers and three right towers or (2) two left towers, four middle towers, and two right towers.

Procedure

Participants completed the task in-person in the lab, seated at a table. In front of them were 120 blocks (40 per color), a sheet displaying the three tower options and their associated benefits (see Figure 1), a sheet outlining the block costs (see Figure 1), and the offloading materials (i.e., blank paper, a pen, and markers of the same three colors as the blocks).

The experimenter explained the benefits of the towers, the cost of using each block, and the overall goal for each attempt: to build the combination of towers resulting in the highest net benefit. Participants were told they could not take a tower apart once it had been built, and that they must let the timer expire for each attempt, even if they finished early. The experimenter also pointed out the offloading materials, which could be used at any point during the task.

Each participant was then given the rules for scenario A. The experimenter set the timer, visible to the participant, and the participant began their first attempt. Once time elapsed, the experimenter recorded the number of each tower (i.e., left, middle, right) the participant chose to build. The towers were then disassembled, the blocks returned to the pile, and the timer was set for the second attempt. The participant completed two more attempts for scenario A.

After completing scenario A, half of the participants (N = 23) completed scenario B and the other half (N = 22) completed scenario C. The procedure was the same as for scenario A, other than the varying time and material constraints. Participants completed three attempts for their second scenario and then three attempts for their final scenario. They were not provided with feedback at any point during their nine attempts. Afterwards, they completed a questionnaire asking about their strategies and their perception of the difficulty of each scenario. Finally, participants completed a short demographics form and were debriefed. The study lasted approximately 45 minutes.

Results

Data Scoring and Strategy Coding

We first coded whether or not (1 or 0) each participant found the optimal solution on each of their nine attempts. We scored each attempt on scenario C in greater detail by also computing the total benefit of the constructed towers (i.e., with reference to Figure 1). Possible scores for scenario C ranged from 0 to 540. An independent coder repeated this process to ensure accuracy; agreement was 100%.

We also coded the responses to the strategy questionnaire for data on scenario difficulty and participants' strategies. For scenario C, the first author performed an initial coding of each participant's strategies. A review of these 45 responses led to the identification of three strategy progressions toward the optimal solution: (1) offload with *prospective planning* by checking their math before their first attempt, (2) offload with *retrospective planning* by checking their math after the first attempt, and (3) offload with *concurrent planning* by laying out the blocks into patterns before building. One of these three was assigned to each response based on the primary strategy the participant reported using. If none of these strategies were used, a 0 was assigned. Two independent raters coded the same 45 responses using a coding key developed by the first author. Overall agreement between all three coders was 82.4% and disagreements were resolved via discussion.

Optimal Performance Under Time and Material Constraints

Optimality of Performance by Scenario and Attempt The first and second research questions concerned optimality of human problem solving. The first asked whether optimality decreases as the number of constraints increases. The second asked if people learn to perform more optimally across successive attempts on a scenario. Again, performance was defined as whether or not a participant found the optimal solution for a given attempt. Average performance on the task was 5.3 out of 9 attempts solved *(SD* = 2.0).

To address the first research question, we conducted a one-way repeated measures ANOVA with factor scenario (A, B, C) and with dependent variable the number of optimal solutions across the three attempts. There was a main effect of scenario ($F(2, 88) = 6.70$, $p = .002$, $\eta 2 = 0.13$). A series of *post hoc* paired *t*-tests revealed that participants produced more optimal solutions on scenario A ($M = 2.1$, $SD = 1.9$) than scenario C ($M = 1.3$, $SD = 1.0$) ($t(44) = 3.83$, $p = .001$, $d = 0.72$), and on scenario B ($M = 1.9$, $SD = 1.0$) than scenario C ($t(44) = 2.59$, $p = .039$, $d = 0.52$); see Figure 2. Performance on scenarios A and B was comparable ($p = .911$).

