

Running Head: Responding to Problem Behaviors in Cognitive Tutors

Responding to Problem Behaviors in Cognitive Tutors:
Towards Educational Systems Which Support All Students

Ryan S.J.d. Baker, Ph.D.
Learning Sciences Research Institute
University of Nottingham

Albert T. Corbett, Ph.D.
Human-Computer Interaction Institute
Carnegie Mellon University

Kenneth R. Koedinger, Ph.D.
Human-Computer Interaction Institute
Carnegie Mellon University

Corresponding author:
Ryan S.J.d. Baker, Ph.D.
Learning Sciences Research Institute
University of Nottingham, Jubilee Campus
Wollaton Road
Nottingham, NG8 1BB
UNITED KINGDOM
Ryan.Baker@Nottingham.ac.uk

To appear in NADD Bulletin.

Introduction.

In the last twenty years, interactive learning environments and computerized educational supports have become a ubiquitous part of students' classroom experiences, at all levels of the educational system. As interactive learning environments play an increasingly large role in the educational system, it becomes increasingly important that they be made appropriate for the full spectrum of students, including students with learning and developmental disabilities, behavioral problems, and language deficits. For this reason, in recent years, there has been increasing research into how interactive learning environments can accommodate students with disabilities (Stiles, 2001; Cole et al., 2003; Moore & Taylor, 2003). The presence of a developmental disability along with a mental illness heightens the need for effective educational interventions (Prout, 2005). Prout eloquently framed the need for increased attention to diagnosis and intervention for school-aged youth with dual diagnoses (2005). The scarce attention to this population has tended to focus on treatment rather than educational intervention.

In this paper, we will discuss work towards ensuring that one type of educational system, Cognitive Tutors (Koedinger, Anderson, Hadley, & Mark, 1997), help all students achieve. Cognitive Tutors are a type of interactive learning environment which use cognitive modeling and artificial intelligence in order to adapt to individual differences in student knowledge and learning. Cognitive Tutors are designed to be integrated into full-year curricula; students use the Cognitive Tutor software two days a week, and engage in group work and classroom lecture on other class days. Within the Cognitive Tutor environment, each student individually completes mathematics

problems. The Cognitive Tutor environment breaks down each mathematics problem into the steps of the process used to solve the problem, making the student's thinking visible. The Cognitive Tutor environment also supports the student in representing a problem in multiple ways, including equations, tables, and graphs (Koedinger et al., 1997).

As the student works through a problem, their work is analyzed by a model of student cognition in the problem domain (cf. Anderson, Corbett, Koedinger, & Pelletier, 1995). Incorrect answers turn red; if the student's answer is not just incorrect but indicative of a larger misconception, the student is given a text explanation (see Figure 1). If a student is struggling, he or she can also request a hint. When the student requests a hint, the system first gives a conceptual hint. The student can request further hints, which become more and more specific until the student is given the answer (see Figure 2).

[Place Figure 1 around here]

[Place Figure 2 around here]

As the student works through the problems in a specific curricular area, the system uses Bayesian knowledge-tracing (Corbett & Anderson, 1995) to determine which skills that student is having difficulty with – calculating the probability that the student knows each skill based on that student's history of responses within the tutor. Using these estimates of student knowledge, the tutor can give each student exercises which are relevant to the skills which that student is having difficulty with. The tutor does not move on to new material until the student has reached mastery (defined as a 95% probability of

knowing each skill) on all of the skills relevant to the current curricular area. Hence, each student works at a pace appropriate to his/her learning, and when the student reaches each section of the tutor curriculum, he or she is completely prepared with all of the pre-requisite knowledge necessary to succeed on the new material.

Cognitive Tutors have been highly successful at improving the learning outcomes of at-risk students. In a large-scale study in El Paso, TX, previously failing students who re-took Algebra using a Cognitive Tutor curriculum had significantly better performance on state standardized exams than students who re-took Algebra using a traditional curriculum (Carnegie Learning, 2002). In another large-scale study in Kent, WA, previously failing students using a Cognitive Tutor Algebra curriculum had significantly higher learning gains on a standardized exam than students using a traditional mathematics curriculum who had passed their previous mathematics course (Plano, 2004). Because of their documented and replicated success at improving learning outcomes, Cognitive Tutor mathematics curricula are currently in use in around 6% of U.S. high schools.

