Learning to Distinguish Between Representations of Data: A Cognitive Tutor that Uses Contrasting Cases

Ryan Shaun Baker, Albert T. Corbett & Kenneth R. Koedinger
Human-Computer Interaction Institute, Carnegie Mellon University, 5000 Forbes Ave, Pittsburgh, PA, 15213
Tel: (412) 268-1208, Fax: 412-268-1266
Email: rsbaker@cmu.edu, corbett@cmu.edu, koedinger@cmu.edu

Abstract: Students often fail to learn crucial distinctions between different representations of data. For instance, many students learning about scatterplots consistently create representations which have the surface features of scatterplots but with informational content more appropriate for discrete bar graphs. Schwartz and Bransford (1998) have found that combining feature-based conceptual instruction with contrasting cases is an effective way to help students make conceptual distinctions. We adapt their approach to the domain of data representation and incorporate it into a cognitive tutoring curriculum. We show that this new curriculum improves learning more than a curriculum where the contrasts are not present.

Introduction

How can we design curricula in which students learn to make appropriate distinctions between different ways of representing data? In recent years, educational standards have placed increasing emphasis on learning to analyze and work with data, as early as in middle school (NCTM 2000). A major part of learning to work with data is learning how to represent it in a number of different forms. Some representations are particularly useful to learn, both because they support certain types of inference particularly well (Tufte, 1983; Tabachneck-Schijf, Leonardo, and Simon, 1997) and because they facilitate communication with other individuals who also understand those representations. Participation in activities that involve external representations, wherever and with whomever they occur, requires having a rich understanding of those representations and their affordances.

Unfortunately, students often fail to learn the informational and functional distinctions between different representations, hampering their ability to use them effectively. In this paper, we present an approach for helping students develop an understanding of such distinctions, focusing specifically on student understanding of scatterplots and bar graphs. Our approach, adapted from Schwartz and Bransford's (1998) contrasting cases method, combines conceptual instruction on the informational differences between representations with a scaffold that helps students explicitly link the type of information available to appropriate types of representations. We present two studies: in study 1, we compare a contrasting cases approach to an approach where students are taught scatterplots without reference to their differences from other representations, in a homeschool setting; in study 2, we test whether the contrasting cases approach is also effective in a classroom setting, and investigate whether an intervention to focus students' time on more difficult parts of each problem results in even more successful learning.

Student difficulties with scatterplots

One of the major challenges to developing a rich understanding of representations is learning to distinguish between different representations. Students often fail to learn what situations a newly learned representation is appropriate for, preferring to use the first and simplest representations of data they learn, even in situations where those representations are not useful (Hancock, Kaput, and Goldsmith, 1992). Students also develop schemas of different representations that are focused solely on their surface features rather than their structural or informational properties (McGatha, Cobb, and McClain, in press). When attempting to draw scatterplots, for instance, students tend to draw a representation with the surface features of a scatterplot but which is the informational equivalent of another representation, a discrete bar graph. The graph created has a quantitative variable on one axis and a nominal or categorical variable (1) on the other axis, instead of a quantitative variable on each axis. This behavior manifests itself in two ways, termed the *variable choice error* (Baker, Corbett, and Koedinger, 2001) and the *nominalization* error (Baker, Corbett, and Koedinger 2002; cf.

Lehrer and Schauble, 2001). When a student makes the variable choice error, she or he places a nominal or categorical variable rather than a quantitative variable along one axis. An

