

Exploring the Assistance Dilemma: The Case of Context Personalization

Candace Ann Walkington (cwalkington@wisc.edu)

Wisconsin Center for Educational Research, 1025 W. Johnson Street
Madison, WI 53706 USA

Keith Maull (Keith.Maull@colorado.edu)

Institute of Cognitive Science, Department of Computer Science, 1777 Exposition Drive
Boulder, CO 80301 USA

Abstract

Context personalization, the practice of matching features of an instructional component to a learner's interests and experiences, has been framed in the literature as a means by which to improve learning by enhancing motivation. However, a related perspective could consider personalization a form of instructional assistance, with the potential to support the learning of new concepts. In this paper, the *assistance dilemma*, known to be "a fundamental unsolved problem in cognitive science" (Koedinger, Pavlik, McLaren, & Alevan, 2008, p. 2159), is investigated for context personalization. Two research studies explore whether personalization can be considered a form of assistance, and how this intervention mediates performance measures.

Keywords: personalization; assistance dilemma; cognitive tutor; algebra

Background

The *personalization hypothesis* (learnlabs.org/research/wiki) posits that matching up features of an instructional program to a learner's interests, experiences, or typical language usage will lead to more learning, compared to when instruction is not personalized. The benefits of personalization are often framed in terms of motivation or interest (e.g., Heilman, Collins, Eskenazi, Juffs, & Wilson, 2010); however personalization could also be considered a form of instructional assistance. This may be a useful perspective, given that research results for personalization are mixed, with some studies showing positive effects (Heilman et al., 2010) and others showing no effect (McLaren, Lim, Gagnon, Yaron, & Koedinger, 2006). Framing personalization as assistance makes explicit that benefits should only be expected in some cases, namely when the assistance is both needed and substantive. In this paper, the idea of personalization is explored from the perspective of an important issue in the science of instruction, the assistance dilemma.

The Assistance Dilemma

The *assistance dilemma*, or how to balance information-giving with information-withholding in learning environments, has been framed as "a fundamental unsolved problem in cognitive science" (Koedinger, Pavlik, McLaren, & Alevan, 2008, p. 2159). Assistance is considered not only to be direct hints or scaffolds, but any modification to the learning environment that enhances performance or reduces

mental effort. Cognitive load theory differentiates between *extraneous* cognitive load, stemming from activities not related to schema acquisition, *intrinsic* cognitive load, or inherent difficulty from interactivity of knowledge elements, and *germane* cognitive load, or effort related to schema acquisition (Sweller, Merriënboer, & Paas, 1998). Forms of assistance that reduce extraneous cognitive load should enhance learning by freeing up cognitive resources, if the schemas being learned are sufficiently challenging. However, also central to the assistance dilemma is the notion of *desirable difficulties* – research has shown that modifications that reduce performance during instruction (like decreasing feedback) can actually increase learning (Schmidt & Bjork, 1992), suggesting that mental effort can be germane to learning. The assistance dilemma considers when to give assistance and how much assistance to give in order to facilitate learning and learning efficiency, while also acknowledging that assistance can serve as a "crutch" or a "scaffold" (Koedinger et al., 2008).

Context Personalization

Here the primary topic of concern is one specific type of personalization - *context personalization*. In context personalization interventions, features of an instructional program are matched to individual learner's personal interests and experiences. For example, in mathematics, rather than being given a generic story problem on harvesting wheat from a field of grain, a learner might receive a variation of this problem based on their individual interests, perhaps a mathematical scenario about playing a video game or shopping at the mall. The idea that such personalization of mathematics problems may enhance learning is prevalent in the culture of schooling (Fives & Manning, 2005); however little research has empirically examined its impact.

The assistance dilemma has sometimes been framed with respect to the "education wars," or the struggle to strike a balance between giving more assistance (i.e. direct or traditional instruction) and less assistance (i.e. problem-solving and discovery learning) in learning environments. Context personalization is an especially fascinating instance of the assistance dilemma, because although it could be seen as a form of information-giving, it is widely supported by reform movements (e.g. National Council of Teachers of Mathematics, 2000).