Though Cognitive Tutors were designed for mainstream mathematics classes, rather than special education classes, recent research has indicated that Cognitive Tutors are also extremely effective for students with learning disabilities and English-language deficits. In a large-scale experimental study conducted throughout the Miami-Dade County school district, students with learning disabilities and language deficits performed significantly better on a year-end state standardized exam after completing an Algebra I course taught using Cognitive Tutors, instead of a course taught using a traditional mathematics curriculum (Sarkis & The Reliability Group, 2004). In this study, the

performance gap between students with learning disabilities and students in mainstream classes was reduced by 3/4 when students used Cognitive Tutors instead of a traditional curriculum. Similarly, in Kent, WA, students with English-language deficits performed almost twice as well if they used a Cognitive Tutor curriculum rather than a traditional curriculum (Plano, 2004).

These reports have not speculated on *why* Cognitive Tutors are more effective for students with learning disabilities and English-language deficits. Many of the same features that make Cognitive Tutors beneficial for mainstream students may explain why special-needs students also benefit from using Cognitive Tutors. For instance, the software's ability to select exercises tailored to each student's current knowledge is likely to benefit special-needs students in much the same fashion as it benefits mainstream students. The individual pacing is likely to benefit students with learning disabilities, who may need extra time to mastery the material. The use of multiple representations may benefit students whose learning disabilities hinder their ability to learn verbally or visually. The use of multiple graphical representations may also explain part of the software's success with students who have limited English or linguistic ability. An expansion of educational strategies for learners with the dual diagnoses of mental illness and developmental disability may be of great benefit (cf. Prout, 2005).

Gaming the System

Though there is evidence that Cognitive Tutors are effective for students with learning disabilities and English-language deficits, a recent study has suggested that Cognitive Tutors are less effective for students who engage in some problem behaviors. In particular, students who engage in a behavior termed "gaming the system" have

significantly poorer learning than other students (Baker, Corbett, Koedinger, & Wagner, 2004). Gaming the system is defined as attempting to complete problems by exploiting properties of the system rather than by learning the material and trying to use that knowledge to answer correctly. Gaming has been documented in a variety of types of interactive learning environments (Baker et al., 2004; Cheng & Vassileva, 2005; Magnussen & Misfeldt, 2004) and even in interactions between students and human teachers (Arbreton, 1998). In the specific context of Cognitive Tutors, gaming the system consists of guessing systematically and requesting hints until the answer is given, without reading the earlier hints. While there are no data to document this phenomenon in the population of students with dual diagnosis, it can be postulated to exist at levels similar to those in the population of students with learning disabilities.

Students who game Cognitive Tutors tend to start out with lower prior knowledge of the material than the rest of the class, and tend to have lower scores on general tests of academic achievement. Though many students with low prior knowledge catch up to the rest of the class, students who game fall further behind over the course of the tutor lesson. (Baker et al., 2004). Gaming is particularly detrimental for students who tend to game on the specific problem steps that they do not know (Baker, Corbett, & Koedinger 2004).

Hence, existing Cognitive Tutors appear to be effective for many students with learning disabilities but not for students who engage in a specific problem behavior, gaming the system. In order to adapt appropriately to the full spectrum of students, Cognitive Tutors will need to adapt appropriately when a student games the system. In this paper, we detail a system that adapts when students game, and present a case study showing how this system's adaptive behavior can improve gaming students' learning.

Responding Automatically to when Students Game the System.

To address gaming, we developed a new component for Cognitive Tutoring software – an animated agent named “Scooter the Tutor”. Scooter is designed to reduce the incentive to game, by giving gaming students extra exercises, while at the same time helping students learn the material they were avoiding by gaming, by targeting those extra exercises to the exact steps a student gamed. Scooter was also designed to affect non-gaming students as minimally as possible – if a student does not game, their experience using the tutor changes very little.

