Learning from the folly of others: Learning to self-correct by monitoring the reasoning of projective pedagogical agents

Sandra Y. Okita, Teachers College, 525 West 120th Street, New York, NY 10027, okita@tc.columbia.edu
Azadeh Jamalian, Teachers College, 525 West 120th Street, New York, NY 10027, aj2334@columbia.edu

Abstract: Students find it relatively natural to catch other people's errors, even if they are not inclined to catch their own. A total of sixty-two nine- to eleven-year-old students participated in two studies that tested the hypothesis that monitoring the reasoning of a pedagogical agent solving math problems, can help students learn the skill of monitoring, and eventually self-correct when solving math problems. A projective pedagogical agent "ProJo" was designed to openly displays its reasoning when solving math problems allowing children to "look for mistakes". Two testing environments "Doodle Math", and "Puzzle Math" were implemented to directly compare self-monitoring and self-other monitoring treatment. The results showed initial evidence that self-other monitoring may be an effective way to help students develop metacognitive skills to self-correct and accurately solve problems.

Introduction

Careless mistakes can easily be avoided if children learn to check their answers. One way to minimize mistakes is for people to monitor their thinking and activity. This way, they can catch potential mistakes or confusions before they lead to problematic errors. There are various viewpoints on how the ability to self-monitor develops in childhood and how to promote it (Markman, 1977). One viewpoint is that children develop abilities to self-monitor as they work to resolve their own confusions and internal conflicts (Karmiloff-Smith, 1974; Karmiloff-Smith 1979; Brown 1987). From this perspective, one way to help children learn to self-monitor is to put them in situations that will lead to cognitive conflict. Often times people turn to others to understand themselves. In this viewpoint children learn to self-monitor by observing other people who self-monitor (Palincsar & Brown, 1984; Okita & Schwartz, 2006). In this case, presenting children models of others who overtly self-monitor should lead to improvements. This research examines a third alternative: Children learn to self-monitor by monitoring other people, and with practice, they turn the external monitoring inwards. Children may find it difficult to be attentive to their own mistakes when they are concentrating on a problem or task. However, they may find it relatively easy to catch other people's mistakes. Gelman and Meck (1983) found that children, who had difficulty checking their own work, found it easier to point out mistakes of others (e.g., puppets). By focusing children on catching another's mistakes, children may learn to the skill of monitoring, and eventually self-monitor.

Related Works

Studies have shown that learning among peers and comparing ourselves to peers can be quite effective (Chi, 2001; Graesser, Person, & Magliano, 1995). For example, comparing test scores helps people learn where they stand academically. Another way is observing others to better understand the self. For example observation can be better than doing when students cannot solve a math problem, but observing others solve a math problem may help you learn how to solve problems. This is because the person they are observing can provide a model of competent performance. Research has also shown that observing somebody who knows less (or knows about the same as you) can be beneficial. Other studies in the area have shown (e.g., reciprocal teaching) that students may spontaneously compare their understanding to what they observe in another person, and any discrepancies can alert them to think more deeply about who is right (Seifert & Hutchins, 1992). This implies that observing an imperfect peer, under the right circumstances, might be a great trigger for learning and reflection. It is expected that the effect of self-monitoring may not be immediate while the task is still challenging and unfamiliar, and may take some time for the behavior to turn inward. Norman (1983) suggests that at first, students who self-other monitor may do worse than students who do not, but over time they may surpass those students without the self-other monitoring treatment. An alterative approach to a human peer is the use of a pedagogical agent. Studies also demonstrates that students can learn by observing, an agent that knows less than they do (Schwartz, Chase, Chin, Oppezzo, Kwong, Okita, Roscoe, Jeong, Wagster, & Biswas, 2009).

An ideal learning environment would be an interactive pedagogical agent designed to solve math problems that would engage the student in catching mistakes and monitoring agent behavior. Students with practice can turn this external monitoring inwards and learn to self-correct. The Projective Pedagogical Agent, "ProJo," was developed so that students could learn to monitor and check for potential mistakes. ProJo displays the reasoning behind the problem solving so that the student can follow ProJo's "lead" and check the answer. The paper introduces two studies that test the hypothesis of whether modeling a pedagogical computer agent solving math problems may help students learn to self-monitor and solve math problems with better accuracy.

Two testing environments "Doodle Math" (study1) and "Puzzle Math" (study 2) were created to allow a direct comparison between self-other monitoring to self-monitoring.