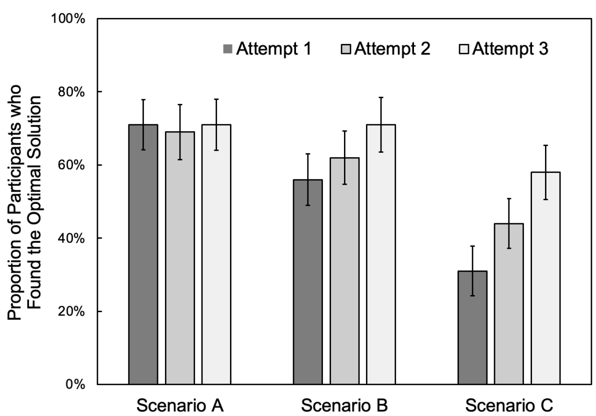

Figure 2. Proportion of participants finding the optimal solution by scenario and attempt. Error bars are SEs.

We addressed the second research question using a parallel one-way repeated measures ANOVA with factor attempt (1, 2, 3) and with dependent variable the number of optimal solutions across the three scenarios. The prediction was that optimality would increase across the three attempts. This was in fact the case. There was a main effect of attempt ($F(2, 88) = 4.31$, $p = .016$, $\eta 2 = 0.09$). A series of *post hoc* paired *t*-tests revealed that participants produced more optimal solutions on the third attempt ($M = 2.0$, $SD = 0.8$) compared to the first attempt ($M = 1.6$, $SD = 0.8$) ($t(44) = 2.80$, $p = .023$, $d = 0.49$). Performance on the second attempt ($M = 1.8$, $SD = 0.9$) was comparable to performance on both the first attempt ($p = .896$) and the third attempt ($p = .078$).

Closer inspection of Figure 2 revealed that performance on scenario A remained relatively constant across its three attempts. This aligns with its intended design to be the easiest of the three scenarios. Scenario B performance tended to increase over its three attempts. This was designed to be the scenario of moderate difficulty, given that it imposed only a temporal constraint. Nevertheless, scenarios A and B had the same number of people who found the optimal solution on the third attempt. Scenario C was designed to be the most difficult given its material constraint. Like scenario B, performance increased over its three attempts, but it had the lowest overall performance. This supports our prediction that increasing constraints results in fewer people finding the optimal solution.

Effects of Offloading on Optimal Performance The third research question asked whether people who chose to offload during problem solving (i.e., use paper, pen, and markers) performed more optimally than those who did not. We predicted that there would be a positive effect of offloading. Surprisingly, there was not. An independent samples *t*-test revealed that performance on scenario A was comparable for participants who chose to offload versus those who didn't ($p = .274$). The same was true for scenario B ($p = .449$) and scenario C ($p = .911$); see Figure 3.

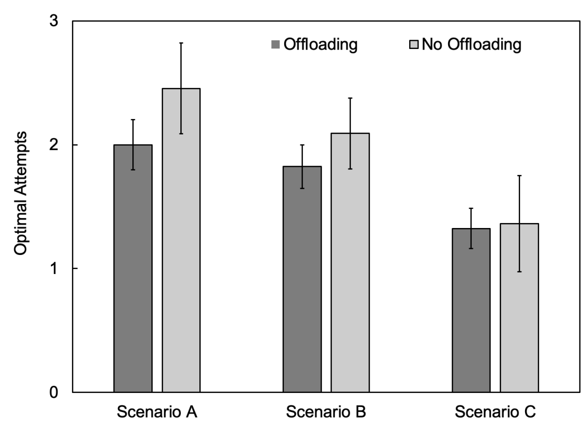

Figure 3. Average number of optimal attempts for each scenario, grouped by use of offloading. Error bars are SEs.

We were surprised by these results as they are inconsistent with some prior studies (Risko & Gilbert, 2016). Still, recall Zahner and Corter (2010) found lower performers tended to rely on the resources of the external environment more than higher performers. These mixed findings across studies prompted further examination of the specific types of offloading participants engaged in on scenario C; see section Scenario C: Optimal Strategy Progressions.