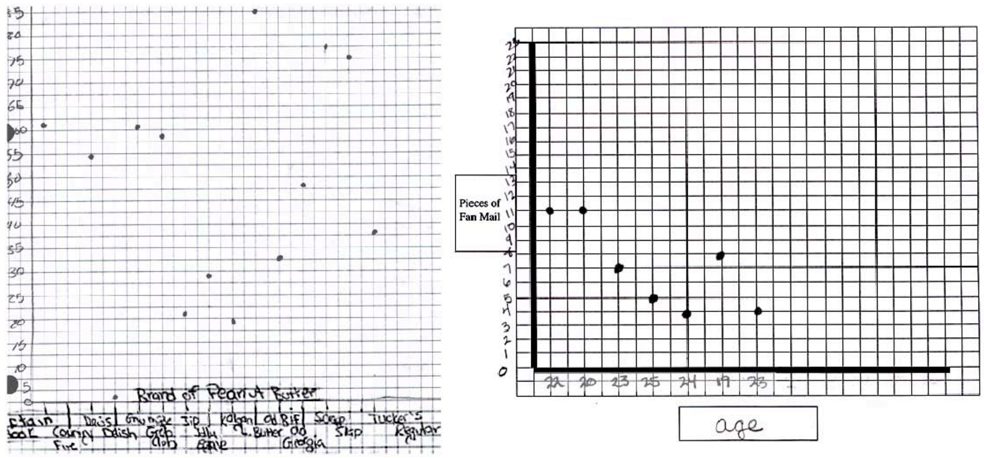

Figure 1: The variable choice (L) and nominalization (R) errors.

example of this is shown in Figure 1L, where the student has placed the names of different brands of peanut butter rather than their quality ratings along the X axis. Students making the related nominalization error treat a quantitative variable already chosen for one axis as if it were a nominal variable For instance, in Figure 1R, the student should have plotted the values on the X axis as an interval scale in numerical order: 19, 20, 21, 22, 23, 24, 25. Instead, that student plotted the individual values of the variable in the same order as they appeared in the original data set, with one value (23) appearing twice: 22, 20, 23, 25, 24, 19, 23. In both cases, the student's actions would be correct if the student was attempting to create a bar graph with a nominal X axis and a quantitative Y axis, but are inappropriate when generalized to scatterplots.

Research has shown that students' difficulty at distinguishing between scatterplots and bar graphs is resilient both to simple supports such as labeling axes with the variables students should use (Baker et al, 2002), and to direct feedback. Baker, Corbett, Koedinger, and Schneider (2003) developed a cognitive tutor lesson which detected and attempted to remediate the variable choice error using model-tracing (Anderson, Corbett, Koedinger, and Pelletier, 1995). As a student worked through exercises, their actions were matched to models both of correct performance and "buggy" performance. If the buggy performance model was a better match to the student's actions than the model of correct performance, then the system gave the student feedback telling the student why his/her action was incorrect. For instance, if a student chose a categorical variable for the X axis, the tutor immediately informed the student: "You have chosen a categorical variable as your X axis. This would be appropriate for a BAR GRAPH rather than the type of graph you've been asked to draw. Ask for help if you're not sure what kind of variable you need." Using this model-tracing tutor for one and a half class sessions moderately improved performance at generating correct scatterplots, but did not significantly reduce the number of students who made the variable choice error on the post-test. In order to develop a tutor that responds effectively to the variable choice and nominalization errors, we must understand why the errors occur. Given the errors' resilience to scaffolding and instruction, two possible accounts are that the errors are procedural misconceptions ("bugs") (Brown and Burton, 1978; vanLehn, 1990), or that they are manifestations of a schematic, "conceptual" misconception (Clement,1982; Minstrell, 1989). A student with a bug might use the right procedure, but make a characteristic error during one step of it; a student with a conceptual misconception might use the wrong schema and procedures flawlessly.

The comparison of computational models, written in ACT-R (Anderson and Lebiere, 1998), provides insight into which account better explains these phenomena (a preliminary discussion of this comparison can be found in Baker, Corbett, and Koedinger, 2003). In this comparison, the variable choice and nominalization errors were modeled by one negative transfer account similar in character to those in Singley and Anderson (1989), and by two accounts derived from the impasse account for bugs given in vanLehn (1990). In the negative transfer model, the correct association between the goal of drawing a bar graph and placing a nominal variable on the X axis is overgeneralized such that an association is instead made between the goal of drawing *any* graph and placing a nominal variable on the X axis. This model was compared to two impasse models: one where the students did not know the difference between nominal and quantitative variables, and one where the students did not have any idea what kind of variables would go on the two axes for any type of graphs. In both of these models, the simulated student reached an impasse when choosing variables or how to represent them, and used the information in the question, selected a variable/variable type randomly, or gave up. Although the impasse models both achieved reasonable fits to the data (r^2=0.90, Mean Absolute Deviation (MAD)=0.26; r^2=0.916, MAD=0.11), the model which accounted for the pattern of student behavior as stemming from overgeneralization of knowledge of bar graphs achieved a better fit to the data (r^2=0.972, MAD=0.06) without introducing more flexibility of fit. In other words, our modeling suggests that it is more likely that the variable choice and nominalization errors stem from applying steps from a correct procedure in an incorrect situation rather than from the generation of an ad-hoc solution to solve an impasse during the execution of a procedure. The errors stem from the failure to distinguish between concepts, and need to be remediated in a fashion that takes this into account. In the next section, we will discuss a new cognitive tutor for scatterplots that incorporates a method successfully used in other curricular settings to remediate this sort of conceptual misconception.