Previous Studies Perhaps the most well-known study of context personalization in mathematics is Cordova and Lepper (1996). Elementary school students were given computer activities on order of operations, presented in different instructional formats. Results showed that students in the condition where the activity was individually personalized to interests (as previously assessed by questionnaires) had significantly higher learning gains than those in a generic condition. Anand and Ross (1987) conducted a similar study where elementary school children learned about division of fractions in a computer-assisted environment. They found an overall positive performance effect for students who received context personalization, and a significant interaction indicating that personalization was most beneficial for low- and middle-achieving students. In a third study where 7th grade students were given story problems involving division of whole numbers, Lopez and Sullivan (1992) found that both individual and group-level context personalization enhanced post-test performance, but only for more difficult problems involving two operations. However, other studies of elementary mathematics have found no effect for context personalization (e.g. Bates & Weist, 2004).

In the domain of mathematics, the few studies that have been conducted on personalization have been with elementary grade students solving simple arithmetic problems, and outcomes for personalization are mixed. Whether context personalization can support performance in higher-level mathematics and for more mature learners is unknown. The body of research on context personalization is so small, that from the literature it is not even apparent whether personalization could generally be considered a form of assistance. Thus the primary goal of the studies reported here was first to establish if context personalization is a form of assistance among a more advanced group of learners, and then to begin to explore how personalization may mediate different performance measures.

Study 1

The first study examined the impact of personalization on student performance during face-to-face problem-solving sessions. An abbreviated version of the analyses for Study 1 is given to lead into the larger experimental study, Study 2.

Method

Twenty-four high school Algebra I students were given 3 four-part story problems on linear functions to solve using pencil-and-paper. Two of the problems had been personalized to out-of-school interests and experiences the student had discussed during a pre-interview, while one problem was a normal story problem from the Cognitive Tutor Algebra (carnegielearning.com) curriculum. The first two parts of each problem were *result unknowns* (Koedinger & Nathan, 2004), where the student was given a specific x-value to plug into a linear process like “ $y=2x+11$,” and had to solve for y. In the third part of each problem, the student had to write an algebra rule, and the final part was a *start*

unknown, where the student must solve for x in a linear process given a specific value of y. All problems had one of 13 linear functions (i.e. $y=2x+11$) as their underlying structure. See Table 1 for example problems.

Table 1: Example of normal and personalized problems

Problem Type	Example
Normal	Some early Native Americans used clam shells called Wampum as a form of currency. Tagawininto had 80 wampum shells, and spends 6 of them every day. a. How many shells did Tagawininto have after 10 days? b. How many shells did he have after a week? c. Write an algebra rule that represents this situation using symbols. d. After how many days did he have 8 shells?
Personalized	You are playing your favorite war game on the Xbox 360. When you started playing today, there were 80 enemies left in the locust horde. You kill 6 enemies every minute. a. How many enemies are left after 10 minutes? b. How many enemies are left after 7 minutes? c. Write an algebra rule that represents this situation using symbols. d. If there are only 8 enemies left, how long have you been playing today?

Student responses were coded for accuracy by two coders ($\kappa = 0.96$), and performance data was analyzed using a mixed-effects logistic model. The dependent variable was whether the student got a problem part correct or incorrect. Random effects included which student was solving the problem, and which linear function was being solved. Fixed effects included whether the problem was normal or personalized, which problem part was being solved, whether the student was classified as low-performing, medium-performing, or high-performing¹, and whether the linear function was classified as easy, medium, or hard².

Results

Analyses showed that personalization was a significant predictor of performance. Main effects for all independent variables were significant ($p < .05$), along with a number of interaction terms, including the interaction between problem type and student level, problem type and problem level, and

¹ This classification was based on performance during the session, however the mathematics standardized test scores of the students from the low-performing group were significantly lower than students in the other two groups ($t=2.73, p<.05$).

² This classification was based on performance during the session, but classifications were also reviewed based on mathematical difficulty of the functional form.

problem level and student level. The results are summarized in Figure 1, which also gives the p -value for each corresponding regression coefficient.

	Easy Problem	Medium Problem	Hard Problem
Low-Performance Student	No significant effect	No significant effect	Increased performance** (2% to 28%)
Medium-Performance Student		No significant effect	No significant effect
High-Performance Student	Decreased performance* (96% to 73%)	No significant effect	Increased performance* (59% to 90%)

Figure 1: Impact of personalization on performance, relative to performance on normal problems, by student and problem levels (* $p < .05$, ** $p < .01$)

Figure 1 shows how personalization increased performance for low-performance students solving hard problems and high-performance students solving hard problems. The figure also shows an *expertise reversal* (Kalyuga, Chandler, & Sweller, 2003) effect for personalization when easy problems were being solved by high-performance students. The figure gives the success rates as estimated by the logistic model for both personalized and normal problems.

