International Society of the Learning Sciences

Promoting Students' Self-Regulated Learning Choices with Diagrams in Intelligent Tutoring Software

Tomohiro Nagashima, Saarland University, Saarland Informatics Campus, nagashima@cs.uni-saarland.de
Bin Zheng, Carnegie Mellon University, binzheng@andrew.cmu.edu
Stephanie Tseng, Carnegie Mellon University, stseng2@andrew.cmu.edu
Elizabeth Ling, Harvard University, elizabethling@college.harvard.edu
Vincent Aleven, Carnegie Mellon University, aleven@cs.cmu.edu

Abstract: Although students' self-regulated learning has been studied extensively, past research has not investigated students' fine-grained, self-regulated choice-making processes during learning with visual representations and strategies to support such processes. We conducted design and experimental studies with 148 students to develop and evaluate an intervention package for supporting students' self-regulated choice-making in using diagrammatic scaffolding in algebra tutoring software. A classroom experiment showed that students with the intervention learned greater conceptual and procedural knowledge in algebra than students in the control condition whose choices were not supported. Also, students with the intervention chose to use diagrams less frequently overall but showed distinctive use patterns that changed over time, indicating a form of self-regulated diagram use. This study demonstrates the importance of understanding and supporting choice behaviors that change over time during learning, going beyond simply measuring the frequency of choice behaviors and encouraging students to engage in these behaviors more frequently.

Introduction

In modern society, where we have access to abundant information and countless resources, learners need to proactively and strategically choose to use available resources so that they can handle tasks effectively and efficiently (Schwartz & Arena, 2013). Thus, one goal of education is to foster learners who can make strategic choices in using resources (Chin et al., 2019; Cutumisu et al., 2019).

One important choice for learners to make involves whether and how to use visual representations during learning activities. Visual representations are often used as instructional scaffold that can facilitate learners' sense-making processes during learning and problem solving. From a cognitive perspective, visual representations can help learning by, for example, reducing cognitive effort and making relevant information salient (Ainsworth, 2006). However, to use visual representations effectively, learners must (come to) understand when it is appropriate to proactively use visual representations, so that they can navigate problem solving effectively and efficiently (e.g., over-use of visual scaffolding might not help learning very much) (Schwonke et al., 2013). Acquiring such understanding is critical since learners will not always be presented with visual information in everyday problem-solving situations. Instead, they need to proactively choose to (or not to) use/create visual information to aid their problem solving when it is appropriate. Therefore, to use visual representations strategically, learners need to make self-regulated choices. In doing so, they need to judge if the use of visual scaffolding would help their problem solving or not, and make a choice accordingly. Such self-regulated choice behaviors could presumably lead to enhanced domain-level learning (Long & Aleven, 2016; Roll et al., 2011).

Despite the importance of understanding learner choices with visual representations, past research has rarely allowed for, and measured, these choices. Studies have mostly focused on learning *with* visual representations and student learning when students are *given* visual representations (e.g., Rau et al., 2015). A few studies investigated learners' spontaneous use of visual representations in problem solving (e.g., Uesaka et al., 2010). These studies found that, when the use of a visual representation is introduced as an option during problem solving, students generally tend not to choose to use it. Prompts can help learners use visual information more frequently (Wu et al., 2020), but these past studies use aggregated data in measuring choices (e.g., calculating the sum of choice behaviors) or post-hoc self-report data, making it unclear what specific choice behaviors students engaged in and how their choice patterns *changed over time* during their learning, a critical aspect of self-regulated learning (Greene et al., 2021; Roscoe et al., 2013). As learners develop skills and knowledge during learning, it is reasonable to expect that their use of certain strategies may change over time (Greene et al., 2021).

This paper reports on a design study and an experimental study we conducted with a total of 148 middle-school students in the U.S. to generate new scientific knowledge regarding learner's self-regulated choice behaviors during learning with visual representations, how these behaviors change over time, and how they might be supported with technology. We employed a user-centered design approach with eight students to create an intervention package for supporting students in self-regulated diagram use. This design phase was followed by an experimental study with 140 middle-school students in actual school classrooms in the U.S. that tests the learning that results from this intervention, when added to an intelligent tutoring system for algebra problem solving.