[Place Figure 3 around here]

Scooter the Tutor is a puppy character, using graphics from the Microsoft Office Assistant, modified to enable Scooter to display a wider range of emotions. Scooter uses an artificial intelligence-based detector to determine whether a student is gaming (on steps the student does not already know), and what problem steps they are gaming on (Baker, Corbett, & Koedinger 2004). If a student is not gaming, Scooter looks happy and occasionally gives the student positive messages. If there is moderate evidence of gaming, Scooter looks upset, and gives the student a warning message. If there is strong evidence of gaming, Scooter responds in a more serious fashion, giving the student a set of supplementary exercises designed to give the student another chance to cover the material that the student may bypassed by gaming. These supplementary exercises include items which require understanding one of the concepts required to answer the step the student gamed past, and questions about what role the step the student gamed

past plays in the overall process of solving the problem. Examples of Scooter's expressions and supplementary exercises are given in Figure 3.

Case Studies.

In this section we will investigate, in detail, how the tutoring system's adaptation to gaming changes students' experience and learning within the tutor. To this end, we present a pair of case studies, comparing the behavior, experiences, and learning of two students who each engaged in a considerable amount of gaming. One of these students used a tutoring system which adapted to gaming (i.e. with Scooter); the other student used a tutoring system without such adaptation.

Each student was drawn from a set of middle-school classrooms in the Pittsburgh suburbs which used Cognitive Tutors as their regular curriculum, year-round, and which were participating in an experimental study comparing a tutor with Scooter to a tutor without Scooter (cf. Baker et al., to appear). Within the case studies that we will present, each student worked with the same tutor lesson on creating and interpreting scatterplots of data.

These two students were selected because of their similarity on the pre-test for the scatterplot unit, and the similarity of their behavior within the tutor, allowing us to concretely examine how Scooter affects students' experiences and learning.

Using a Tutor That Did Not Adapt to Gaming. "Pat" (not his/her real name; neither names nor gender were recorded) performed poorly on his unit pre-test. In particular, Pat's answers indicated the presence of a pair of common misconceptions (cf. Baker, Corbett, & Koedinger, 2002), where students select variables which would be more appropriate for a bar graph than a scatterplot (one nominal variable and one

quantitative variable), and represent quantitative variables on the axis as if they were nominal variables.

Pat was observed gaming 10% of the time while using the tutor (by a human observer using a published quantitative observational method [Baker et al., 2004]). According to the detector, Pat never gamed on the steps involving the skill of representing a quantitative variable on the axis. However, the detector indicated that Pat gamed 19 times (4.75 times per problem completed) on the problem steps involving skills in selecting appropriate variables for a scatterplot.

On the unit post-test, Pat did not make any errors in representing quantitative variables on the axis. However, Pat was not able to select appropriate variables for a scatterplot, continuing to select one nominal variable and one quantitative variable, as he did on the pre-test. Hence, Pat learned a difficult skill he did not game on, but did not learn the difficult skill he did game on.

Using a Tutor That Adapted to Gaming. Like Pat, “Chris” performed poorly on the unit pre-test. In particular, Chris’s answers indicated the presence of one of the same two misconceptions that Pat’s answers demonstrated – Chris also selected variables which would be more appropriate for a bar graph than a scatterplot (one categorical variable and one quantitative variable). However, Chris was able to accurately represent a quantitative variable on the axis, on the unit pre-test.

Chris was observed gaming the same proportion of time as Pat gamed – 10% of the time. While using the tutor, Chris gamed 32 times (5.4 times per problem completed) on steps that involved the skills necessary to choose appropriate variables. Because Chris

gamed on these steps, the tutor gave Chris six sets of supplementary exercises on this material (one exercise per problem completed).

