

A First Evaluation of the Instructional Value of Negotiable Problem Solving Goals on the Exploratory Learning Continuum

Carolyn ROSÉ, Vincent ALEVEN, Regan CAREY, & Allen ROBINSON
*Human-Computer Interaction Institute and Department of Mechanical Engineering
Carnegie Mellon University,
5000 Forbes Avenue, Pittsburgh PA 15213*

Abstract. We evaluate the effectiveness of a tutorial-dialogue based approach to guided exploratory learning involving problem solving goals that are negotiated between tutor and student rather than dictated by the tutor or freely chosen by the student. This approach, referred to as Negotiable Problem Solving Goals (NPSG), is located on a previously untested space on what we call The Exploratory Learning Continuum. The results of our empirical classroom investigation provide design recommendations for a new type of tutorial dialogue system and strong evidence in favor of tutorial dialogue support in exploratory learning environments.

Introduction

The tutorial dialogue literature provides us with many convincing proofs of the technical feasibility of tutorial dialogue systems [e.g., 7,13,6,1]. What is needed now is insight on how to wield that technology to benefit student learning beyond what is possible with more standard forms of interaction supported by state-of-the-art tutoring systems. Looking at naturalistic human tutorial dialogue inspires us to broaden our view of what intelligent tutoring systems can provide to students, and to consider forms of interaction that are not typically supported by current intelligent tutoring systems. One of the major research goals of the CycleTalk project [14] has been to investigate the instructional effectiveness of novel ways of using tutorial dialogue technology in an exploratory learning environment.

We investigate two separate dimensions that have framed much of the literature on exploratory learning. We evaluate the effectiveness of a tutorial-dialogue based approach involving problem solving goals negotiated between tutor and student rather than dictated by the tutor or freely chosen by the student. This approach, which we refer to as Negotiable Problem Solving Goals, is located on a previously untested space on what we call The Exploratory Learning Continuum. Although our experimental manipulation involves the use of human tutors, the results of our investigation provide design recommendations for a new type of tutorial dialogue system that holds promise for demonstrating the potential contribution tutorial dialogue technology can make to the field of Intelligent Tutoring. In the remainder of the paper we review the exploratory learning literature, the specifics of our experimental design, an analysis of our results, and development plans.

1. Review of Background Literature on Exploratory Learning

A popular conceptualization of exploratory learning is that what distinguishes “exploratory learning” from “non-exploratory learning” is the level at which goals are provided to the learner. Exploratory learning is associated with “high level goals” such as “survive in this simulation environment” or under specified goals such as “find all implications that can be drawn from these premises”. In contrast, non-exploratory learning is associated with “low level goals” or “fully specified goals” such as “solve this equation” or “verify whether this implication is true”. Non-exploratory learning in this conceptualization may involve means-ends analysis; however the search is directed down a small number of “correct” paths. We argue that all learning is exploratory, and alternative learning tasks or learning environments can be placed along a continuum, which we refer to as The Exploratory Learning Continuum.

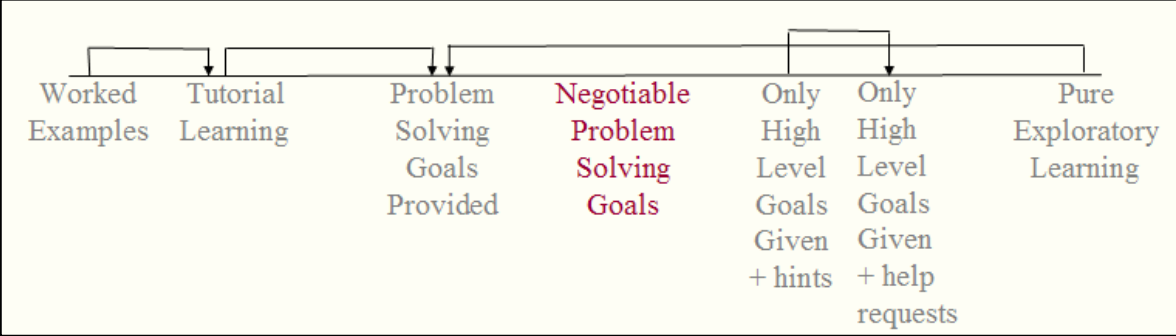


Figure 1 The Exploratory Learning Continuum: Arrows connect approaches that have been experimentally compared in published works. The arrow points from the less effective approach towards the approach shown to be more effective.