The Diagram Choice Tutor

To investigate and appropriately measure students' choices in using visual representations, we designed and developed the *Diagram Choice Tutor*, an Intelligent Tutoring System for middle-school algebra in which students can choose, for each problem-solving step, whether or not to use a visual representation (called "tape diagrams," Murata, 2008) to help with their problem solving. As shown in Figure 1, when students opt to use diagrams, they need to select from three diagrams that depict what to do in the next step, one optimal diagram and two suboptimal or incorrect ones. This anticipative form of interactive visual scaffolding has been shown to support learning and problem-solving performance by having learners *self-explain* their problem-solving steps in a visual form, before they do these steps in the standard format (in this case, symbolic equations) (Nagashima et al., 2021).

The Diagram Choice Tutor presents an appropriate learning environment for learners to exercise self-regulated choice behaviors. Chin et al. (2019) argue that an ideal choice-based learning environment presents learners with options they would not select naturally (e.g., seeking negative feedback, Cutumisu et al., 2015). In the Diagram Choice Tutor, students (aged around 10-15) may naturally avoid using diagrams because doing so would require additional problem-solving steps, and tape diagrams introduce a new representation that students may not be familiar with (Murata, 2008). However, engaging with diagrammatic scaffolding during algebra learning has been shown beneficial for student learning, even with the expected cognitive effort required (Nagashima et al., 2021). Our Diagram Choice Tutor is also instrumented to collect students' learning process data to help overcome the lack of understanding regarding how students make choices.

Figure 1
In the Diagram Choice Tutor, (a) students start with an interface that shows a symbolic problem. They can choose to request a diagrammatic scaffold or not. (b) If they choose to use the scaffold for the given step, they then need to select, from three options, the diagram that shows the correct and optimal next step. After selecting the correct diagram option (e.g., the one on the left in the figure), students are prompted to do the step symbolically by referring to the diagrammatic scaffold that they have. (c) Students can also choose not to use diagrams to solve the step. This process repeats (i.e., students have the choice for every problem-solving step).

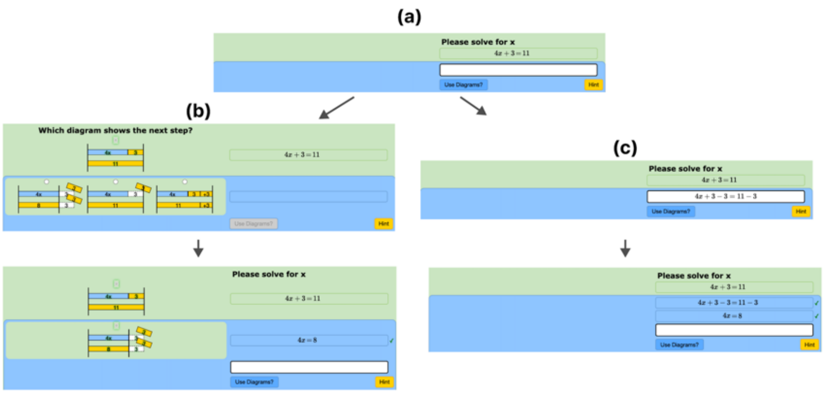

Designing Support for Self-Regulated Diagram Use

Idea generation with students

Promoting self-regulated choices in the Diagram Choice Tutor presents a challenging design problem. It is difficult to define a criterion for when students should use diagrams (i.e., for when diagrams are most helpful for learning), as opposed to, for example, simply promoting more frequent use of diagrams. Drawing on the literature on help seeking in Intelligent Tutoring Systems, one might surmise that self-regulated learners would ideally neither over-use nor under-use the visual scaffolding but rather use it *when needed* (i.e., when the use might be most beneficial, Aleven et al., 2016). Although more frequent use of visual scaffolding means that students would get more exposure to the advantages that the visual representation has on learning (e.g., it might promote conceptual learning) (Uesaka et al., 2010), in the context of our tutor, students could overly rely on using the visual scaffolding in solving equations. Such over-use of diagrams in our tutor may lead to the acquisition of rather superficial "diagram-to-symbols translation knowledge" (i.e., copying what is shown in a diagram into the symbolic problem-solving step in the tutor, Nagashima et al., 2022, p. 1752). Because, in our Diagram Choice Tutor, tape diagrams as instructional scaffolding supplement the canonical representation (i.e., algebraic equations in symbolic notation), such superficial knowledge may not help learners when solving more advanced equation problems (for which tape diagrams are no longer useful, e.g., equations with negative numbers). On the other hand, self-regulated learners may use the visual scaffold to understand how to solve equations that they are not familiar with, but as they practice more, they might choose to practice their problem-solving/procedural skills without relying on the visual aid too much (Aleven et al., 2016). During this learning process, they may actively engage in key iterative stages of self-regulation (Zimmerman & Campillo, 2003), which are comprised of: self-assessment ("Can I solve this problem without help?"), self-monitoring ("Am I doing well with/without diagrams?"), and self-reflection activities ("How well did I do with/without diagrams?").