Study 1: Will Contrasting To Bar Graphs Help Students Learn Scatterplots?

Curricular Design

In study 1, we compare the effectiveness of a cognitive tutor which focuses on the contrasts between scatterplots and bar graphs (condition CONTRASTING-CASES), to a cognitive tutor which teaches about scatterplots, but does not explicitly focus on the differences between scatterplots and bar graphs (condition SCATTERPLOT-ONLY).

An approach which has been successful at remediating conceptual misconceptions is the contrasting cases method (Schwartz and Bransford, 1998). In the contrasting cases method, students learn by comparing between cases which have been selected because they differ on specific important features. This assists students in learning which characteristics are important to attend to in distinguishing between categories. Contrasting cases are most effective when combined with direct instruction on why the features that differentiate the cases are relevant. In Schwartz and Bransford (1998), contrasting cases were given first, with conceptual instruction drawing on the distinctions students learned during the contrasting cases. Schwartz and Bransford suggested that this ordering is especially effective for promoting understanding of conceptual instruction. We hypothesized that conceptual instruction might also assist students in understanding a set of contrasting cases. Reversing Schwartz and Bransford's ordering offered several additional advantages for our instructional context: First, this ordering enabled us to combine the conceptual instruction on the features differentiating scatterplots from bar graphs with conceptual instruction on how to generate a scatterplot. Additionally, this ordering enabled us to integrate the contrasting cases into the process of creating a scatterplot used to answer specific questions, with every exercise implicitly demonstrating the practical utility of the distinction between the informational content of bar graphs and scatterplots. Hence, in our study 1, we gave students conceptual instruction (which included instruction on how to generate a scatterplot) before the contrasting cases.

Conceptual instruction was given via a PowerPoint presentation with voiceover and some simple animations. Students went through the PowerPoint presentation at their own pace, although the presence of voiceover tended to keep the students to approximately the same total time. The instruction given differed based on whether the student was in the CONTRASTING-CASES or SCATTEPRLOT-ONLY condition; see Table 1 for a detailed comparison of the conceptual instruction in each condition.

Next the student used the cognitive tutor. In each tutor exercise, the student was given a set of variables (including the variables to use in drawing the graph, and distractor variables of both data types), and a set of questions to answer. Students then generated and interpreted a scatterplot using the interfaces in Figure 2 and 3; see Table 2 for the details of this process. Errors made during tutor use were remediated by the same model-tracing feedback used in Baker, Corbett, Koedinger, and Schneider (2003), with the exception that the feedback for choosing the wrong variable in the SCATTERPLOT-ONLY condition no longer mentioned bar graphs.

TABLE 1: The conceptual instruction in the two conditions of Study 1

	Condition CONTRASTING-CASES	**Condition SCATTERPLOT-ONLY**
1	Introduction to data analysis	Same
2	Definition and examples of categorical and quantitative data	Definition and examples of quantitative data
3	Definition of bar graphs The types of information bar graphs contain The kinds of questions bar graphs can be used to answer The kinds of questions bar graphs cannot be used to answer	Nothing
4	Definition of scatterplots The types of information scatterplots contain The kinds of questions scatterplots can be used to answer	Same
5	How to generate scatterplots (included choosing scale and plotting points)	Same
6	How to use scatterplots to answer questions	Same
7	Review of the differences between scatterplots and bar graphs	Review of scatterplots
8	Introduction to cognitive tutor interface	Same (except for lack of c.c. scaffold)
Min Time	27 minutes, 11 seconds	23 minutes, 6 seconds