On the unit post-test, Chris was able to select appropriate variables for a scatterplot. Hence, Chris learned a difficult skill, even though he gamed on it. In addition, Chris continued to succeed at accurately representing quantitative variable on the axis.

Discussion and Conclusions.

This pair of case studies demonstrate how a Cognitive Tutor can constructively respond to a student's choice to game the system. Pat and Chris started out with the same misconception, that scatterplots should contain the same variables as bar graphs. Both Pat and Chris repeatedly gamed the system on problem steps relevant to that misconception, attempting to complete problems without having to learn that specific skill. Pat used a tutor that did not respond to gaming, and Pat did not learn the material by the post-test. Chris's tutor, on the other hand, responded to gaming by giving Chris supplementary exercises covering the skills Chris had gamed on. By the post-test, Chris had learned the skills that he was attempting to bypass by gaming.

It is worth noting that Chris did not stop gaming, even after receiving multiple supplementary exercises. Hence, Chris's learning gains likely came from receiving more opportunities to learn the material he bypassed by gaming, rather than from the exercises changing the software's incentive structure and leading Chris to stop gaming. This suggests that in some cases it may not be necessary for a system to actually eliminate problem behaviors. Instead, it may be sufficient for a system to mitigate those behaviors' negative effects.

However, the choice to mitigate or eliminate a problem behavior may depend in part on the frequency of the problem behavior. Though Chris and Pat each gamed heavily on the steps they had the most difficulty with, their overall frequency of gaming was moderate (10%). In our past work, we have observed very few students gaming more than a quarter of the time. A problem behavior which students engage in to the complete exclusion of appropriate software usage may need to be addressed differently than a problem behavior which students engage in more selectively.

More generally, it appears that Cognitive Tutors can effectively and appropriately respond to problem behaviors such as gaming the system. An important area of future work will be determining whether supplementary exercises given on gamed steps will be effective at improving gaming students' learning within a year-long curriculum.

At this point, Cognitive Tutors have been shown to be effective for students with learning disabilities and English-language deficits; by modifying Cognitive Tutors so that they can adapt to problem behaviors, we can move towards making Cognitive Tutors appropriate for, and useful to, all students. Gaming the System has been shown to be the primary problem behavior that affects student learning in mainstream Cognitive Tutor classes, and the study presented here shows that Cognitive Tutors can be modified to adapt to gaming the system in a way that benefits learning. It is not yet proven, however, that the supplemental exercises Scooter uses to benefit learning will be appropriate for all students; in particular, Scooter's exercises may be less effective for students with some types of learning disabilities, if those students' gaming reduces the benefits of the multi-representational support offered in the rest of the Cognitive Tutor lesson.

Another question for future research is whether gaming will be the predominant problem behavior among students with mental illnesses, especially those students with a combination of mental illness and developmental or learning disabilities, or whether other problem behaviors will dominate these students' interactions with Cognitive Tutors. Different problem behaviors are likely to require different adaptations on the part of the software. In order to determine what problem behaviors students with the dual diagnoses of a mental illness and a developmental disability engage in when using Cognitive Tutors, it may be necessary to repeat the observational studies that established the link between gaming and poorer learning (cf. Baker et al., 2004), in classes with dually-diagnosed students.

In the long-term, it will be valuable to determine whether the sort of intervention detailed in this paper can be incorporated into other types of interactive learning environments. It is not clear whether the exact type of supplementary exercises we use here will be useful within all types of learning environments; at minimum, though, other types of learning environments can probably be improved using an approach similar to the approach we used. That approach starts by identifying which problem behaviors are associated with poorer learning, developing systems that can automatically detect those behaviors, and finally re-designing systems to either guide students to reduce their problem behaviors, or give students different opportunities to learn the material their problem behaviors cause them to miss. As school education increasingly incorporates interactive learning environments, making sure those environments adapt to relevant problem behaviors will ensure that the benefits of those environments extend to all students.

Acknowledgements.