Figure 1. ProJo agent in Doodle Math (Left), Puzzle Math & Experimental Setting (Right)

Study 1: Self vs. Self-Other Monitoring

Study Design and Procedure

In the first study, forty elementary school students between the ages 9- to 11- years participated in a one-to-one forty-five minute session. The students were randomly assigned to the self-monitoring treatment or self-other monitoring treatment. Self-other treatment was compared directly to self-treatment as a between subjects design. Students were taught a math trick on divisibility rules for multiplication. In the self-other treatment, ProJo appears in the learning environment (See Figure 1), whereas in self-treatment, ProJo is not present. The treatment session has ten trials (5 pair trials). Children in the self-treatment solve all ten trials on their own. In the self-other treatment, the student and ProJo, each take turns solving five trials each. In other words the student in the self-other treatment condition will be solving 5 trials on their own, and 5 trials monitoring the agent ProJo solve problems. The students were responsible for monitoring and catching any mistakes ProJo made. ProJo solved problems that were sometimes correct and other times incorrect. When it was ProJo's turn, he would say, "I think I learned a lot by watching you play. Let me try. I don't like mistakes, so if you can catch me when I make a mistake that would be great. We are a team, and I wouldn't want you to lose a point because of my mistake." ProJo would make a few intentional "slow motion" mistakes so the student could monitor and catch (or not catch) the error. The relative effects of self-other monitoring were evaluated using calculation time and accuracy measures during the treatment sessions, and pre- post-tests were administered at the beginning and end of the study. To see the effect of each treatment (self vs. self-other) the analysis looked only at the 1st of the pair for the self and self-other treatments. Therefore the first analysis included 1st of the pair (five problems) from the self-other treatment, and the 1st of the pair (five problems) from the self-treatment condition.

Materials and Measures

The testing environment "Doodle Math" (See Figure 1 Left) was created to allow students to practice monitoring. Students used a pen and tablet device to write out their calculations on the screen. Doodle Math was developed to run in a web-based environment. The key measures from this study include a pretest and posttest, math problems used in the 10 trials during treatment session, screen shots during treatment, and a play back log file of student's behavior. A pretest was used to examine the current math level of the students and to identify the calculation errors students made. The errors were later implemented as ProJo's calculation mistakes, and displayed during the self-other monitoring treatment sessions. The relative effects of self-other monitoring were evaluated using time and accuracy measures during the treatment sessions, as well as on pre- and post-tests. The posttest showed if self-other treatment had an overall effect on math accuracy. The screen shots showed student's reasoning and method when problem solving (e.g., did they follow the math trick, where did they make a mistake, did they accurately catch ProJo's mistakes). The log file recorded the hand written calculation for each student by keeping a log of the data points (e.g., x, y points of the white drawing area) from the pen and tablet. The log file allowed ProJo to play back the reasoning process and calculation of the student in the student's original handwriting and timing. This playback log file was an important assessment tool for teachers because they could playback the recording and see where the student hesitated during problem solving.

Results

The initial thought was that students in the self-other treatment would "slow down" when solving problems on their own. During the treatment session, students in the self-other treatment (Figure 3 dotted line) showed a more gradual decrease in time than the self-treatment (See Figure 2 dotted line). The students in the self-treatment increased calculation speed immediately after the first trial. Students in both treatments spent most time on the first problem, possibly getting acquainted. The effect of accuracy and time was short lived where the two treatments showed little difference by the posttest. Comparing calculation time and accuracy across the trials, results showed that students in the self-treatment immediately picked up speed and did worse over time.

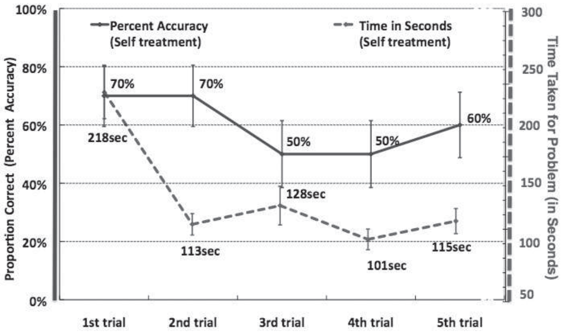

Figure 2. Accuracy and time for self-treatment.