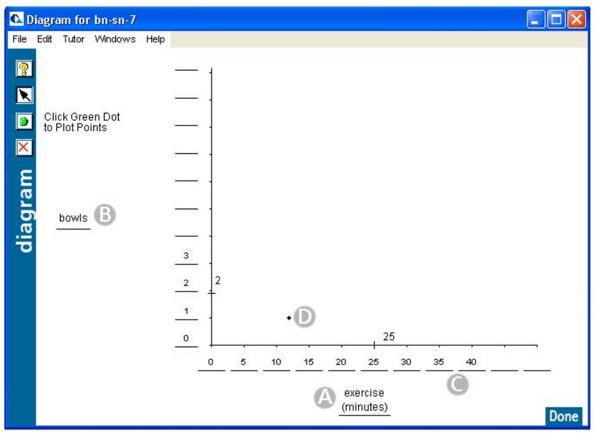

Figure 2: The Cognitive Tutor User Interface

The contrasting cases scaffold was inspired in part by a data format scaffold in Lovett's Statistics Tutor (2001), adapting that scaffold to our different educational context. In the Statistics Tutor, the student is guided to identify each variable's type and to select a representation for those variables. In order to remediate the variable choice error, we add an explicit contrast of the suitability of each variable for each representation, based on the type of information contained in that variable and the type of information used in the representation. Our contrasting cases scaffold is shown in Figure 3. In this scaffold, each variable in the data set is listed, and for each variable the student must first identify whether it is a quantitative ("numerical") variable or a categorical variable. After doing so, the student must identify whether that variable is appropriate or inappropriate for a scatterplot (quantitative variables are appropriate, categorical variables are not), and whether that variable is appropriate or inappropriate for a bar graph (a bar graph uses one variable of each type, so taken individually, a variable of either type is appropriate for use in a bar graph). By having the student decide whether each variable would be appropriate for a scatterplot and/or a bar graph, the scaffold assists the student in understanding the distinction between these two representations of data. Moreover, the student makes this distinction immediately after considering the feature (variable type) that distinguishes the cases, reinforcing the connection between the contrasting cases and the feature that contrasts them. This connects to and reinforces the distinction that was introduced during the conceptual instruction. Since the contrasting cases scaffold can be quite lengthy when the student has a large data set to work with, we introduced an additional new feature to the tutor, termed problem-step fading. In the version of mastery learning used in past cognitive tutors (Anderson et al, 1995), an assessment of which skills the student has mastered drives the selection of new problems. With problem-step fading, the same assessments of student mastery are used to fade a specific step of a process when the student has reached mastery, transforming that step from student-completed exercise to worked-example. Hence, the learner is required to complete fewer and fewer steps as they demonstrate increased mastery. The goal of this approach is to maintain the overall structure of the problem for the student while focusing their time and attention on the more difficult parts of the problem. Other approaches have successfully blended worked-examples and problem-solving within a single exercise (Renkl, Atkinson, Maier, and Staley, 2002), but without using assessments of an individual student's skills to select which steps are worked-examples and problem-solving. We will discuss the effects of problem-step fading in study 2.

TABLE 2: The process of generating and interpreting a scatterplot in each condition of Study 1

Step	Condition CONTRASTING-CASES	Condition SCATTERPLOT-ONLY	See
1	Use the contrasting cases scaffold	Nothing	Fig. 3
2	Choose variables for the X and Y axes	Same	Fig. 2, Labels A&B
3	Choose bounds and scale for each axis	Same	Not shown
4	Label axes with bounds and scale	Same	Fig. 2, Label C
5	Plot points on the graph, by clicking the point tool and then clicking on the graph	Same	Fig. 2, Label D
6	Answer interpretation questions	Same	Not shown