This research was supported by IERI grant number REC-0437794 to " Learning-Oriented Dialogue in Cognitive Tutors: Towards a Scalable Solution to Performance Orientation". We would like to thank Angela Wagner, Ido Roll, Steve Ritter, Tom McGinnis, Jay Raspat, Meghan Naim, Dina Crimone, Russ Hall, Sue Cameron, and Francis Battaglia for assisting in the administration of the study discussed here.

For further information, contact Dr. Baker at: Ryan.Baker@Nottingham.ac.uk

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.
- Arbreton, A. (1998) Student Goal Orientation and Help-Seeking Strategy Use. In S.A. Karabenick (Ed.), *Strategic Help Seeking: Implications For Learning And Teaching*, pp. 95-116, Mahwah, NJ: Lawrence Erlbaum Associates.
- Baker, R.S., Corbett, A.T., & Koedinger, K.R. (2002) The Resilience of Overgeneralization of Knowledge about Data Representations. Paper presented at the Annual Meeting of the American Educational Research Association, New Orleans, LA, USA.
- Baker, R.S., Corbett, A.T., & Koedinger, K.R. (2004) Detecting Student Misuse of Intelligent Tutoring Systems. In J.C. Lester, R.M. Vicari, and F. Paraguaçu (Eds.) *Lecture Notes in Computer Science 3220: Intelligent Tutoring Systems*, 531-540. Berlin: Springer-Verlag.
- Baker, R.S., Corbett, A.T., Koedinger, K.R., Evenson, S., Roll, I., Wagner, A.Z., Naim, M., Raspat, J., Baker, D.J., & Beck, J.E. (to appear) Adapting to When Students Game an Intelligent Tutoring System.
- Baker, R.S., Corbett, A.T., Koedinger, K.R., & Wagner, A.Z. (2004) Off-Task Behavior in the Cognitive Tutor Classroom: When Students "Game The System". *Proceedings of ACM CHI 2004: Computer-Human Interaction*, 383-390.
- Carnegie Learning, Inc. (2002) *Results from El Paso, Texas* (Cognitive Tutor Research Report TX-01-01). Pittsburgh, PA.

- Cheng, R., & Vassileva, J. (2005) Adaptive Reward Mechanism for Sustainable Online Learning Community. *Proc. of the International Conference on Artificial Intelligence in Education*, 152-159. Amsterdam: IOS Press.
- Cole, R., van Vuuren, S., Pellom, B., Hacıoglu, K., Ma, J., Movellan, J. Schwartz, S., Wade-Stein, D., Ward, D., & Yan, J. (2003) Perceptive Animated Interfaces: First Steps Toward a New Paradigm for Human-Computer Interaction. *Proceedings of the IEEE: Special Issue on Multimodal Human Computer Interface*, 91 (9), 1391-1405.
- Corbett, A.T., & Anderson, J.R. (1995) Knowledge Tracing: Modeling the Acquisition of Procedural Knowledge. *User Modeling and User-Adapted Interaction*, 4, 253-278.
- Koedinger, K.R., Anderson, J.R., Hadley, W.H., & Mark, M. (1997) Intelligent Tutoring Goes to School in the Big City. *International Journal of Artificial Intelligence in Education*, 8, 30-43.
- Magnussen, R., & Misfeldt, M. (2004) Player Transformation of Educational Multiplayer Games. *Proceedings of Other Players*. Available at <http://www.itu.dk/op/proceedings.htm>
- Moore, D.J., & Taylor, J. (2000) Interactive multimedia systems for people with autism, *Journal of Educational Media*, 25 (3), 169-177.
- Plano, G.S. (2004) *The Effects of the Cognitive Tutor Algebra on Student Attitudes and Achievement in a 9th Grade Algebra Course*. Unpublished Doctoral Dissertation, Seton Hall University.

Prout, H.T. (2005). Dual diagnosis in children and adolescents: issues and opportunities.

The NADD Bulletin, 8(1).