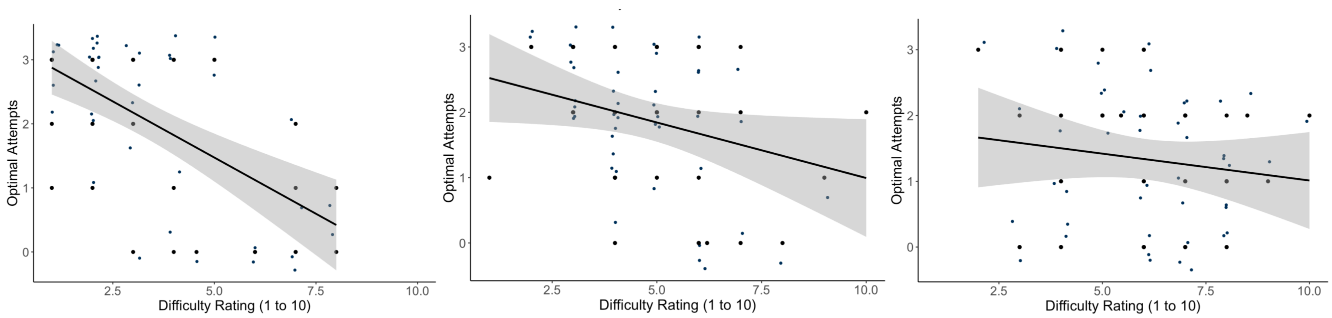

Figure 4. Correlation between difficulty rating and the number of optimal attempts for scenario A (left), B (middle), and C (right). Note that the jitter plot spatially separates the points around each coordinate for ease of interpretation.

Difficulty Ratings and Optimal Performance

The fourth research question asked whether people's metacognitive rating of the difficulty of a scenario is associated with the optimality of the solutions they generated. A one-way repeated measures ANOVA with factor scenario (A, B, C) on participants' difficulty ratings revealed a main effect of scenario ($F(2, 88) = 29.79$, $p < .001$, $\eta 2 = 0.41$). A series of post-hoc paired *t*-tests found that, as expected, participants rated the difficulty of scenario A ($M = 3.2$, $SD = 2.1$) as lower (i.e., less difficult) than both scenario B ($M = 4.7$, $SD = .8$) ($t(44) = 4.10$, $p = < .001$, $d = 0.62$) and scenario C ($M = 6.1$, $SD = 1.8$) ($t(44) = 7.27$, $p < .001$, $d = 1.10$). Scenario B was also rated as less difficult than scenario C ($t(44) = 3.86$, $p < .001$, $d = 0.58$). These results corroborate our design goal of making scenario C the most difficult one.

Next, we investigated the relationship between difficulty rating and performance (i.e., number of optimal solutions across the three attempts separately for each scenario). A correlational analysis revealed the expected significant negative relationship between difficulty ratings and number of optimal solutions on scenario A ($r = -.63$, $p < .001$) and on scenario B ($r = -.31$, $p < .001$), but not for scenario C ($r = .12$, $p = .492$); see Figure 4.

Scenario C: Strategy Progressions

The final set of analyses addressed the fifth research question: Across successive attempts on the most difficult scenario, C, do people exhibit different strategy progressions toward the optimal solutions?

In total, 33 of the 45 participants found the optimal solution of 540 on at least one of their three attempts; see Figure 5. (Recall that two different tower configurations achieved that score.) Of these 33 participants, 27 of them found the optimal solution in a monotonic manner. The first group ($N = 3$) found an optimal solution on attempts 1 and 3, the second group ($N = 11$) found an optimal solution on attempts 2 and 3, and the third group ($N = 7$) found an optimal solution on attempt 3. A fourth group ($N = 6$) found an optimal solution on all three attempts. The remaining of the 33 participants ($N = 6$) behaved in a non-monotonic manner: They found an optimal solution on either attempt 1 or 2 but not on attempt 3.

The first three groups are shown in black in Figure 5. Interestingly, they all used specific forms of offloading to find the optimal solution. These 21 participants were thus assigned to one of three strategy progressions.

Participants ($N = 3$) who used *Retrospective Planning* and participants ($N = 11$) who used *Prospective Planning* found the optimal solution on two of the three attempts. Participants in each of these groups adapted their strategies, using either prospective or retrospective planning to check their math. Participants did this by looking at the sheet containing the tower benefits while using the provided pen and paper to try to find the optimal solution. Those who used prospective planning tended to mathematically calculate the optimal solution before building their towers (i.e., during the first attempt). Those who used retrospective planning tended to calculate the value of the towers they had already built, to determine their net benefit. They tended to do this after their first attempt. For both of these groups, their allocation of scarce resources (i.e., the time they spend checking their math) likely took up time where they would have otherwise been building towers, resulting in the two optimal attempts.