To approach this design challenge, we conducted one-on-one virtual idea-generation sessions with eight school students in the U.S. (one 4th grade, one 5th grade, one 6th grade, four 7th grade, and one 8th grade). In each session, the students first practiced a few problems with the Diagram Choice Tutor. Then, sharing a virtual whiteboard, the researcher and the student generated ideas in response to several prompting questions on self-regulated use of diagrams in the tutor (e.g., "What would be some features that would help you think carefully about whether or not to use diagrams for solving equations?"). A total of approximately eight hours of video recordings from the sessions were analyzed by three researchers using the Affinity Diagramming method. Affinity Diagramming is a standard technique used in the field of Human-Computer Interaction to analyze qualitative data to produce shared themes through multi-step synthesizing of codes and themes (Lucero, 2015). This process produced 117 codes, which we grouped into 13 themes. Then, of the 13 themes, those that share similar ideas were grouped together, resulting in five high-level ideas for promoting self-regulated diagram use in the tutor. Due to the page limit, we present just a short statement of four of the five main ideas that directly informed the design of our intervention: (Idea 1) *Tell me that diagrams are there to help, they are not there for no reason* (i.e., students want to understand how diagrams can be useful), (Idea 2) *I want to be prompted to consider using diagrams when they can be useful* (i.e., students want to know and be reminded when diagrams can be useful), (Idea 3) *Show me how diagrams are helping or not helping me* (i.e., students want to know if diagrams help their own problem solving and learning), and (Idea 4) *A diagram badge can help me think about using diagrams* (i.e., students want motivational features such as a badge for using diagrams).

Designing an intervention package for supporting self-regulated diagram use

Based on the four ideas generated by students, we (researchers) designed an intervention *package*, which consists of (a) an interactive tutorial that teaches students how diagrams can be useful (to address Idea 1), (b) an adaptive recommendation pop-up screen that prompts students to think about whether or not to use diagrams when they seem to have trouble solving problems (to address Idea 2), and (c) a student-facing learning analytics dashboard that shows how well students have been performing with and without using diagrams through visualization and badges (to address Ideas 3 and 4). See Figure 2 for more information.

Figure 2

(a) An interactive tutorial that explains how to use diagrams and research evidence on the benefits of diagrams on problem solving. (b) Adaptive recommendations, prompting students to think about diagram use, appear after any three consecutive problem-solving mistakes, after pausing for more than 90 seconds, and on the first problem in each problem set/level (more on problem levels described later). (c) A personalized learning dashboard that presents a graph showing the student's problem-solving performance (i.e., percent correct) when they used diagrams (in blue) and when they did not use diagrams (in red) in the most recent problem level. Students are asked to answer a 5-scale "smiley" question on how they feel about the usefulness of diagrams. The dashboard also provides badges for students, based on their problem-solving performance and the use of diagrams.