Sarkis, H., & The Reliability Group (2004) *Cognitive Tutor Algebra 1 Program*

Evaluation. Presentation to Miami-Dade County Public Schools. Available at

http://www.carnegielearning.com/web_docs/sarkis_2004.pdf

Stiles, M.J. (2001) Disability Access to Virtual Learning Environments. Case study

published by *Disability and Information Systems in Higher Education*, available

at <http://www.dmag.org.uk/resources/casestudies/stilesfull.asp>

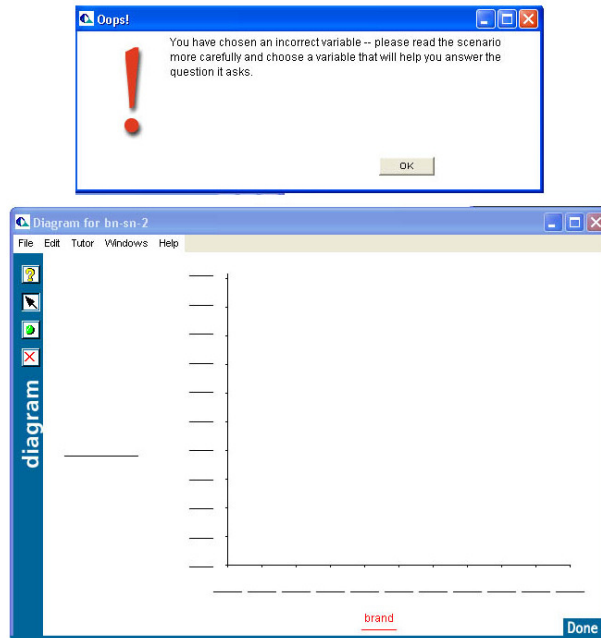


Figure 1: Within a Cognitive Tutor environment, the student has made an error associated with a specific misconception, so they receive a text explanation of their misconception (top window). The student's answer is labeled in red, because it is incorrect (bottom window).

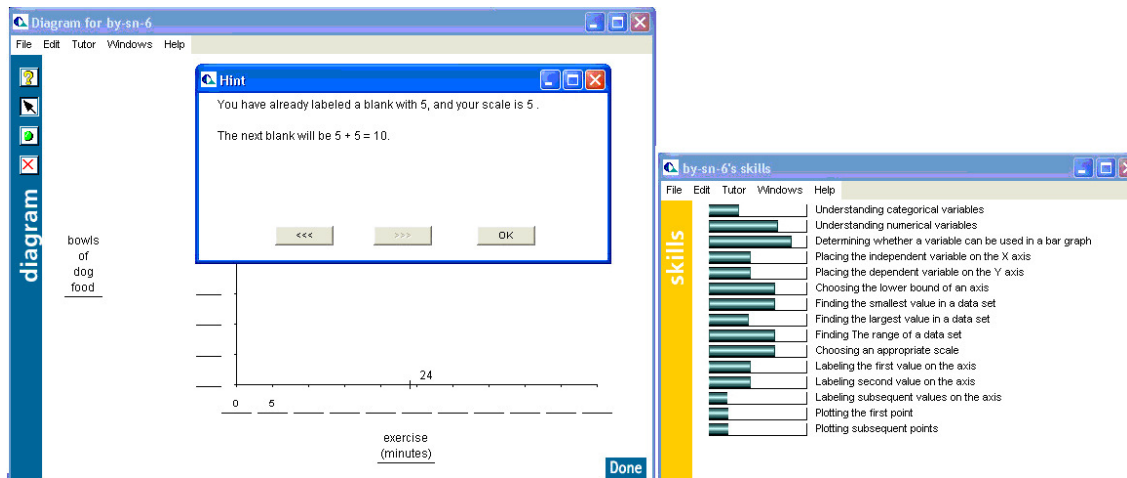


Figure 2: The student has requested multiple hints, and has now received the answer. The window on the left (beneath the hint window) shows the student's answers up until this point; the Cognitive Tutor's estimates of the student's skills are shown in the window on the right. Other problem windows (such as the problem scenario and data table) are not shown.

