Does choosing the concept on which to solve each practice problem in an adaptive tutor affect learning?

Amruth N. Kumar [0000-0002-1951-3995]

Ramapo College of New Jersey, Mahwah NJ 07430, USA
amruth@ramapo.edu

Abstract. We conducted a controlled study to investigate whether having students choose the concept on which to solve each practice problem in an adaptive tutor helped improve learning. We analyzed data from an adaptive tutor used by introductory programming students over three semesters. The tutor presented code-tracing problems, used pretest-practice-post-test protocol, and presented line-by-line explanation of the correct solution as feedback. We found that choice did not increase the amount of learning or pace of learning. But, it resulted in greater improvement in score on learned concepts, and the effect size was medium.

Keywords: Choice of practice problem, Adaptive tutor, Controlled study.

1 Introduction

An adaptive tutor presents practice problems on the concepts students have not yet mastered. If the tutor requires students to select the concept on which they would like to solve each practice problem, it would give students choice, and thereby, a sense of agency [2], which is known to improve performance on learning tasks (e.g., [3]). It would give them a sense of control over their path through the learning material, the type of choice typically associated with enhanced learning [3,13]. In response, if the tutor honors the student's choice when presenting the next practice problem, it facilitates self-directed learning [6]. If it does not, i.e., it presents the next problem on a concept other than the one chosen by the learner, it promotes cognitive dissonance [4], which is known to help learning [1]. So, the act of choosing the practice concept might itself enhance learning, whether or not the tutor subsequently honors the choice. We investigated whether having the student choose the concept of each practice problem helped improve learning in an adaptive tutor.

For this study, we used a problem-solving software tutor on `for` loop, a programming concept. The tutor presents code-tracing problems, wherein, the student is asked to read a program and identify its cumulative output, one output at a time along with the line in the program that produces that output. The tutor provides line-by-line explanation of the correct answer as feedback, which has been shown to improve learning in prior evaluations [8]. This explanation is in the style of a worked example [12]. The tutor is adaptive [9] and covers 10 concepts in C++, Java or C#. It is part of a suite of problem-solving tutors for introductory programming topics called problets (www.problets.org).

The tutor is accessible over the web. Instructors of introductory programming courses use the tutor typically for after-class assignments. So, typically, students use the tutor on their own time, and often, multiple times till they master all the concepts. In this study, the institutions that used the tutor were randomly assigned to control or experimental group each semester. Data was collected over 3 semesters from Fall 2017 - Fall 2018. During that time, the number of students who used the tutor and granted permission for the use of their data in the study was 202 in control group and 179 in experimental group.

Every time the tutor was used, it administered pretest-practice-post-test protocol [7]. During pretest, the tutor presented one problem per concept to prime the student model. If the student solved a problem partially correctly, incorrectly, or opted to skip the problem without solving it, the tutor presented line-by-line explanation as feedback. Once the student had solved all the pretest problems, the tutor presented practice problems on only the concepts on which the student had skipped solving the problem or solved the problem partially/incorrectly during pretest. For each such concept, the tutor presented multiple problems until the student had mastered the concept, i.e., solved a minimum percentage (e.g., 60%) of the problems correctly. After each incorrectly solved problem, the tutor presented line-by-line explanation of the correct answer. Finally, during adaptive post-test, which was interleaved with practice, the tutor presented a test problem on each and only the concepts that had already been mastered by the student during practice. Pretest, practice and post-test were administered by the tutor back-to-back without interruptions, entirely over the web. The entire protocol was limited to 30 minutes.

The only difference in treatment between control and experimental groups was during adaptive practice. Before each practice problem, the tutor presented a list of all the concepts and the percentage of problems the student had solved correctly on each concept. The tutor used the same pre-determined order of concepts for both the groups. Experimental subjects were asked to pick the next concept that they wanted to practice, *but only when at least two concepts remained un-mastered*. In contrast, control subjects just viewed the list of concepts before moving on to the next problem. As a result, the sequence of practice problems solved by experimental subjects differed from that of control subjects. For control subjects, the tutor presented problems in the pre-determined order of concepts using round-robin algorithm, taking care not to present more than two problems back to back on any one concept. For experimental subjects, it used the subject's choice as the seed to pick the next concept not yet mastered in the pre-determined order of concepts, and presented a problem on it. *The resulting problem may or may not have been on the concept chosen by the student*. According to learning theory, the student could benefit whether the practice problem matched the chosen concept or not: when the two matched, the student could benefit from cueing [11] and self-directed learning [6]. When the two mismatched, the student could benefit from cognitive dissonance [1].