These results suggest that context personalization is indeed a form of assistance, and that it can allow students some level of immediate success and opportunity for learning when they are struggling to solve difficult problems. One striking result is that there were 6 students (5 of whom were low-performance) in the sample that got no parts of their normal problem correct, but were able to have varying levels of success (25-100%) on their personalized problems. The reverse was never true. This suggests that personalization may have the most potential as a form of assistance when students are near the edge of their capabilities, and like other forms of assistance, should be faded out as expertise develops.

Study 2

Ultimately a larger sample size was needed to better understand how personalization impacts performance and problem-solving behaviors. Thus preliminary results from Study 2, a large randomized-control experiment, are reported next. In this study, there is also discussion of *learning efficiency* (Koedinger, Corbett, & Perfetti, 2010), or the idea that because instructional time is so valuable, completing instructional activities in less time (without reducing learning) can be considered an important outcome of an intervention. Koedinger et al. (2010) observe that “too many theoretical analyses and experimental studies do not address the time costs of instructional methods” (p. 34).

Method

Participants and Materials One hundred and forty-five high school Algebra I students participated in Study 2. The students’ school used the Cognitive Tutor Algebra curriculum. Cognitive Tutor is an interactive, software-based intelligent tutoring system that presents multi-part algebra problems to students and offers customized problem-selection as well as hints and feedback. The program individualizes problem selection by using *knowledge-tracing* approaches to determine mastery of the concepts being learned. The program also uses *model-tracing* to relate the learner’s problem-solving actions to a cognitive model in order to diagnose errors and offer feedback. For a more in-depth discussion of Cognitive Tutor, see Koedinger and Aleven (2007).

Procedure and Setup Students were randomly assigned to experimental and control groups. When students in both groups entered Unit 6 of the software (Linear Models and Independent Variables), the computer administered an interests survey where they rated their interest in 9 topics (sports, music, movies, TV, games, food, stores, art, and computers). Students in the control group then received the standard algebra story problems in Unit 6, while students in the experimental group received problems selected by the computer to be personalized to their interests.

Each of the 27 problems in the unit had 4 variations corresponding to different interest categories. These variations were written based on a prior survey of student interests ($N = 50$). The variations were similar to the manipulation shown in Table 1, although in this study the changes to the story were sometimes considerably more simple. For each problem, students were asked to fill in different cells of a table as they solved result unknowns, start unknowns, and wrote algebra rules. Figure 2 shows a screenshot of the questions posed and the answer key for the scenario, “You are jogging on the school track to train for the sports team you are on. You are jogging at a rate of 2.9 meters per second, and have already gone 100 meters.”

- 1 How many meters total will you have jogged in 30 more seconds?
- 2 How many meters total will you have jogged in four more minutes?
- 3 In how many more seconds will you have jogged 5000 meters total?
- 4 In how many more seconds will you have jogged 1,927 meters total?

Quantity Name	the time	distance jogged
Unit	seconds	meter
Expression	X	$2.9X + 100$
Question 1	30	187
Question 2	240	796
Question 3	1689.6552	5000
Question 4	630.0	1927

Figure 2: Screenshot of questions and solution values for a story problem in Unit 6 of Cognitive Tutor

Methods of Analysis The unit of analysis was one student solving one part of one problem, which corresponds to

filling in one cell in Figure 2. The data included the 73,953 problem parts being solved by the 145 students as they were presented with normal or personalized versions of the 27 base story problems in Unit 6. Performance was modeled with a mixed-effects logistic regression. Given the manner in which Cognitive Tutor’s artificial intelligence selects problems and assesses learning based on knowledge-tracing algorithms, students generally received both different problems and different numbers of problems. Thus in the statistical analysis, base problem and student are both random effects. Fixed effects included what condition the student was in (experimental or control), as well as what *knowledge component* (Koedinger & Alevan, 2007) or skill was being addressed in the cell they were filling in.

Knowledge components were placed in three groups based on overall difficulty within Unit 6. The first group was easy knowledge components; these included filling in names for quantities and units (first two rows in Figure 2) and entering a given value. The second group was medium-difficulty knowledge components, such as solving result and start unknowns and working with different types of numbers (small, large, decimal). The third group was the most difficult knowledge components assessed in the tutor, writing algebraic expressions of various forms. Along with an analysis of performance, measures of learning efficiency, hint-seeking, and gaming the system were also analyzed.