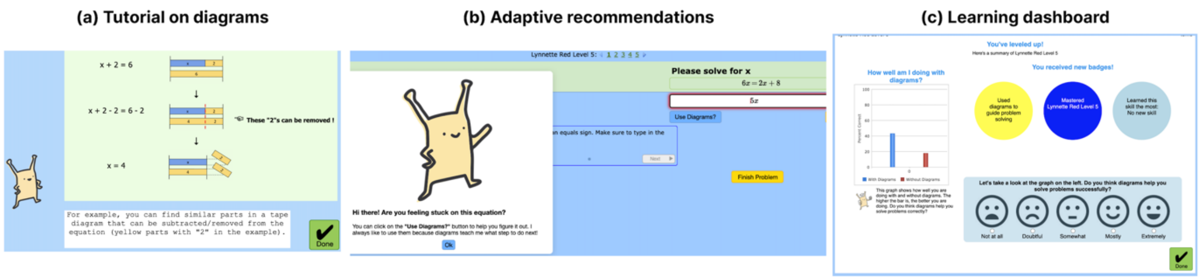

We designed these multiple intervention components and combined them as an intervention *package* to support the different stages of self-regulation mentioned earlier, rather than aiming to design one intervention. This decision was made because students expressed various needs and ideas for using diagrams during the idea-generation sessions. It is also suggested in the literature on self-regulated learning (Oppezzo & Schwartz, 2013). Specifically, (a) the tutorial is meant to help students think whether diagrams are useful for them (*self-assessment*), (b) the adaptive recommendations are to help students during the *self-monitoring*, and (c) the dashboard is designed to help students *self-reflect* on their choice behaviors of using or not using diagrams.

Method: Classroom Study

We then conducted a controlled classroom experiment to test 1) whether the intervention package helps students gain better domain-level knowledge and skills in algebra and 2) whether and how the intervention package helps students demonstrate self-regulated use of diagrams. We compared two conditions in which students practiced algebra problem solving with the Diagram Choice Tutor (i.e., students in both conditions had control over whether to use diagrams for each problem-solving). The conditions differed in whether students had access to the intervention components (Supported Choice condition) or not (Unsupported Choice condition).

Participants

A total of 179 students participated in the study in their in-person classroom (38 5th graders, 37 6th graders, 86 7th graders, and 18 8th graders). Participants came from 11 classes in two schools in the U.S., taught by two teachers. Students in each class were randomly assigned to either the Supported Choice condition (n = 87) or the Unsupported Choice condition (n = 92). The participating teachers noted that their students' experience with tape diagrams was minimal (e.g., their instruction had never focused on tape diagrams).

Materials

All students in both conditions in the study used the Diagram Choice Tutor (Figure 1) to practice equation solving during the study. Students in both conditions were assigned the same sets of problems (Table 1). Students in the Supported Choice condition additionally had access to the self-regulated learning (SRL) intervention components (Figure 2) embedded in the tutor. In both conditions, students had control over when to use diagrams. The only difference was whether or not students had the additional SRL components to support their choice making. Figure 3 illustrates how we integrated the intervention components into the Diagram Choice Tutor.

Table 1
Types of equation problems assigned in the tutor (in both conditions)

Problem level	Problem type	Problem level	Problem type
1	x + a = b	7	ax + b = cx + d
2	ax + b = c	8	ax + b = c (bonus content)
3	ax + b = c	9	ax = bx + c (bonus content)
4	ax = bx + c	10	ax + b = cx + d (bonus content)
5	ax = bx + c	11	ax = bx + c (bonus content)
6	ax + b = cx + d	12	ax + b = cx + d (bonus content)

Figure 3
Students in both conditions had the same practice problems from Level 1 to Level 12. Students in the Supported Choice condition additionally engaged with (a) the tutorial on diagrams before Level 1, and at the end of each Level, they were shown (c) the learning dashboard. (b) The adaptive recommendations were available from Level 1 to Level 12. Students in the Unsupported Choice condition solved algebra problems from Level 1 to Level 12 with no SRL intervention. Students in both conditions spent the same total time in the study.

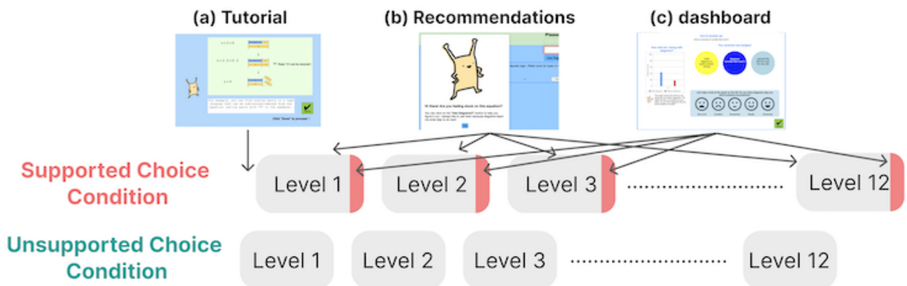

We developed a web-based pretest and posttest on conceptual knowledge and procedural knowledge in early algebra based on items in the literature (e.g., Booth et al., 2013). Each test had 16 multiple-choice conceptual knowledge items and five open-ended procedural knowledge items. Two isomorphic versions of the test were developed and assigned in a counter-balanced way across pretest and posttest.