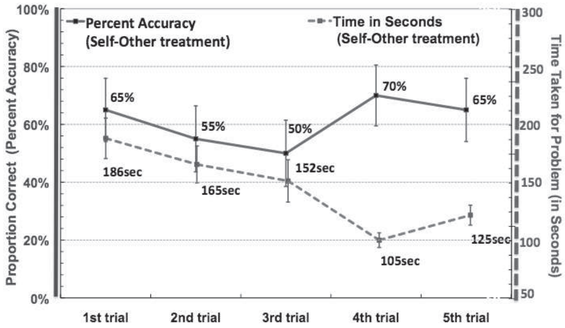

Figure 3. Accuracy and time for self-other treatment.

Students in the self-other monitoring treatment slowed down after the first exposure to the agent, then gradually picked up speed and improved in accuracy over time (Figure 3). One possible interpretation is that students are working hard to internalize the monitoring, which cashed out on later problems. The results were not sufficiently strong to guarantee this interpretation, but showed a promising trend.

As mentioned earlier, the pretest identified some calculation errors students made. The student's errors were later implemented as ProJo's calculation mistakes, and displayed during the self-other monitoring treatment. During monitoring, the agent ProJo would randomly make three mistakes and correctly solve two problems, so the student did not automatically expect ProJo to always make an error. The three mistakes included two similar calculation mistakes that the student made on the pretest (mix-up procedure and losing track in procedure), and one general calculation mistake (different from the student's mistake). The results showed that similar mistakes (as their own) were difficult to catch, while students had no trouble catching general calculation mistakes, or acknowledging when the ProJo solved the problems correctly. One possible explanation is that general calculation errors are "matters of fact" types of problems, and students can easily remember their math facts and compare them to ProJo's computation. In contrast, for the mix-up and lose track problems, the errors are procedural, and students need to follow the procedure and check it against their own answer. This may be more difficult because they are new to the procedures, whereas math facts are familiar. An alternative explanation is that the calculation errors are different from the mistakes the students would make, so it is easy for them to catch. The critical limitation was that ProJo was successful at replaying errors, but unsuccessful in making the students notice their own mistakes.

The study provided initial evidence that self-other treatment using ProJo may be an effective way to help students learn to self-correct. The studies were short in duration and designed to see if there were any short-term effects. A longer intervention may determine if more sustained practice at monitoring would lead to stronger effects. Another limitation is that that current study falls short in determining whether the students were monitoring or just copying behavior. The second study extended the session to an hour and continued treatment for two days to see if this would sustain monitoring practice and improve accuracy. A new testing environment "Puzzle Math" (See Figure 1 Right) was created for the second study so that it included a detection features to measure student monitoring.

Study 2: Extended Treatment and Detecting Monitoring and Behavior

Study Design and Procedure

Twenty-two elementary school students from Northern Manhattan participated in a one-hour treatment session. The students were between the ages 9- to 11-years and participated for one session a day for two days. The students were randomly assigned to one of two conditions, self-monitoring treatment or self-other monitoring treatment. Self-other treatment was compared directly to self-treatment as a between subjects design. The treatment session had seven puzzles. Children in the self-treatment solved all seven puzzles on their own. In the self-other treatment, the student took turns with ProJo, the student solving five puzzles, and monitored ProJo solving two puzzles. The relative effects of self-other monitoring were evaluated using accuracy measures, a monitor log that detects if students were monitoring (i.e., self correcting) their own work, and a pre- and post-tests administered at the beginning and end of the study.

Materials and Measures

The second test environment "Puzzle Math" incorporated many of the lessons learned from the first study. Puzzle Math placed emphasis on, 1) extended treatment over two days, 2) a puzzle that helps identify student's procedural calculation mistakes, and 3) the addition of a log monitor that detection when students are self correcting. In the first study, students made a lot of procedural mistakes (i.e., mixing up numbers, losing track and calculating the same number twice), which were not calculation mistakes, but procedural mistakes. In the first study, we found that the calculation steps could differ by the student's math level. A student who may not be able to use the math efficiently, may look at 25 x 240 and write out 5 x 5 x 2 x 2 x 2 x 2 x 3 x 5 and start calculating each number one by one increasing their chances of losing track and mixing up numbers and getting the answer wrong. A student, who is more efficient, may look at 25 x 240; see that 240 can be broken down to 4 x 60 to make the calculation easier because 25 x 4=100, and 100 can be easily calculated by 60 = 6000. To address this issue, Puzzle World has a set of six to eight problems where the divisible numbers are small (1~2 digits), requiring fewer factoring (e.g., 4 = 2 x 2, 12 =3 x 4), and fewer chances of losing track. More details on the problem sequence are described in the next section.