Results

Performance Main effects for both condition and difficulty of knowledge component were significant predictors of performance, as was the interaction between condition and knowledge component difficulty (Table 2). The raw coefficients are in logit form, and in the third column they are transformed to odds. “Condition-E” is the experimental group (received personalized problems), and “KC” stands for knowledge component. The reference groups are the control condition and easy knowledge components.

Table 2: Regression coefficients for mixed effects logistic model predicting problem performance (fixed effects only)

	Raw Coeff	Std Err	Exp (Coeff)	z-value
(Intercept)	1.95	0.111	7.03	17.47 ***
Condition-E	0.35	0.091	1.42	3.48 ***
KC-Easy	Ref.			
KC-Medium	-0.85	0.029	0.43	-29.59 ***
KC-Hard	-2.07	0.055	0.13	-37.66 ***
Condition-E× KC-Medium	-0.19	0.041	0.83	-4.68 ***
Condition-E× KC-Hard	0.03	0.077	1.03	0.38

‘***’ $p < .001$

Being in the experimental condition increased predicted performance on easy knowledge components (odds = 1.42, $p < .001$) and hard knowledge components (odds = 1.42 ×

1.03 = 1.46, $p < .001$). The effect was only marginally significant for medium knowledge components (odds = 1.42 × 0.83 = 1.17, $p = .089$). The increased performance on easy knowledge components, like entering in labels for quantities and units, suggests that personalization facilitates learners’ ability to get through less mathematically-relevant portions of the problem in a timely manner. The increased performance on writing symbolic expressions (hard KCs) shows that context personalization could be considered a form of assistance for this challenging algebraic task.

Placing the raw coefficients in Table 2 into the logistic model, it can be seen that for hard knowledge components, the control group had a predicted performance of 46% correct, while the experimental group had a predicted performance of 56% correct. Thus the model estimates that personalization improves performance on hard knowledge components by 10%, which is a considerable increase for the domain of algebraic expression-writing.

Results suggest that personalization provides assistance to students as they learn the difficult skill of algebraic expression-writing. Personalization may also increase learning efficiency by facilitating performance on easy knowledge components. There was no evidence of an expertise reversal effect in this study – as these problems were based on general survey data rather than individual interviews, they may have been less distracting or seductive.

Problem Step Duration Learning efficiency was examined by looking at the time a learner spent completing each problem part (i.e. filling in one cell in Figure 2), as it varied by condition. In terms of average step duration, the pattern of difficulty for knowledge components was slightly different than it was for performance. The knowledge components with the lengthiest step durations included writing expressions with both slope and intercept terms, and solving a start unknown when the expression had a positive slope. For the purposes of the duration analysis, these are considered “hard” knowledge components. The knowledge components with medium average step durations included solving all other start unknowns, working with different types of numbers (small, large, decimal), and writing expressions with only slope terms. The easy knowledge components were the same as they previously were.

A linear mixed-effects regression was conducted with base problem and student as random effects and condition and knowledge component difficulty as fixed effects. The dependent measure was the number of seconds the learner spent on the problem step. Results are shown in Table 3.

The main effect for condition was not significant; however, the interaction between condition and knowledge component difficulty was significant. As can be seen from the table, personalization significantly decreased step duration when the problem part being solved was associated with a hard (time consuming) knowledge component. The estimated size of the effect is a reduction of 8.6 seconds ($3.06 + 5.54 = 8.6$, $p < .01$). This suggests that personalization assists learners solving hard problems, and has the potential to increase learning efficiency. Overall,

students in the experimental group had 1.88 correct answers per minute in the unit, while control group students had 1.56 correct answers per minute. Looking only at the hard knowledge components, the efficiency scores for the experimental and control group were 0.59 correct per minute and 0.42 correct per minute, respectively.