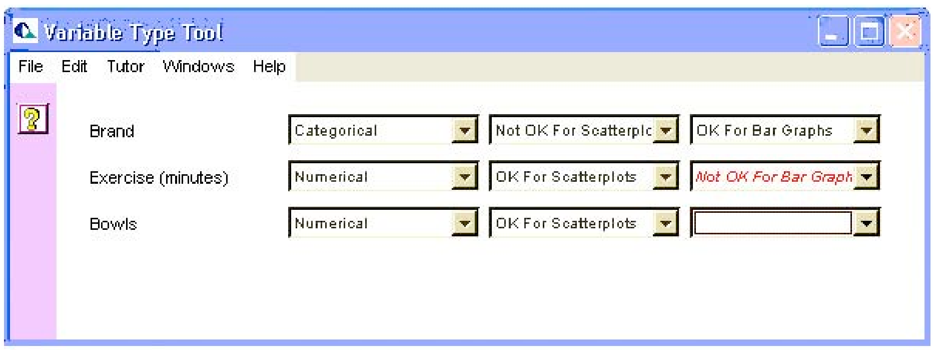

FIGURE 3: The contrasting cases scaffold

Study Design

In study 1, we gave students conceptually different instruction in different conditions. In order to avoid a situation where students in different conditions knew each other and told each other what they had learned, while still working within an authentic learning setting, we conducted this study with homeschool students. The students used the software on their own computers, at home; we requested that the students' parents and siblings not interact with them as they used the software. If two children from the family used the software, they were placed in the same condition, and families with multiple children were distributed randomly between the conditions. We recruited homeschool families by posting ads on homeschooling newsgroups and internet mailing lists for homeschooling families in Pennsylvania. Although the newsgroups and lists were targeted to parents in Pennsylvania, parents from other states read these newsgroups as well, and 44% of the eventual respondents were from other states. After reading our ads, parents went to a webpage and signed up to receive the software through the mail. Participating students took the pre-test (administered by their parents), viewed the PowerPoint presentation, used the cognitive tutor, took the post-test, and finally returned the tests and tutor log files to us by self-addressed stamped envelope. We controlled the amount of time each student spent on the tutor. The tutor allowed students to work for 75 minutes, and at that point, let them complete the problem they were working on, and then quit, telling them they had completed their work with the tutor. A sufficient number of problems were included in the tutor that no student ran out of problems to work on. The pre-test and the post-test each consisted of one of two nearly isomorphic problems, counterbalanced between the pre-test and post-test. In each problem, students were given a data table with three quantitative variables and one nominal variable, and were asked to generate a scatterplot to show the relationship between two of the quantitative variables. The pre-test and post-test were graded in terms of what percentage of the steps of the process of creating a scatterplot were correct – students were given two points for choosing the correct variables (one point for choosing variables of the correct type), two points for correctly labeling each axis with an appropriate axis (one point for only one axis), and two points for making at most one error while plotting points. The pre-test and post-test were similar to the pre-test and post-test in Baker, Corbett, Koedinger, and Schneider (2003), but had a quantitative distractor as well as the nominal distractor found in that study.

Results

Of the 203 students who received the software in 2 mail-outs, 39 eventually returned their pre-test and post-test to us. Response rates were fairly low, but not significantly different between conditions: 22% in CONTRASTING-CASES versus 16% in SCATTERPLOT-ONLY, $\chi^2(1, N=203)=1.31$, $p\sim0.25$. After completing the conceptual instruction, students in the two conditions spent the same time using the tutor (87 minutes in the SCATTERPLOT-ONLY condition, 88 minutes in the CONTRASTING-CASES condition). Students in the CONTRASTING-CASES condition spent 13% of their time using the contrasting cases scaffold, and 20% of their time choosing variables; students in the SCATTERPLOT-ONLY condition spent 22% of their time choosing variables. Average performance improved substantially from pre-test to post-test in both conditions. In the SCATTERPLOT-ONLY condition, average performance improved significantly, from 44% to 88%, $t(15)=3.34$, $p<0.01$. Performance also improved significantly in the CONTRASTING-CASES condition, from 43% to 99%, $t(22)=7.53$, $p<0.001$. There was significantly greater improvement in the CONTRASTING-CASES condition, $F(36,1)=6.41$, $p=0.02$, controlling for pre-test score; hence, the time spent studying the differences between bar graphs and scatterplots seems to have been beneficial.