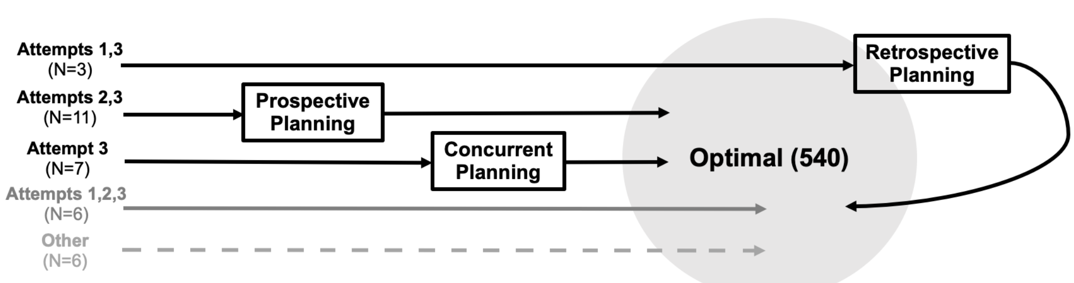

Figure 5. All progressions toward the scenario C optimal solution. The first three are strategy progressions.

The third group (N = 7) found the optimal solution only once. Unlike the first two groups, these participants adapted their strategy by offloading with materials besides pen, markers, or paper. Instead, they laid out the blocks on the table in front of them before building the towers. Consistent with Tversky's (2011) proposal, this may have helped participants visualize their tower combinations before building, since a tower could not be disassembled once it was built. We termed this *Concurrent Planning* since this strategy progression spanned all three attempts.

The final two groups in Figure 5 did not utilize the three strategy progressions during their three attempts. The fourth group of participants (N = 6) found the optimal solution on all three attempts, without the use of offloading. This is consistent with Zahner and Corter's (2010) finding that higher performers relied on external resources less than their lower performing counterparts.

Discussion

The goal of the current study was to investigate optimal problem solving under time and material constraints. The novel task developed for this study was inspired by theories of bounded rationality (Simon, 1955) and how optimization costs impact human problem solving (Conlisk, 1988).

For each scenario, participants faced varying constraints with the goal of finding the optimal solution. The results showed that they found the optimal solution most often on scenario A, which was rated as the least difficult scenario. They found the optimal solution least frequently on scenario C, which was rated as the most difficult scenario. Surprisingly, use of general offloading did not help participants find the optimal solution. The correlations between the number of optimal solutions found and perceived difficulty ratings were significant for scenarios A and B, but not scenario C. We further investigated scenario C using participants' responses on the strategy questionnaire. We documented three adaptive strategy progressions for those who found the optimal solution on scenario C: *Prospective Planning, Concurrent Planning,* and *Retrospective Planning.* These terms were used because each type of planning occurred at a clear timepoint during the scenario. Participants planned either prospectively (i.e., during attempt one), concurrently (i.e., throughout the scenario) or retrospectively (i.e., after attempt one).

These strategies are dynamic in that they were used at different time points over the course of the scenario. This is unlike general offloading (Risko & Gilbert, 2016) which is typically described as static. In most of these problem solving environments, participants have one attempt to find a solution, and thus one opportunity to use offloading. In contrast, the problem solving environment used in the current study allowed participants three attempts per scenario. Thus, participants could learn from prior attempts and adjust their plan on successive attempts to find the optimal solution. Those who used one of the three strategy progressions described how their strategy changed over the three scenario C attempts. Importantly, this was often in response to their increasing understanding of the task environment and its constraints.

Here, we propose that optimality may be better understood as a trend over time, and thus best investigated in studies that allow multiple attempts. Maximization (Smith, 1982) is used when the constraints are known and constant. The problem solver is assumed to have unbounded computational power, ensuring finding the solution with the highest net benefit. Satisficing (Simon, 1955) recognizes the bounds of human computational power. The problem solver stops their search once they meet a predefined threshold. Studies with goals of maximizing or satisficing tend to allow participants only one attempt to find a solution. The current study reveals that success at finding an optimal solution may be associated with experience with a task environment. Importantly, understanding of the most complex scenario improved over successive attempts, resulting in more participants finding the optimal solution on attempt three than attempt two, and on attempt two than attempt one; see Figure 2. Thus, investigations of optimal problem solving under constraints may be better suited for multi-attempt task environments. This might also promote research into adaptive use of strategies by examining how they change over time.