The concepts covered by the tutor can be classified as known, attempted, practiced or learned by the student. A concept is **known** if the student solves the pretest problem on the concept correctly. A concept is **learned** if the student solves the pretest problem on the concept partially/incorrectly or skips solving it, solves enough problems during practice to master it, and solves the post-test problem correctly. On the other hand, if the student solves the post-test problem incorrectly, the concept is **practiced**, not yet learned. If so, the tutor schedules additional practice problems on the concept for the student. If the student runs out of time because of the 30-minute limit placed on the duration of the tutoring session and does not complete mastering the concept during practice, the concept is categorized as **attempted**.

During grade calculation of code-tracing problems, the outputs identified in the correct sequence (c) were credited and any incorrectly identified outputs thereafter (i) were penalized. The grade was calculated as max((c – i) / n, 0), where n was the total number of outputs in the program. Therefore, the score on each problem was normalized to 0 → 1.0.

If a student used a tutor multiple times, we considered data from the session when the student had learned the most number of concepts. If the student did not learn any concepts, we considered data from the first session when the student had solved the most number of problems. Since the only difference in treatment between control and experimental groups was during practice stage, and that too, when students solved problems on two or more concepts, only students who had solved practice problems on two or more concepts were retained for the study in both control and experimental groups. As a result, 98 students each remained in control and experimental groups.

We considered 7 variables for the study: (1) The **score per pretest problem** to verify that control and experimental groups were comparable; (2) The **number of practice problems** solved during the session. This included practice problems solved on learned concepts as well as concepts merely practiced or attempted. This and the next two variables were used to evaluate the impact, if any, of having to choose the concept before each practice problem; (3) The **score per practice problem**; (4) The **time spent per practice problem**; (5) The **number of concepts learned** as a measure of the amount of learning; (6) The **number of practice problems solved per learned concept**, as a measure of the pace of learning. It was calculated by dividing the number of practice problems solved on all the learned concepts by the number of concepts learned; and (7) **Pre-post change in grade per problem on learned concepts** as a measure of improvement in learning. The fixed factor was treatment: experimental group students had to choose the concept underlying each practice problem, whereas control group students did not.

2 Results and Discussion

We conducted one-way ANOVA of each of the variables with treatment as the fixed factor.

We found no significant main effect for treatment on the **score per pretest problem** [(F1,195) = 1.841, p = 0.176] or the time spent per pretest problem [F(1,195) < 0.001, p = 0.991]. *So, the two groups were comparable*. We did not find any significant difference in the **number of practice problems** solved [F(1,195) = 0.991, p = 0.321], but the **score per practice problem** was significantly different [F(1,195) = 7.897, p = 0.005]: control group subjects scored a mean of 0.835 ± 0.029 points per practice problem whereas experimental subjects scored 0.776 ± 0.029 points. *So, control subjects scored significantly more per practice problem.* One explanation for this difference is that experimental subjects suffered cognitive dissonance when the problem they were presented was not on the concept they chose. Hence, they scored significantly less on practice problems. We found no significant main effect for treatment on the **time spent per practice problem**.

We did not find a significant difference between control and test groups on **the number of concepts learned**. *So, the treatment did not affect the amount of learning.* We found a significant main effect for treatment on the **number of practice problems solved per learned concept** [F(1,109) = 4.965, p = 0.028]: control subjects solved a mean of 2.92 ± 0.249 problems to learn each concept whereas experimental subjects solved 3.30 ± 0.223 problems. *So, the pace of learning was significantly slower for experimental subjects than control subjects.* This may also be ascribed to cognitive dissonance: the practice problem presented by the tutor matched the concept chosen by the student in only 40.65% of the cases. *So, the treatment of merely providing choice without always honoring it did not benefit the pace of learning.*

We also found a significant main effect for treatment on **pre-post change in grade per problem on learned concepts** [(F1,109) = 5.028, p = 0.027]: the change was 0.716 ± 0.056 for control subjects compared to 0.802 ± 0.051 for experimental subjects. *So, the improvement in score was significantly greater for experimental group than control group.* The effect size (Cohen's d) is 0.43, corresponding to medium effect.