Procedure

The study was conducted during the schools' regular class periods across five consecutive days. Researchers joined the class through a video conferencing system to facilitate the study sessions. Students in both conditions first completed the web-based pretest. Then, they watched a brief video illustrating how to use the tutor. From the second day up to (and including) the fourth day, students solved algebra problems using the Diagram Choice Tutor. Students in both conditions spent the same total time in the study sessions; students in the Unsupported Choice condition spent the time exclusively on equation solving, while those in the Supported Choice condition spent it on equation solving *and* the SRL intervention components combined (Figure 3). On the final day, students completed the web-based posttest. After the posttest, all students were given access to both versions of the tutor so that they could experience the software that had been used in both conditions.

Results

Of the 179 students, 168 students completed all parts of the study. We excluded students who scored 100% on the pretest and those who did not complete more than 50% of the test items on the pretest and/or posttest, decided before testing treatment effects (Supported Choice condition: $n = 4$, Unsupported Choice condition: $n = 4$, Chan et al., 2022). Further, after a discussion with the participating teachers on students' exposure to equation solving prior to the study, all students from two advanced classes were excluded from the sample (Supported Choice: $n = 11$, Unsupported Choice: $n = 9$; these students were originally included in the study as the teachers wanted them to experience a research study, which they said they would not experience in the regular course of schoolwork). The final sample consisted of 140 students (Supported Choice: $n = 69$, Unsupported Choice: $n = 71$). No statistically significant difference was observed between the conditions in the dropout/exclusion rate, $\chi 2$ (1, N = 179) = .05, $p = .92$.

How did the intervention influence students' learning outcomes?

Table 2 shows students' mean pretest and posttest scores on conceptual knowledge and procedural knowledge in algebra. It also shows the average total number of algebra problems students solved in the tutor within the same time given to both conditions. To examine if the intervention package enhanced students' conceptual and procedural learning, we ran two separate linear regressions with condition as an independent variable, and conceptual and procedural knowledge as a dependent variable, respectively. In both models, students' pretest score (on conceptual and procedural knowledge, respectively) was added as a covariate to control for students' prior knowledge before the study. We found that students in the Supported Choice condition learned greater conceptual ($\beta = .95$, $t(137) = 2.52$, $p = .01$) and procedural knowledge ($\beta = .53$, $t(137) = 2.08$, $p = .04$) than those in the Unsupported Choice condition.

We also compared, between the conditions, the number of problems solved. We used a linear regression model with the same independent variable, and combined pretest score (conceptual and procedural) as a covariate, but with the average number of problems solved as the dependent variable. The model showed no statistically significant difference on the average number of problems solved in the tutor, $\beta = .67$, $t(137) = .32$, $p = .75$.

Table 2
Students' pretest and posttest scores (standard deviations in parentheses), and the number of problems solved

Condition	Conceptual Knowledge (*max* = 16)		Procedural knowledge (*max* = 5)		Number of problems solved
	pretest	posttest	pretest	posttest	
Supported	9.78 (2.72)	10.48 (2.49)	1.74 (1.74)	2.88 (2.03)	32.64 (13.52)
Unsupported	9.28 (2.22)	9.28 (2.15)	1.73 (1.74)	2.35 (1.88)	31.25 (13.10)

How did the intervention influence students' self-regulated diagram use?

To investigate students' diagram use across the two conditions, we examined tutor log data on the frequency of diagram use (i.e., how many times students requested to use diagrams). On average, students in the Supported Choice condition chose to use diagrams 0.12 times ($SD = .10$) per problem-solving step while those in the Unsupported Choice condition used diagrams 0.16 times ($SD = .12$). A two-sample t-test with the condition as an independent variable showed that students in the Supported Choice condition, on average, chose to use diagrams less frequently, $t(135.3) = 2.03$, $p = .04$.