Previous investigations of exploratory learning have compared student learning in conditions such as (1) passive worked example studying, (2) active but totally guided tutorial learning, (3) problem solving, and (4) unguided exploration. On the macro-level, what is manipulated is the amount of structure provided for students. In (1) and (2), for example, students make no choices whatsoever, although students in (2) are more active than students in (1). High level goals are set, and low level steps are provided. In (3), problem solving goals are made for the student, but the student chooses how to satisfy those goals through means-ends analysis. In (4), the student sets problem solving goals and chooses how to satisfy those goals. Thus, in (4) the student has the greatest autonomy, but the student is limited by their own conception of what is possible and valuable to explore. In (3) the student is prompted to explore areas in the space of possibilities that they may not have thought of by themselves. Furthermore, they reap the benefits of exploring alternative ways of achieving those goals. However, they do not get the practice setting goals for themselves that students in (4) get.

Many state-of-the-art tutoring systems fall into the problem solving category where problem solving goals are dictated. It is no coincidence since published investigations along the Exploratory Learning Continuum have typically shown this place on the continuum to be particularly effective. For example, Charnay & Reder (1986) compare Worked Examples, Tutorials, Problem Solving, and Pure Exploration. Worked examples mixed with problem solving was the best combination, consistent with other similar published results [16]. Along similar lines, Klahr & Nigam (2004) have shown in an empirical investigation of children learning the scientific method that tutorial based learning mixed with problem solving is more efficient than pure exploratory learning. Other work has explored a part of the continuum in between problem solving and pure exploratory learning. In the light of a series of previous results showing the benefits of guided exploration over pure exploration [e.g., 9], the Smithtown work [e.g., 15] and the

Computer-Based Simulation Games work [10] involve guidance provided by high level goals such as learning about a model or survival in a simulation environment. Leutner (1993) demonstrates the importance of students with prior domain instruction actively requesting help rather than help being provided in an unsolicited manner during their interaction with a simulation environment. Note that in contrast to other published results that consistently point towards problem solving as the most promising point on the Exploratory Learning Continuum, these results point in the opposite direction, towards a less strongly guided approach, although they do not explicitly evaluate these two approaches in comparison with problem solving. In this paper we empirically evaluate a new place on the continuum that we refer to as Negotiable Problem Solving Goals, which falls in between problem solving and the types of guided exploratory learning evaluated in the past [e.g., 9,15]. Our empirical investigation compares Negotiable Problem Solving Goals with two approaches that mix tutorial learning and problem solving. In all three conditions, students interact with a simulation environment.

Related to the distinction between “high level goals” and “low level goals” is the distinction between “learning oriented goals” and “performance oriented goals”, which is the second conceptualization of exploratory learning that we investigate in this paper [3,11]. Some have argued that the distinction is identical and that under specified goals are inherently more learning oriented and correspondingly more conducive to learning. Others have argued that learning orientation is more of a characteristic of the learner than the task, and that even in connection with the same goals provided to the learner, learners with different orientations will approach the problem differently, and that difference in orientation may be responsible for the contributing to or detracting from the depth with which the learner absorbs the material [11].

2. Method

We are conducting our research in the domain of thermodynamics, using as a foundation the CyclePad articulate simulator [4]. CyclePad offers students a rich, exploratory learning environment in which they apply their theoretical thermodynamics knowledge by constructing thermodynamic cycles and performing a wide range of efficiency analyses without expense or danger.

Materials. The domain specific materials used in the study, which consisted of a take-home assignment, pre/post test, introductory reading material about rankine cycles, and focused readings with suggested illustrative analyses to perform using the CyclePad simulator for three forms of rankine cycles, were all developed by a Carnegie Mellon University mechanical engineering professor with the help of three of his graduate students and minimal input from our team. These domain specific materials were exactly the same across conditions, with the exception of the manipulation specific instructions described below. Thus, we strictly controlled for information presentation in all written materials. Additionally, we used a questionnaire to assess student attitudes after their participation.

Experimental procedure common to all conditions. The study consisted of two labs involving work with CyclePad that were assigned to the whole class. The first lab was a self-paced take-home assignment done during the first week of the study. The second lab was a 3-hour on-campus lab session completed during the second week of the study. Although the labs were mandatory assignments, participation in the study was optional. We strictly controlled for time between conditions. The 3-hour lab session was divided into 8 segments: (1) After completing the consent form, students were given 20 minutes to work through a 50 point pre-test consisting of short answer and multiple choice questions

covering basic concepts related to rankine cycles, with a heavy emphasis on understanding dependencies between cycle parameters. (2) Students then spent 15 minutes reading an 11 page overview of basic concepts of rankine cycles. (3) Next they spent 25 minutes working through the first of three focused materials with readings, suggested problem solving goals, and analyses to help in meeting those goals. (4) Next they spent 20 minutes working through the second set of focused materials. (5) They then spent 20 minutes through the third set of focused materials. (6) They then spent 40 minutes in a Free Exploration phase creating the most efficient rankine cycle they could with no instructional support either from the tutor or any of the instructional materials they had been given previously. (7) They then spent 20 minutes taking a post-test that was identical to the pretest. (8) Finally, they filled out the questionnaire. The experimental manipulation took place during steps (3)-(5).