A similar effect was seen with the number of students who chose correct variables for the X and Y axis in both conditions. In the SCATTERPLOT-ONLY condition, performance improved from 47% to 88%, $\chi^2(1, N=32)=6.78$, $p=0.01$. In the CONTRASTING-CASES condition, performance improved from 48% to 100%, $\chi^2(1, N=46)=16.24$, $p<0.001$. For this measure, there was marginally significantly greater improvement in the CONTRASTING-CASES condition, $F(36,1) = 3.20$, $p=0.08$, controlling for pre-test score. The number of students who committed the variable choice error was significantly different between conditions at pre-test, $\chi^2(1, N=39)=3.99$, $p=0.05$, and there was a similar trend for the nominalization error, $\chi^2(1, N=39)=2.26$, $p=0.13$, so we cannot compare the two conditions directly in terms of these errors (both errors were substantially less common at pre-test in the SCATTERPLOT-ONLY condition). However, considered by itself, the CONTRASTING-CASES condition was effective at remediating these specific errors. The variable choice error occurred 22% of the time at pre-test in the CONTRASTING-CASES condition, and 0% of the time at post-test, $\chi^2(1, N=46)=5.61$, $p=0.02$. Similarly, the nominalization error occurred 13% of the time at pre-test in the CONTRASTING-CASES condition, and 0% of the time at post-test, $\chi^2(1, N=46)=3.21$, $p=0.07$. We will re-consider the difference between the CONTRASTING-CASES and SCATTERPLOT-ONLY conditions in remediating these errors in our discussion of Study 2.

Study 2: Problem-Step Fading

Study Design

Study 2 was designed to determine what role problem-step fading had played in the success of the tutor in study 1, and specifically to see if it had had helped reduce the cost of adding a lengthy scaffold to the tutor. Study 2 also allowed us to see if the contrasting cases-based tutor would be as effective in a classroom setting as it was in a homeschool setting. To investigate these questions, we compared the effectiveness of a cognitive tutor which used problem-step fading (FADING) to the effectiveness of a cognitive tutor where the student had to complete the entire problem on his/her own (NO-FADING). Since the conditions varied only in how each student's time was focused as they used the tutor, the two conditions could be (and were) presented to students in the same class without risk of contamination between conditions. Both conditions were drawn from the CONTRASTING-CASES condition in the previous study; that is, students viewed PowerPoint-based declarative instruction discussing scatterplots in relation to bar graphs, and used the contrasting cases scaffold. Seventy students at two middle schools in suburban Pittsburgh participated in the entirety of this study. All of the participating students were enrolled in a year-long cognitive tutor mathematics course. Each student was given a pre-test, a post-test, viewed the conceptual instruction, and completed at least four exercises in the tutor; students absent during any day of the study were eliminated from analysis. Every student used the tutor for approximately the same amount of time, so some students were able to complete more problems than others. Problems were given in the same order for each student. The pre-test and post-test were the same as in study 1;

which meant that they were identical to the pre-test and post-test in Baker, Corbett, Koedinger, and Schneider (2003), except for different cover stories and the presence of one new quantitative distractor variable.

Results

In this study, students had different experiences in the two tutor conditions (FADING and NO-FADING), but learned about the same amount from pre-test to post-test. In both conditions, students showed significant gains in average performance and made the variable choice error significantly less frequently. Students spent approximately the same time (49.6 minutes in FADING, 46.4 minutes in NO-FADING) actively using the tutor in the two conditions, $t(68)=1.39$, $p=0.17$, but students in the FADING condition completed 61% more problems on average (5.4 versus 3.4), which was significant, $t(68)=4.07, p<0.0001$. Students in the FADING condition spent 48% more time on the difficult skill of choosing axis variables, $t(68)=4.04, p<0.0001$, but 32% less time labeling values along the axis after choosing bounds and scale, $t(68) =2.88$, $p<0.01$. The amount of time spent plotting points, using a scaffold to choose bounds and scale, and using the contrasting cases scaffold were not significantly different between conditions, respectively a 20% difference, $t(68)=1.17$, $p=0.24$, a 5% difference, and a 1% difference. The absence of any difference between conditions in the time spent using the contrasting cases scaffold suggests that the addition of problem-step fading did not reduce the contrasting case scaffold's time cost. In both conditions, there was a significant improvement in the average percentage of each problem correct from pre-test to post-test. In the NO-FADING condition, performance improved from 40% to 69%, $t(37)=-4.36$, $p<0.001$, for a paired t-test. In the FADING condition, performance improved from 41% to 74%, $t(32)=-4.01$, $p<0.001$, for a paired t-test. There was not a significant difference in learning gain between the two conditions, $F(67,1)=0.42$, $p=0.52$, controlling for pre-test. In both conditions, there was a significant increase in the number of students who chose variables of the correct type for both the X and Y axes. In the NO-FADING condition, performance improved from 53% to 87%, $\chi^2(1, N=76)=10.53$, $p=0.001$, and in the FADING condition, performance improved from 56% to 91%, $\chi^2(1, N=64)=9.69$, $p=0.002$; there was not a significant difference in learning gain between the two conditions, $F(67,1)=0.17$, $p=0.68$, controlling for pre-test.
Finally, across the two conditions, the prevalence of the variable choice error decreased from 19% to 6%, which was significant, $\chi^2(1, N=140)=5.42$, $p=0.02$. The nominalization error, on the other hand, did not change in frequency from pre-test to post-test (going from 0% to 1%) (2).