To further investigate how students' choice behaviors *changed over time*, we then looked at the frequency of diagram use for each problem in the tutor. Figure 4 shows the proportions of students who chose to use diagrams at least once on each problem in the tutor, grouped by problem level. The graph reveals several notable patterns in their choice behaviors. First, many students in both conditions chose to use diagrams when they saw a new problem type (the first problem in Levels 1, 2, 4, and 6). Second, for the first problem in Levels 3, 5, and 7, in which students saw the same type of problems that they practiced in the previous level, students in the Supported Choice condition show a relatively low level of diagram use whereas at these levels students in the Unsupported Choice condition used diagrams at the rate that is not very different from their diagram use in the previous levels (Levels 2, 4, and 6, respectively). Finally, students in the Supported Choice condition show significant drops in their use of diagrams from the first to the second problem in Levels 1, 2, 4, and 6. This pattern indicates that many of the students in the Supported Choice condition chose not to use diagrams when they solved the same type of problem for the second time. This trend is also observed in the Unsupported Choice condition but is not as pronounced as it is in the Supported Choice condition (e.g., most notably in Level 6).

Figure 4

Proportions of students who chose to use diagrams at least once on each problem across problem levels. Only Levels 1-7 are shown, as Levels 8-12 were bonus problems with the same content introduced in Levels 1-7.

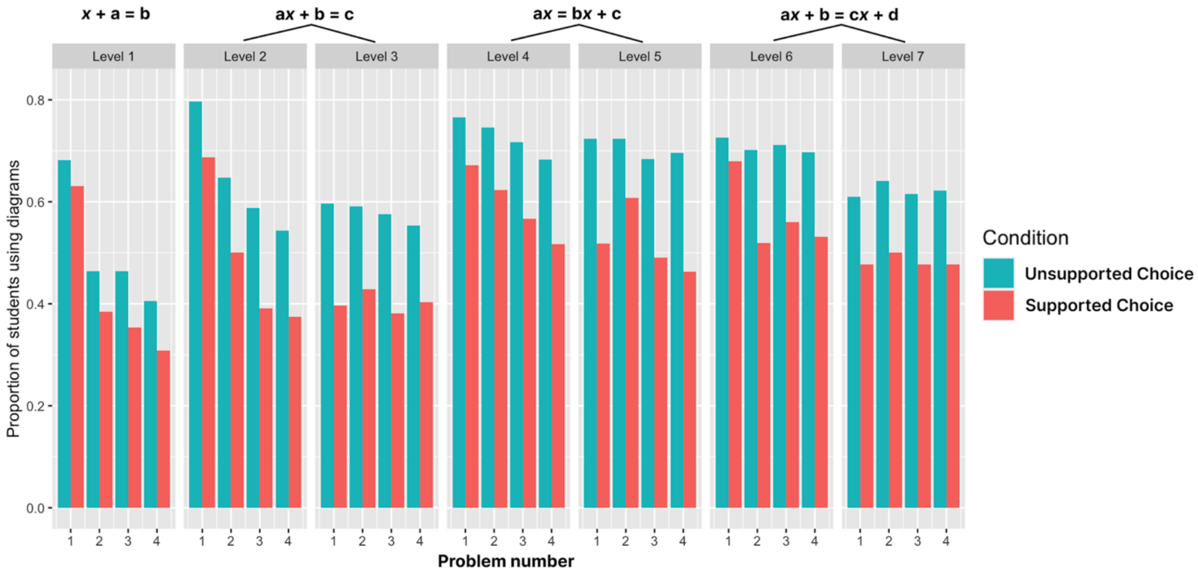

Discussion and Conclusion

Prior work on choice behaviors in using visual representations does not offer insights into students' *choice behaviors that change over time* during learning processes despite its importance in understanding self-regulated learning processes. In our study, we first worked with middle-school students to design several intervention components to support self-regulation during learning with diagrams in the Diagram Choice Tutor. These components included a tutorial on benefits of using diagrams, adaptive recommendations that encourage students to think about using diagrams or not, and a personalized learning dashboard showing the student's recent problem-solving performance with and without diagrams. We then conducted a classroom experiment to evaluate the effectiveness of the intervention. We found that students who received the intervention learned greater conceptual and procedural knowledge, even though they chose to use diagrams less frequently overall than those without the intervention. How could less diagram use lead to greater learning in our study?