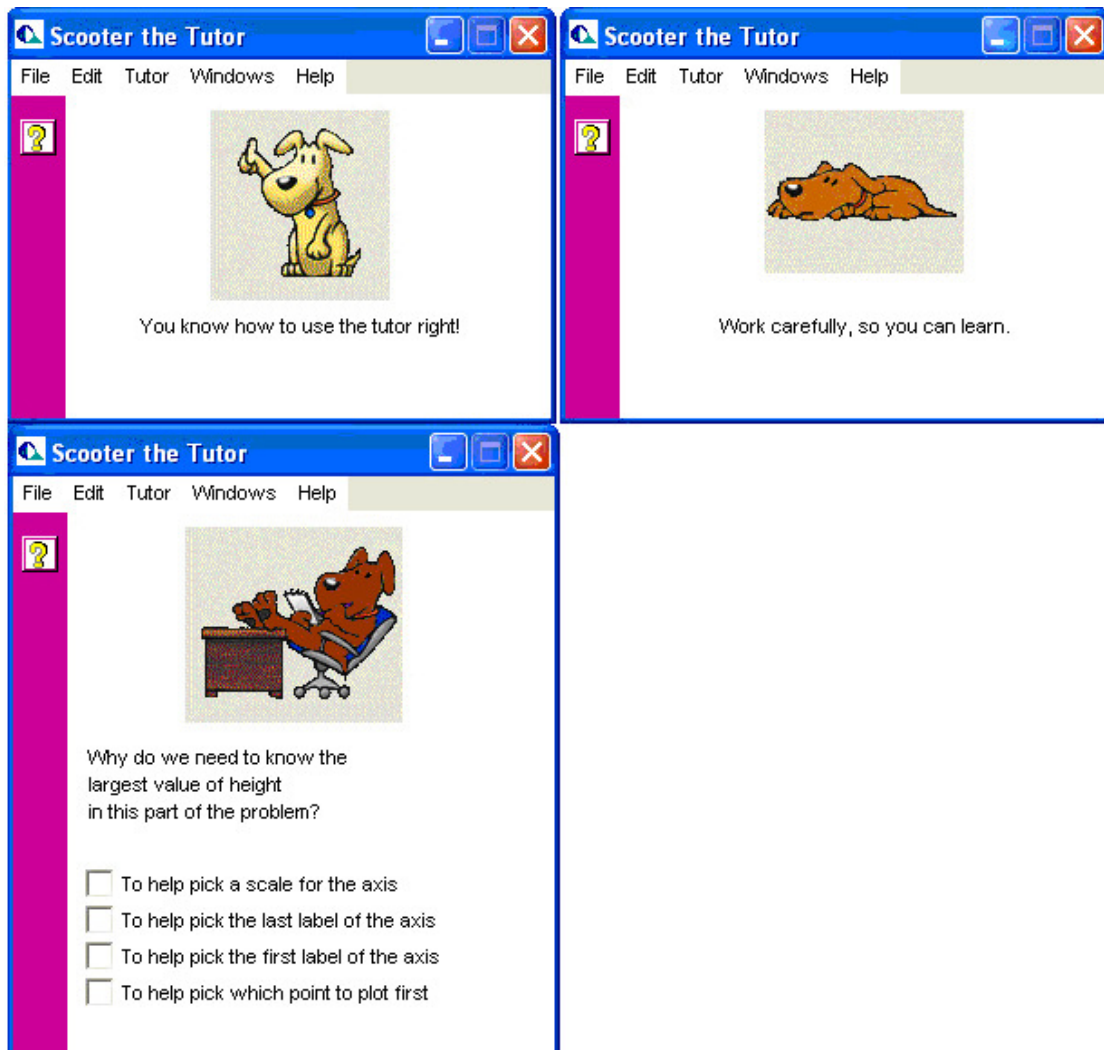


Figure 3. Scooter the Tutor: Top-left, looking happy when the student has not been gaming. Top-right, looking unhappy when the student has been gaming moderately. Bottom-left, intervening to give a supplementary exercise to a student who has been gaming a considerable amount.