We close with a puzzle about offloading. Consistent with Risko and Gilbert (2016), we predicted that offloading would aid participants in finding the optimal solution by distributing their internal cognitive computations into the external environment. This supports Simon's (1955) proposal that human problem solving is bounded, both internally by computational limitations and externally by experience with the task environment. However, our study revealed that those who chose not to offload tended to outperform those who chose to offload; recall Figure 3. Furthermore, in scenario C, the most difficult scenario, none of the six participants who found the optimal solution on every attempt engaged in offloading. Nor did they did not use any of the provided materials during any attempt. This finding challenges Simon's theory of bounded rationality. It raises the question of whether enough experience with a task environment can result in humans no longer being bound by that environment's constraints. If this is true, offloading could potentially harm these individuals in terms of the time and effort it takes to utilize the external environment while problem solving (recall Zahner & Corter, 2010).

Future studies should further address this proposal using multi-attempt studies where each attempt has identical constraints. This research program could corroborate whether people are able to find the optimal solution more successfully over time (i.e., multiple attempts) and provide additional evidence that expertise with certain task environments might enable the human problem solver to find optimal solutions without the use of offloading.

Acknowledgments

We thank Erin Kingsley and Ashley Sasser for help with data collection and data coding.

References

Boddy, M., & Dean, T. L. (1994). Deliberation scheduling for problem solving in time-constrained environments. *Artificial Intelligence, 67(2)*, 245-285.

Bhui, R., Lai, L., & Gershman, S. J. (2021). Resource-rational decision making. *Current Opinion in Behavioral Sciences, 41*, 15-21.

Conlisk, J. (1996). Why bounded rationality?. *Journal of economic literature, 34(2)*, 669-700.

Conlisk, J. (1988). Optimization cost. *Journal of Economic Behavior & Organization, 9(3),* 213-228.

Hawkins, G. E., Brown, S. D., Steyvers, M., & Wagenmakers, E. J. (2012). An optimal adjustment procedure to minimize experiment time in decisions with multiple alternatives. *Psychonomic bulletin & review, 19(2)*, 339-348.

Howes, A., Duggan, G. B., Kalidindi, K., Tseng, Y. C., & Lewis, R. L. (2016). Predicting short-term remembering as boundedly optimal strategy choice. *Cognitive Science, 40(5),* 1192-1223.

Howes, A., Lewis, R. L., & Singh, S. (2014). Utility maximization and bounds on human information processing. *Topics in cognitive science, 6(2)*, 198-203.

Hutchins, E. (1995). How a cockpit remembers its speeds. *Cognitive science, 19(3)*, 265-288

Lewis, R. L., Howes, A., & Singh, S. (2014). Computational rationality: Linking mechanism and behavior through bounded utility maximization. *Topics in cognitive science, 6(2)*, 279-311.

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.*

Payne, J. W., Bettman, J. R., & Johnson, E. J. (1988). Adaptive strategy selection in decision making. *Journal of experimental psychology: Learning, Memory, and Cognition, 14(3)*, 534.

Pingle, M., & Day, R. H. (1996). Modes of economizing behavior: Experimental evidence. *Journal of Economic Behavior & Organization, 29(2)*, 191-209.

Pingle, M. (1992). Costly optimization: an experiment. *Journal of Economic Behavior & Organization, 17(1)*, 3-30.

Risko, E. F., & Gilbert, S. J. (2016). Cognitive offloading. *Trends in cognitive sciences, 20(9)*, 676-688.

Simon, H. A. (1955). A behavioral model of rational choice. *The quarterly journal of economics, 69(1*), 99-118.

Smith, J. F., Mitchell, T. R., & Beach, L. R. (1982). A cost-benefit mechanism for selecting problem-solving strategies: Some extensions and empirical tests. *Organizational Behavior and Human Performance, 29(3),* 370-396.

Tversky, B. (2013). Visualizing thought. *In Handbook of human centric visualization* (pp. 3-40). New York, NY: Springer New York.

Zahner, D., & Corter, J. E. (2010). The process of probability problem solving: Use of external visual representations. *Mathematical Thinking and Learning, 12(2)*, 177-204.