A closer look at students' changing choice behaviors in the tutor log data allowed us to understand what one form of self-regulated use of diagrams might look like. Specifically, the data suggested that students in the Supported Choice condition used diagrams more frequently when they saw a new problem type (at the outset, many students were not familiar with the problem types used in the study, according to participating teachers) but chose to use them less frequently when they kept seeing the same type of problem in the tutor. On the other hand, this pattern was not observed for students in the Unsupported Choice condition, where students made choices without the help of the intervention. These insights suggest that students whose choices were supported used diagrams when the use of diagrams may have been most helpful for learning, neither over-using nor under-using them. We conjecture that students in the Supported Choice condition, by engaging with the intervention components, were able to monitor and reflect on their own use of diagrams and were able to choose to use diagrams when they thought doing so would be helpful. For instance, it is possible that students, when seeing the graph on the dashboard, were able to deeply reflect on how useful it is to use diagrams in each problem level and on whether or not to use diagrams on the next level. Such informed choice-making practices may have contributed to the greater conceptual and procedural learning that was observed for the students in the Supported Choice condition. They focused their practice on learning how to solve new types of problems with diagrams (which may have led to conceptual learning) and then chose to practice problems without relying on the visual aid too much (which may have led to procedural learning). These insights would not have been gained only with aggregated data points but were possible with the temporally fine-grained log data collected with the technology. Such learning processes might also have contributed to faster, more efficient problem solving. Despite several additional activities (e.g., the tutorial and dashboard) that students in the Supported Choice condition had to complete (Figure 3), no difference was found on the average total number of problems completed in the tutor between the conditions.

Due to the design of the experiment, it is not possible to tease apart the effects of individual intervention components on student learning and diagram use. However, we argue that the design of the intervention as a *package* was aligned with and promoted different stages of self-regulation, as opposed to designing and testing a single intervention element. This decision was informed both by idea generation sessions with school students and a theoretical view on self-regulation. Still, we acknowledge that these speculations cannot be fully validated using the data from the current study and that other interpretations are possible. Also, the changing choice behaviors we observed in the tutor were not validated with statistical significance testing. Finally, self-regulated use of diagrams can take other forms than that addressed in the current study, depending on the context and domain of diagram use (e.g., self-constructing diagrams, Uesaka & Manalo, 2006; Uesaka et al., 2010).

The current study makes several contributions to the field of the learning sciences. First, the study contributes novel insights into how students' choice behaviors change over time, when students can choose to use or not use diagrams in the context of problem-solving practice in an intelligent system. Aggregated measures of choice behaviors (e.g., mean frequency of diagram use) do not afford the same insight. This insight is key to understanding and supporting self-regulated choice behaviors that lead to greater domain learning. The study also illustrates, by investigating students' choice behaviors, what effective use of diagrams during learning might look like. Further, we demonstrate that an intervention informed by ideas generated by students supports not only an effective self-regulated behavior but also greater domain learning, which is not typically achieved by interventions designed to support self-regulated learning with technology (but see Long & Aleven, 2016).

References

Ainsworth, S. (2006). DeFT: A conceptual framework for considering learning with multiple representations. *Learning and Instruction, 16*(3), 183-198.

Aleven, V., Roll, I., McLaren, B., & Koedinger, K. R. (2016). Help helps, but only so much: Research on help seeking with Intelligent Tutoring Systems. *International Journal of Artificial Intelligence in Education, 26*, 205-223.

Booth, J. L., Lange, K. E., Koedinger, K.R., & Newton, K. J. (2013). Using example problems to improve student learning in algebra: Differentiating between correct and incorrect examples. *Learning and Instruction, 25*, 24-34.

Chan, J. Y. C., Lee, J. E., Mason, C. A., Sawrey, K., & Ottmar, E. (2022). From Here to There! A dynamic algebraic notation system improves understanding of equivalence in middle-school students. *Journal of Educational Psychology, 11*4(1), 56-71.

Chin, D. B., Blair, K. P., Wolf, R. C., Conlin, L. D., Cutumisu, M., Pfaffman, J., & Schwartz, D. L. (2019). Educating and measuring choice: A test of the transfer of design thinking in problem solving and learning. *Journal of the Learning Sciences, 28*(3), 337-380.

Cutumisu, M., Chin, D. B., & Schwartz, D. L. (2019). A digital game-based assessment of middle-school and college students' choices to seek critical feedback and to revise. *British Journal of Educational Technology, 50(*6), 2977-3003.

Greene, J. A., Plumley, R. D., Urban, C. J., Bernacki, M. L., Gates, K. M., Hogan, K. A., Demetriou, C., & Panter, A. T. (2021). Modeling temporal self-regulatory processing in a higher education biology course. *Learning and Instruction, 72*, 101201.