Study 2 replicates our results regarding the effectiveness of the CONTRASTING-CASES condition of study 1. In a different educational setting, with a different population, students showed substantial gains between pre-test and post-test, both in terms of overall learning and reduction in the frequency of the variable choice error. Furthermore, as in study 1 but quite unlike the curriculum in Baker, Corbett, Koedinger, and Schneider (2003), the contrasting cases curriculum used in this study significantly reduced the frequency of the variable choice error. Nonetheless, problem-step fading did not seem to substantially help or hinder students. It shifted time between skills, and allowed students to complete more problems, but did not result in significantly higher post-test performance. It is somewhat surprising that students did not perform significantly different at post-test between conditions, given the difference in their experiences with the tutor; however, given that students in both conditions experienced the same conceptual instruction, and spent a rather short time using the tutor, one possibility is that the difference between the two conditions would have become more articulated given more time.

Conclusion

Overall, the studies presented here suggest that cognitive tutoring curricula can be augmented by the inclusion of contrasting cases, where appropriate. In these studies, we developed a cognitive tutoring curriculum which combined instruction on the salient features differentiating scatterplots and bar graphs with a series of comparisons of variables appropriate and inappropriate for use in these representations. This curriculum succeeded both in substantially improving understanding of scatterplots and in greatly reducing the frequency of errors where students incorporated bar-graph features into scatterplots. This curriculum was effective both in homeschool and classroom settings. Students in 4% of public U.S. secondary schools now use one or more years of mathematics instruction based upon cognitive tutors with model-tracing remediation. These tutors have been quite successful (Koedinger et al 1997), producing over 1 SD improvements over traditional classroom instruction, but some specific lessons have been less successful at remediating errors than others. Remediating such resilient errors will be a focus of future work at improving existing cognitive tutoring curricula. The comparison of computational models is a powerful way to determine which student errors stem from conceptual misconceptions, and to develop a precise and well-founded account of what distinctions the students are failing to make. Once we have done this, we may be able to substantially improve our existing cognitive tutoring curricula by extending the contrasting cases/conceptual instruction approach to these situations.

In the long-term, determining what types of pedagogies are most effective for different types of educational challenges will enable the development of tutoring systems which can adapt not just to the fact that a student is having difficulty, but to why the specific student is most likely to be having that difficulty. Then the systems can more effectively aid the student in mastering that difficulty. More broadly, the studies presented here show that the combined contrasting cases/salient features approach advanced in Schwartz and Bransford (1998) can be effectively incorporated into existing curricula which are designed for a very different educational domain, student population, and overall instructional approach. In addition, we show that the specific ordering of contrasting cases and conceptual instruction used by Schwartz and Bransford is not essential to this approach's success. In Schwartz and Bransford's approach, contrasting cases were given first; in our studies, conceptual instruction was given first. In both sets of studies, the combined approach was successful. This demonstrates that the contrasting cases/salient features approach is a powerful technique for remediating conceptual misconceptions, and one that can be applied in a wide variety of curricular settings.

Endnotes

(1) In discrete bar graphs and most other graphical representations, categorical and nominal variables are represented in the same fashion.

(2) Unlike the variable choice error, the nominalization error's frequency has been fairly unstable across studies – determining why this is the case will be an interesting question for future research.

References

Anderson, J.R., Corbett, A.T., Koedinger, K.R., Pelletier, R. (1995). Cognitive Tutors: Lessons Learned. *Journal of the Learning Sciences, 4(2),* 167-207.

Anderson, J. R. & Lebiere, C. (1998). *The atomic components of thought.* Mahwah, NJ: Erlbaum.