To summarize, even though the two groups were comparable to begin with, experimental group needed significantly more practice problems to learn each concept, but had significantly greater improvement in score on the learned concepts. We found the same results when we considered only less-prepared students, i.e., those whose score per pretest problem was 0.9 or less. Insofar as choice enhances intrinsic motivation and engenders agency, the results of this study warrant the provision of choice, given that we did not find any negative cognitive effects of choice.

Experimental subjects chose the first concept in the list presented to them 65.25% of the time, and second concept 22.88% of the time. Given that the list always contained at least two concepts, this lopsided distribution of choice suggests that students more often than not chose concepts in the order in which they were presented.

Cueing has been shown to improve transfer of learning [11], in particular, when multiple problem-solving examples are presented to the learner [5,10]. In our experimental setup, the name of the concept could have served as a consistent, valid cue if the practice problem had always matched the student's choice. We speculate that providing choice, combined with consistent, valid cueing might lead to better learning outcomes than were observed in this study. This will be the subject of a future study.

Acknowledgments

Partial support for this work was provided by the National Science Foundation under grant DUE-1432190.

References

1. Aïmeur, E., Frasson, C. and Lalonde, M.: The role of conflicts in the learning process. SIGCUE Outlook, vol. 27 (2). pp. 12-27. (2001).
2. Bandura, A.: Social cognitive theory: an agentic perspective. Annual Review of Psychology, vol. 52(1), pp. 1-26. (2001).
3. Cordova, D., and Lepper, M.: Intrinsic motivation and the process of learning: Beneficial effects of contextualization, personalization, and choice. Journal of Educational Psychology, vol. 88, pp. 715–730. (1996).
4. Festinger, L. and Carlsmith, J.M.: Cognitive consequences of forced compliance. Journal of Abnormal and Social Psychology, vol. 58 (2), pp. 203-210. (1959).
5. Gentner, D., Loewenstein, J., and Thompson, L.: Learning and transfer: A general role for analogical encoding. Journal of Educational Psychology, vol. 95(2), pp. 393–408. (2003).
6. Hammond, M. and Collins, R.: Self-directed learning: Critical practice. London: Kogan Page Limited. (1991).
7. Kumar, A.N.: A Model for Deploying Software Tutors. In Proc. IEEE 6th International Conference on Technology for Education (T4E), Amritapuri, India, pp. 3-9. (2014).
8. Kumar, A.N.: Explanation of step-by-step execution as feedback for problems on program analysis, and its generation in model-based problem-solving tutors. Technology, Instruction, Cognition and Learning. (TICL) J. Special Issue on Problem Solving Support in Intelligent Tutoring Systems, vol. 4(1). (2006).
9. Kumar, A.N.: A Scalable Solution for Adaptive Problem Sequencing and its Evaluation. In Proc. The Fourth International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems (AH 06), Dublin, Ireland, pp. 161-171. (2006).
10. Norman, G., Dore, K., Krebs, J., and Neville, A. J.: The power of the plural: Effect of conceptual analogies on successful transfer. Academic Medicine, vol. 82(10), S16-8. (2007).
11. Speicher, T., Bell, A. Kehrhahn, M., and Casa, D.: Effect of Cueing on Learning Transfer Among Health Profession Students Engaged in a Case-Based Analogical Reasoning Exercise. The Internet Journal of Allied Health Sciences and Practice, vol. 12(3), Article 4. (2014).
12. Sweller, J., and Cooper, G.A.: The use of worked examples as a substitute for problem solving in learning algebra. Cognition and Instruction, vol. 2, pp. 59–89. (1985).
13. Zuckerman, M., Porac, J., Lathin, D., Smith, R., and Deci, E.: On the importance of self-determination for intrinsically-motivated behavior. Personality and Social Psychology Bulletin, vol. 4, pp. 443–446. (1978).