Table 3: Regression coefficients for linear mixed-effects model predicting problem step duration (in seconds)

	Coeff	Std Err	t-value
(Intercept)	17.07	2.53	6.76 ***
Condition-E	-3.06	2.54	-1.21 ***
KC-Easy	Ref.		
KC-Medium	19.62	0.75	26.16 ***
KC-Hard	43.47	1.32	32.87 ***
Condition-E× KC-Medium	-1.52	1.04	-1.46
Condition-E× KC-Hard	-5.54	1.80	-3.07 **

*** $p < .001$ ** $p < .01$

Reading Time The time students spent reading each story scenario was measured by calculating the elapsed time between when the problem first came up on the student’s screen, and their first interaction with the tutor. A linear mixed-effects regression analysis was conducted with student and item as random effects, and condition as a fixed effect. The dependent variable was reading time in seconds. Results showed that condition was a significant predictor of reading time ($z = -2.06, p < .05$) with personalization of story scenarios estimated to reduce average reading time by 7.6 seconds. The predicted reading time for students in the control group was 33.1 seconds, compared to 25.5 seconds for the experimental group. This again suggests that personalization improves learning efficiency by reducing time spent on less mathematically-relevant problem parts.

Hint-Seeking The hint-seeking behavior of students was analyzed by looking for differences in the average number of hints given to the student by the tutor per problem part. A linear mixed-effects model with student and item as random effects found that condition was a significant predictor of hints requested ($z = -2.33, p < .05$). Personalization reduced the number of hints per problem part by 0.12 hints, from an estimated 0.38 hints per problem part for the control condition to 0.26 hints per problem part for the experimental condition. Personalization seemed to act as a form of assistance that allowed students to use substantially fewer of the built-in hints in the Cognitive Tutor software.

While it is critical that students try to persist in solving problems and learn from making mistakes, it is also important that students have the metacognitive awareness to know when to seek hints, rather than continue to flounder. Using only the transactions where students received hints, an exploratory analysis was conducted to see which problem parts students received the most hints on.

Table 4 shows that the experimental group received fewer of their total hints on easy knowledge components, and more of their total hints on hard knowledge components. The increased hint-seeking for hard knowledge components does not explain the performance differences presented earlier; the dependent measure in those analyses was whether the student got the correct answer on the first attempt, without asking for a hint.

Table 4: Percentage of total hints requested by experimental and control groups for different knowledge components

KC Type	% of total hints received	
	Control	Experimental
Easy	10.4%	7.0%
Medium	54.7%	54.6%
Hard	34.9%	38.4%

The analysis of hint-seeking shows that students in the experimental group received fewer hints, and when they did receive hints, they were on more difficult knowledge components. This suggests that personalization facilitates efficiency on easier knowledge components, but students still seek help on more challenging skills.

Gaming the System Research on intelligent tutoring systems has shown that students sometimes engage in behaviors referred to as *gaming the system*, (Baker, Corbett, Koedinger, & Wagner, 2004) where they take advantage of the tutor’s help and feedback. For example, students may enter in answers quickly and repeatedly, trying to guess the answer that the tutor will accept, or students might click rapidly through the tutor’s hints in order to get to the “bottom out” hint, where the tutor essentially gives the student the answer to the problem part. Gaming the system has been shown to be negatively correlated with learning (Baker et al., 2004), and has been framed with respect to the assistance dilemma as a reaction to information-withholding in instructional environments (Koedinger & Aleven, 2007). Accordingly, research has shown that weaker students are more likely to game the system (Baker et al., 2004).

The Cognitive Tutor Gaming Detector (Baker & de Carvalho, 2008) was run on students’ transactions with the tutor while in Unit 6. The gaming detector collects a variety of quantitative data from students’ transactions with the tutor, including time and hint-seeking measures, and determines how often the student is likely to have been gaming the system. Results showed that students in the experimental group gamed the system significantly less often ($t = -2.33, p < .05$). This suggests that personalization is acting as a form of assistance, reducing gaming and increasing use of learning-focused strategies. Personalization might be especially effective for weaker students, given that they most often game the system.

Conclusion

The results presented here suggest that context personalization can be considered a form of assistance, increasing both performance and learning efficiency. Framing personalization with respect to interest or intrinsic motivation is an important perspective, and is central to our own work; however this explanation alone cannot account for empirical results. Specifically, rather than context personalization boosting performance in a single, generic manner across all problems, it seems to be most effective when students are struggling with a difficult task. For adept problem-solvers, it may introduce extraneous cognitive load, causing accuracy to decrease through an expertise reversal effect. The conceptualization of personalization as assistance also helps to explain why the results in the literature are so mixed. Gains should only be expected when students are solving challenging problems where the type of assistance offered by personalization is helpful.

This then leads to the question of *why* context personalization actually provides assistance. In the domain of mathematics story problems, Nathan, Kintsch, and Young's (1992) conception of a *situation model* seems to be a highly probable explanation. Context personalization may allow students to have a better implicit grasp of the actions and relationships in a story scenario, allowing them to write algebraic expressions and solve result and start unknowns more accurately, as they reason with familiar quantities.