Baker, R.S., Corbett, A.T., and Koedinger, K.R. (2001) Toward a Model of Learning Data Representations. *Proceedings of the Cognitive Science Society Conference*, 45-50

Baker, R.S., Corbett, A.T., and Koedinger, K.R. (2002) The Resilience of Overgeneralization of Knowledge about Data Representations. *Presented at the 2002 Annual Meeting of the American Educational Research Association.*

Baker, R.S., Corbett, A.T., and Koedinger, K.R. (2003) Statistical Techniques for Comparing ACT-R Models of Cognitive Performance. In *Proceedings of the 10th Annual ACT-R Workshop,* 129-134.

Baker, R.S., Corbett A.T., Koedinger K.R., Schneider, M.P. (2003) A Formative Evaluation of a Tutor for Scatterplot Generation: Evidence on Difficulty Factors. *Proceedings of the Conference on Artificial Intelligence in Education*, 107-115.

Brown, J.S., Burton, R.R. (1978) Diagnostic Models for Procedural Bugs in Basic Mathematical Skills. *Cognitive Science, 2,* 155-192.

Clement, J. (1982) Students' preconceptions in introductory mechanics. *American Journal of Physics, 50,* 66-71.

Hancock, C., Kaput, J.J., & Goldsmith, L.T. (1992) Authentic Inquiry With Data: Critical Barriers to Classroom Implementation. *Educational Psychologist*, *27(3),* 337-364.

Koedinger, K.R., Anderson, J.R., Hadley, W.H., & Mark, M.A. Intelligent Tutoring Goes to School in the Big City. (1997) *International Journal of Artificial Intelligence in Education*, *8*, 30-43.

Lehrer, R. and Schauble, L. (2001) *Investigating real data in the classroom: Expanding children's understanding of math and science*. New York: Teachers College Press.

Lovett, M. (2001) A Collaborative Convergence on Studying Reasoning Processes: A Case Study in Statistics. In D. Klahr and S. Carver (Eds.) *Cognition and Instruction: 25 Years of Progress*. Mahwah, NJ: Erlbaum.

McGatha, M., Cobb, P., McClain, K. (in press) An Analysis of Students' Initial Statistical Understandings: Developing a Conjectured Learning Trajectory. To appear in *Journal of Mathematical Behavior*.

Minstrell, J. (1989) Teaching Science for Understanding. In Resnick, L.B., and Klopfer, L.E. (Eds.) *Toward the Thinking Curriculum: Current Cognitive Research*. Alexandria, VA: Asssociation for Supervision and Curriculum Development, 129-149.

National Council of Teachers of Mathematics. (2000) *Principles and Standards for School Mathematics*. Reston, VA: National Council of Teachers of Mathematics.

Renkl, A., Atkinson, R.K., Maier, U.H., Staley, R. (2002) From Example Study to Problem Solving: Smooth Transitions Help Learning. *The Journal of Experimental Education, 70 (4),* 293-315.

Schwartz, D.L. and Bransford, J.D. (1998) A Time for Telling. *Cognition and Instruction, 16 (4),* 475-522.

Singley, M.K. and Anderson, J.R. (1989) *The Transfer of Cognitive Skill.* Cambridge, MA: Harvard University Press.

Tabachneck-Schijk, H.J.M., Leonardo, A.M., Simon, H.A. (1997) CaMeRa: A Computational Model of Multiple Representations. *Cognitive Science, 21 (3),*305-350.

Tufte, E.R. (1983) *The Visual Display of Quantitative Information*. Cheshire, CT: Graphics Press.

vanLehn, K. (1990) *MindBugs: The Origins of Procedural Misconceptions.* Cambridge, MA: MIT Press.

Acknowledgments

We would like to thank Angela Wagner, Jay Raspat, Megan Naim, Katy Getman, Frances Battaglia, Pauline Masley, and Heather Frantz for assistance in conducting the studies reported in this paper, Michael Schneider with assistance in software design and implementation, and Jack Zaientz, Lisa Anthony, Lara Triona, John Kowalski, Marsha Lovett, and Santosh Mathan for helpful discussions and suggestions. This work was funded by an NDSEG (National Defense Science and Engineering Graduate) Fellowship, and by NSF grant 9720359 to "CIRCLE: Center for Interdisciplinary Research on Constructive Learning Environments".