Long, Y. & Aleven, V. (2016). Mastery-oriented shared student/system control over problem selection in a linear equation tutor. In A. Micarelli, J. Stamper, & K. Panourgia (Eds.), *Proceedings of the 13th International Conference on Intelligent Tutoring Systems* (pp. 90-100). Springer, Cham.

Lucero, A. (2015). Using affinity diagrams to evaluate interactive prototypes. In *IFIP Conference on Human-Computer Interaction* (pp. 231-248). Springer, Cham.

Murata, A. (2008). Mathematics teaching and learning as a mediating process: The case of tape diagrams. *Mathematical Thinking and Learning, 10*(4), 374-406.

Nagashima, T., Bartel, A. N., Yadav, G., Tseng, S., Vest, N. A., Silla, E. M., Alibali, M. W., & Aleven, V. (2021). Using anticipatory diagrammatic self-explanation to support learning and performance in early algebra. In E. de Vries, J. Ahn, & Y. Hod (Eds.), *15th International Conference of the Learning Sciences – ICLS 2021* (pp. 474-481). International Society of the Learning Sciences.

Nagashima, T., Ling, E., Zheng, B., Bartel, A. N., Silla, E. M., Vest, N. A., Alibali, M. W., & Aleven, V. (2022). How does sustaining and interleaving visual scaffolding help learners? A classroom study with an Intelligent Tutoring System. In *Proceedings of the 44th Annual Meeting of the Cognitive Science Society.* Cognitive Science Society (pp. 1751-1758).

Oppezzo, M., & Schwartz, D. L. (2013). A behavior change perspective on self-regulated learning with teachable agents. In *International Handbook of Metacognition and Learning Technologies (pp. 485-500)*. Springer, New York, NY.

Rau, M. A., Aleven, V., & Rummel, N. (2015). Successful learning with multiple graphical representations and self-explanation prompts. *Journal of Educational Psychology, 107*(1), 30.

Roll, I., Aleven, V., McLaren, B. M., & Koedinger, K. R. (2011). Improving students' help-seeking skills using metacognitive feedback in an intelligent tutoring system. *Learning and Instruction, 21*(2), 267-280.

Roscoe, R. D., Segedy, J. R., Sulcer, B., Jeong, H., & Biswas, G. (2013). Shallow strategy development in a teachable agent environment designed to support self-regulated learning. *Computers & Education, 62*. 286-297.

Schwartz, D. L. & Arena, D. (2013). *Measuring what matters most: Choice-based assessments for the digital age*. The MIT Press.

Schwonke, R., Ertelt, A., Otieno, C., Renkl, A., Aleven, V., & Salden, R. J. (2013). Metacognitive support promotes an effective use of instructional resources in intelligent tutoring. *Learning and Instruction, 23*, 136-150.

Uesaka, Y., & Manalo, E. (2006). Active comparison as a means of promoting the development of abstract conditional knowledge and appropriate choice of diagrams in math word problem solving. In D. Barker-Plummer, R. Cox, & N. Swoboda (Eds.), *Diagrammatic Representation and Inference. Diagrams 2006. Lecture Notes in Artificial Intelligence,* 4045 (pp. 181–195). Springer. Berlin, Heidelberg.

Uesaka, Y., Manalo, E., & Ichikawa, S. I. (2010). The effects of perception of efficacy and diagram construction skills on students' spontaneous use of diagrams when solving math word problems. In *International Conference on Theory and Application of Diagrams* (pp. 197-211). Springer, Berlin, Heidelberg.

Wu, S. P., Van Veen, B., & Rau, M. A. (2020). How drawing prompts can increase cognitive engagement in an active learning engineering course. *Journal of Engineering Education, 109*(3), 723-742.

Zimmerman, B. J. & Campillo, M. (2003). Motivating self-regulated problem solvers. In J. E., Davidson & R. Sternberg (Eds.), *The Nature of Problem Solving*. New York: Cambridge University Press.

Acknowledgments

This research was supported by NSF Award #1760922 and a Presidential Fellowship by the School of Computer Science at Carnegie Mellon University. We thank Matha Alibali, Anna Bartel, Max Benson, Dreami Chambers, Geoff Kaufman, Timothy Nokes-Malach, Amy Ogan, Octav Popescu, Jonathan Sewall, Elena Silla, Nicholas Vest, Kexin Yang, and all the participating students and teachers for their help with the research.