Ultimately, the assistance dilemma is about learning, and in the present paper, we only look at measures of performance and learning efficiency. As it seems plausible that context personalization can productively be considered a form of assistance, in future work we seek to explore further whether this assistance is a "crutch" or a "scaffold."

Acknowledgments

This work was supported by the Pittsburgh Science of Learning Center which is funded by the National Science Foundation award # SBE-0354420, and Carnegie Learning. Special thanks to Ryan Baker, Sujith Gowda, and Ben Shih.

References

- Anad, P., & Ross, S. (1987). Using computer-assisted instruction to personalize arithmetic materials for elementary school children. *Journal of Educational Psychology*, 79(1), 72-78.
- Baker, R., Corbett, A., Koedinger, K., & Wagner, A. (2004). Off-task behavior in the Cognitive Tutor classroom: When students "game the system." Proceedings of ACM CHI 2004: Computer-Human Interaction.
- Baker, R., & de Carvalho, A. (2008). Labeling student behavior faster and more precisely with text replays. In the *Proceedings of the 1st International Conference on Educational Data Mining*.
- Bates, E., & Wiest, L. (2004). The impact of personalization of mathematical word problems on student performance. *The Mathematics Educator*, 14(2), 17-26.
- Cordova, D., & Lepper, M. (1996). Intrinsic motivation and the process of learning: Beneficial effects of contextualization, personalization, and choice. *Journal of Educational Psychology*, 88(4), 715-730.
- Fives, H., & Manning, D. (2005). Teachers' strategies for student engagement: Comparing research to demonstrated knowledge. Paper presented at 2005 Annual Meeting of American Psychological Association, Washington DC.
- Heilman, M., Collins, K., Eskenazi, M., Juffs, A., & Wilson, L. (2010). Personalization of reading passages improves vocabulary acquisition. *International Journal of Artificial Intelligence in Education*, 20(1), 73-98.
- Kalyuga, S., Ayres, P., Chandler, P., & Sweller, J. (2003). The expertise reversal effect. *Educational Psychologist*, 38(1), 23-31.
- Koedinger, K. R., & Alevan, V. (2007). Exploring the assistance dilemma in experiments with Cognitive Tutors. *Educational Psychology Review*, 19, 239-264.
- Koedinger, K. R., & Nathan, M. (2004). The real story behind story problems: Effects of representations on quantitative reasoning. *Journal of the Learning Sciences*, 13(2), 129-164.
- Koedinger, K. R., Pavlik Jr., P. I., McLaren, B. M., & Alevan, V. (2008). Is it better to give than to receive? The assistance dilemma as a fundamental unsolved problem in the cognitive science of learning and instruction. In B.C. Love, K. McRae, & V. M. Sloutsky (Eds.), *Proceedings of the 30th Annual Conference of the Cognitive Science Society*. Austin, TX: Cognitive Science Society.
- Koedinger, K. R., Corbett, A. T., & Perfetti, C. (2010). The Knowledge-Learning-Instruction (KLI) framework: Toward bridging the science-practice chasm to enhance robust student learning. CMU-HCII Tech Report 10-102. <http://reports-archive.adm.cs.cmu.edu/hcii.html>
- Lopez, C., & Sullivan, H. (1992). Effects of personalization of instructional context on the achievement and attitudes of Hispanic students. *Educational Technology Research and Development*, 40(4), 5-13.
- McLaren, B. M., Lim, S., Gagnon, F., Yaron, D., & Koedinger, K. R. (2006). Studying the effects of personalized language and worked examples in the context of a web-based intelligent tutor. In the *Proceedings of the 8th International Conference on Intelligent Tutoring Systems*.
- Nathan, M., Kintsch, W., & Young, E. (1992). A theory of algebra-word-problem comprehension and its implications for the design of learning environments. *Cognition and Instruction*, 9(4), 329-389.
- National Council of Teachers of Mathematics. (2000). *Principles and standards for mathematics*. Reston, VA.
- Schmidt, R., & Bjork, R. (1992). New conceptualizations of practice: Common principles in three paradigms suggest new concepts for training. *Psychological Science*, 3(4), 207-217.
- Sweller, J., Merrienboer, J., & Paas, F. (1998). Cognitive architecture and instructional design. *Educational Psychology Review*, 10(3), 251-296.