Experimental design. Our experimental manipulation consisted of 6 conditions resulting from a 3X2 full factorial design contrasting 3 goal level conditions and two goal orientation conditions. The three goal level conditions included (1) Negotiable Problem Solving Goals (NPSG), which was human tutoring support + written materials that we refer to as a Script, (2) Problem Solving (PS), which consisted of help provided in the style of typical model tracing tutors + Script, and (3) Script only (S). In the Human tutoring condition (NPSG), students are given the opportunity to take the most initiative. In that condition they are free to select problem solving goals from the list provided, with some guidance from the human tutors, and to select from a provided list of CyclePad analyses to meet those goals. In the problem solving condition (PS), students follow the list of provided problem solving goals in order, but they decide how best to meet those goals from the suggested analyses. In the script condition (S), students follow the list of provided problem solving goals and achieve them by following the list of suggested analyses in the specified order. However, the instructions are specified at a high enough level that some means-ends analysis is still required to successfully follow them. Our goal orientation manipulation was a replication of [11], and was completely determined by manipulation specific instructions, which are described below.

It is important to note that the superiority of the human tutoring based negotiable problem solving goals condition is not a foregone conclusion in the light of recent results in the tutorial dialogue community, and thus presents a valid test of our hypothesis about negotiable problem solving goals. Consider the following series of empirical investigations. First, two evaluations of the AutoTutor system, in the domains of computer literacy and physics, showed an advantage over re-reading of the textbook of about 0.5 standard deviations [12,7]. The textbook re-reading condition itself was no better than a no-treatment control condition. However, in a different experiment the learning results obtained with WHY-AutoTutor were no worse than *a human tutoring condition* and yet not better than those in a control condition in which students read targeted “mini-lessons,” short texts that covered the same content as that presented in the dialogue [6]. The mini-lesson condition is different from reading textbook text in that mini-lessons tend to be focused specifically on the knowledge and potential misconceptions involved in a specific exercise. It appears to be a high standard against which to compare. Note that the experimental procedure in our study involves extensive reading for students in all conditions. As a result, our experimental results can be seen as contributing to this line of investigating the trade-offs between human tutoring and a reading control. However, in order to place our experiment accurately in the context of previous results, it is important to consider the following differences. First, students in all conditions in our study were presented with exactly the same reading materials. Rather than replacing the reading

materials as in [6], the role of the human tutors in our study was to help students navigate and understand the materials. Secondly, the reading materials were neither as brief nor targeted to the test as the “minilessons” nor were they as extensive as a text-book.

Outcome Measures: We looked at three outcome measures of instructional effectiveness. Two outcome measures were assessed by means of a Pre/Post test. 32 multiple choice and short answer questions were used to test analytical knowledge of Rankine cycles, including relationships between cycle parameters. An important aspect of this was a set of prediction questions where students were told to predict the impact of a specific change in one cycle parameter on several other cycle parameters. The other part of the test was a set of 9 open response questions assessing conceptual understanding of Rankine cycles. The third type of outcome measure we looked at was ability to apply knowledge to build and optimize a Rankine cycle using CyclePad during a Free Exploration phase.

Participants. We conducted our study over a two week period of time as part of a sophomore Thermodynamics course at Carnegie Mellon University beginning the week when Rankine cycles were introduced in the lecture portion of their class. Each student in the two NPSG conditions (NPSG-LO and NPSG-PO) who completed the study was tutored by one of three mechanical engineering graduate students during an individual tutoring session. The students in the other 4 conditions (PS-LO, PS-PO, S-LO, and S-PO) completed their 3-hour lab in a group lab session that was specific to their condition. Students were assigned to conditions in such a way as to maximize the evenness in distribution of grade so far between conditions and to respect student availability during 4 lab session times, as indicated on an on-line questionnaire. The average grade so far in the class for each condition was virtually identical. However, only 67 out of 120 students both attempted the take home lab and participated in the experiment. An additional 30 students completed the second lab but did not do the take-home assignment.