The pretest determined the student's math level, and any prior knowledge to similar math puzzles. The new math game involved divisibility rules for multiplication (Figure 1 Right). Taking the form of a puzzle, the math game provided opportunities to monitor the agent's reasoning when problem solving. Each puzzle was created from a small set of six to eight math problems on divisibility rules. As the game progresses, restrictions were imposed on the puzzle (e.g. Type in any number between 1-9 where the product equals the bold numbers). The restrictions got more challenging as the game progressed, where the easiest rule was *Calculation Only*: use any number between one to nine, medium level rule was *Calculation & Rule*: use any number between one to four, and the hardest rule consistency across *Entire Puzzle*: use any number between 1 to 4 whose products equal the blue bold number, and only have one of each number (1,2,3,4) in each row and column.

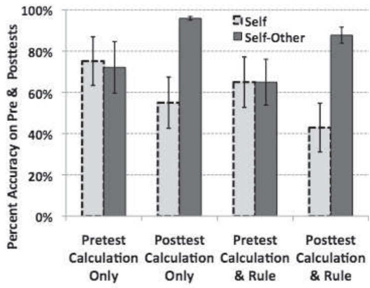

Figure 4. Pre-post-test Accuracy.

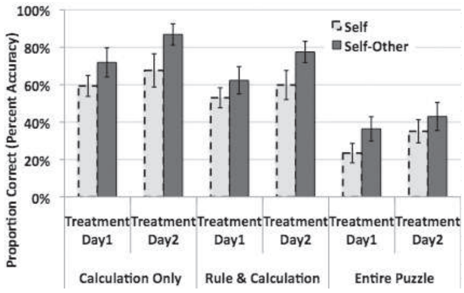

Figure 5. Accuracy over two days.

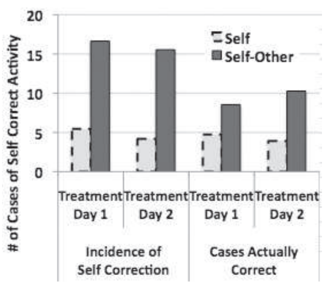

Figure 6. Self correcting activity.

Results

The accuracy data on student performance was categorized into three different scores. The math problems in study 2 involved a puzzle that consisted of a set of problems. The pretest showed that there was no difference in the groups prior to the study (Figure 4). Figure 5 showed the three accuracy scores during treatment session over the two days. Overall, students in the self-other mode had a higher accuracy score compared to the self-mode. The students in the self-other mode seemed to improve at a higher rate from treatment day 1 to day 2. The difference between the two groups on performance (during the treatment session) started to appear by the end of the second day. For accuracy on "calculation" between the two groups was approaching significance at (Self monitoring M=68%, Self-other monitoring M=87%) t(20)= -1.825, p=.083. For accuracy on following "rules & calculation" the difference between the two groups were approaching significance at (Self monitoring M=59%, Self-other monitoring M=78%) t(20)= -1.833, p=.082.

By the time the students took the posttest, there was a significant difference between the two groups, where the Self-other monitoring performed better than self-monitoring. The posttest accuracy score for the calculation-only problems were significantly different between the two treatments (Self monitoring M=55%, Self-other monitoring M=96%) t(20)= -3.28, p<.05. Even with the rule restrictions, there was a significant difference between the two groups where Self-other monitoring scored higher than the self-monitoring at (Self monitoring M=43%, Self-other monitoring M=88%) t(20)= -3.60, p<.05.

The monitor detection log kept of the number of incidents where the student self corrected during problem solving, and changed their answer before finishing the puzzle (Figure 6). In addition to the number of times the student self corrected, of those corrections, we looked at how many were actually correct. The results showed that self-other monitoring had significantly more cases of self correction compared to the self monitoring (Self monitoring M=9.6, Self-other monitoring M=32.6) t(20)= -2.25, p<.05. There was also a significant difference between the two groups on the accuracy of the self correction (Self monitoring M=8.6, Self-other monitoring M=18.8) t(20)= -3.03, p<.05. Although the number of self-corrections slightly decreased from day 1 to 2, the accuracy in self-correction increased from day 1 to 2 for the Self-other monitoring.

The accuracy score on the posttest indicated that students in self-other monitoring performed better compared to the students from self-monitoring (See Figure 4). This was also seen during the treatment session where self-other monitoring had a higher performance rate as the treatment progressed from day 1 to day 2 (See Figure 5). Possibly due to the late-gain effect the difference between the two treatments were not significant until the posttest. As an indicator of monitoring, the number of self-corrections was significantly higher in the self-other monitoring treatment (See Figure 6). Although the number of incidents slightly decreased from day 1 to day 2, the quality of self-correction seemed to increase for self-other monitoring where accuracy increased from day 1 to 2, while in self-monitoring the number of incidents and accuracy decreased.

General Discussion

One challenge that remains is how ProJo is still unsuccessful in making the students notice their own mistakes. However, seeing ProJo's errors still had a modest benefit for student learning. This was seen in the post-test where students solved problems on their own, away from ProJo. If students made an error on their own problem, and then monitored ProJo make a similar error, the students were less likely to make the error on the post-test. If the students had not made an error, but saw ProJo make a mistake on a similar problem, this tended to hurt the student's post-test performance. Students in the self-monitoring treatment who made an error, but did not get to monitor ProJo make a similar error, were likely to make a similar error on the post-test. These tentative results suggest a practical hypothesis that if a student gets a problem right, then the student should solve another problem on their own. If the student gets a problem wrong, have them monitor an agent. These findings have some implication where a more sophisticated version of ProJo could be designed where the system makes real-time decision to have the student monitor an incorrect ProJo, a correct ProJo, or simply continue working on their own. There may be a need to sequence the agent's mistakes to further improve student learning. By identifying unfavorable problem sequencing, ProJo can have important implications on the learner.

Summary

The two studies provided some evidence that self-other monitoring may be an effective way to help students develop metacognitive skills. The present measures showed a late gain effect and some evidence that students were monitoring (self-correcting) rather than copying behavior. The studies also provided two testing environments that isolated and adjusted factors (e.g. kinds of errors, what was monitored by the student, how agent would respond, and the data retrieved), and tracked student progress and achievement in various data forms (screen shot of worked out problems, time and accuracy scores).

References

Brown , A. (1987). Metacognition, executive control, self-regulation and other more mysterious mechanisms. In Weinert, F.E, & Kluwq, R.H. (Eds.), *Metacognition, Motivation and Understanding*, NJ, LEA.

Chi, M.T.H., Silver, S.A., Jeong, H., Yamauchi, T., & Hausmann, R.G. (2001). Learning from human tutoring. *Cognitive Science*, 25, 471-533.

Gelman, R., & Meck, E. (1983) Preschoolers' counting: Principles before skill, *Cognition*, 13, 343-359.

Graesser, A.C., Person, N., & Magliano, J. (1995). Collaborative dialog patterns in naturalistic one-on-one tutoring. *Applied Cognitive Psychologist*, 9, 359-387.

Karmiloff-Smith, A., & Inhelder, B. (1974-1975). If you want to get ahead, get a theory. *Cognition*, 3, 195-212.

Karmiloff-Smith, A. (1979). Micro- and macro- developmental changes in language acquisition and other representational systems. *Cognitive Science,* 3, 91-118.

Markman, E. (1977). Realizing that you don't understand: Elementary school children's awareness of inconsistencies. *Child Development*, 48, 986-992.

Norman, D.A. (1983). Design Rules Based on Analyses of Human Error, *Comm. of the ACM*, 26(4)254-258.

Okita, S. Y., & Schwartz, D. L. (2006). When observation beats doing: Learning by teaching. In S. Barab, K. Hay, & D. Hickey (Eds.), *Proceedings of the 7th International Conference of the Learning Sciences*, Vol. 1 (pp. 509-515). New Jersey: Lawrence Erlbaum Associates. Bloomington, USA.

Palincsar, A. S., & Brown, A. L. (1984). Reciprocal teaching of comprehension fostering and comprehension-monitoring activities. *Cognition and Instruction*, 2(2), 117-175.

Schwartz, D. L., Chase, C., Chin, D. B., Oppezzo, M., Kwong, H. Y., Okita, S. Y., Biswas, G., Roscoe, R., Jeong, H., & Wagster, J. (2009). Interactive metacognition: Monitoring and regulating a teachable agent. In D. J., Hacker, J. Dunlosky, & A. C. Graesser (Eds.). *Handbook of metacognition in education* (pp. 340-358). New York, Routledge.

Seifert, C.M., & Hutchins, E.L. (1992). Error as Opportunity: Learning in a cooperative task, *Human-Computer Interaction*, 7, 409-435.

Acknowledgments

We thank all the student participants in the study, the Sociable Technology and Learning (STL) lab members, Delci Wright, Selen Turkay, Mihwa Kim, and Yanjin Long (now at CMU).