Manipulation specific instructions. Prior to the second lab, students were either told they were assigned to a specific group lab time or that they were to make an appointment for an individual lab time, but they were not told prior to the second lab what type of instructional treatment to expect or how their treatment differed from that of other students. In between segments (1) and (2) and also between segments (2) and (3) of the experimental procedure, students in the Learning Orientation (LO) condition were told that their goal was to learn as much thermodynamics as possible during the lab, and that at the end they would be asked to demonstrate the deep understanding that they acquired. In contrast, students in the Performance Orientation (PO) condition were told that their goal was to achieve the greatest cycle efficiency as possible and that in the end they would be asked to demonstrate their ability to achieve the greatest efficiency possible. Additionally, in between segments (2) and (3) students received instructions specific to their goal level manipulation.

3. Results

First, we verified that the goal orientation manipulation had an effect on student goal orientation. We examined patterns of student responses on two goal orientation manipulation check questions on the questionnaire that were adapted from previous studies investigating student goal orientation [e.g., 3]. In both cases, students in the Learning Oriented (LO) condition were more likely to select the learning oriented response than students in the Performance Orientation condition (PO). We evaluated the reliability of the difference in proportion between conditions using a multinomial logistic regression. In the case of the first question, the difference was marginal $t=1.58$, $p=.11$. For the other question, the difference was

significant, $t=2.33$, $p<.05$. Thus, we concluded that the goal orientation manipulation had an effect on the student population, and if differences in goal orientation do have an impact on student behaviour and learning, we should be able to detect these differences between conditions by examining our outcome measures.

Since not all students who participated in the 3 hour lab completed the take home assignment, we checked to see whether not having completed the assignment had an effect on how successful students were in learning during the lab. There was no significant difference in grade so far in the course between the students who participated in the lab and those who did not. Students who did not do the take-home assignment were evenly distributed across conditions. On average, it was the best students in the class who did the take-home assignment: Mean(no) = 70.26, s.d.= 11.7, Mean(yes)=75.5, s.d.= 9.2, $t(95)=2.4$, $p<.05$. However, controlling for pretest score, there was no reliable difference in post test score between students who did the take-home assignment and those who did not using a 2-tailed paired t-test, $t(24)=1.12$, $p=.27$. Thus, we considered students who did not do the take-home assignment in our analysis of learning gains on the Pre/Post test but not on our assessment of performance with CyclePad during the Free Exploration phase.

We found that there were serious problems with one of our three tutors, namely Tutor 3. He was extremely terse and impatient with students. His transcripts contained almost no conceptual discussion, and in his impatience, he rarely let students complete their work. Instead, he tended to take over and do the lab for them through the VNC connection to their simulation interface. Students who worked with him learned much less than expected based on their pretest scores, as clearly demonstrated in Table 1. Thus, we left the data from the students that he tutored out of the learning gains analysis described below.

	Script (S)		Pseudotutor (PS)		3 Tutors (NPSG)		Tutor 1 (NPSG)		Tutor 2 (NPSG)		Tutor3 (NPSG)	
	Ave	S.Dev	Ave	S.Dev	Ave	S.Dev	Ave	S.Dev	Ave	S.Dev	Ave	S.Dev
Free Exploration Success	63%	n.a	58%	n.a	63%	n.a	100%	n.a	38%	n.a	0%	n.a
Total Test Resid	1%	9%	-2%	9%	2%	10%	3%	6%	3%	10%	-7%	9%
Conceptual Test Resid	3%	13%	-5%	13%	3%	16%	12%	3%	10%	3%	-18%	6%
Analytical Test Resid	-1%	11%	-1%	10%	5%	12%	8%	10%	4%	9%	3%	17%

Table 1 Overview of Outcome Measures from Goal Level Manipulation. Note that test results are reported in terms of residuals resulting from a regression between pretest and posttest score. In other words, this is the portion of the post test score that differs from what is expected purely based on pretest score. A positive value indicates how much higher the post-test score is over and above what is expected based on pre-test score, based on the pattern observed over the whole population. A negative value indicates how much lower the post-test score is than what is expected based on pre-test score.

Overall there was a main effect for the Goal Level manipulation ($F(2,83) = 3.81$, $p < .05$, $MSE = 20.9$), but no main effect for Goal Orientation manipulation or the interaction between the two. Overall the order was $PS < S < NPSG$. Using a Bonferroni post-hoc

analysis, we determined that the difference between NPSG and PS was significant ($p < .05$), whereas the difference between NPSG and S was marginal ($p=.11$). The difference between the S and PS was only a statistical trend. Despite our disappointment at having to drop the data from Tutor 3, we consider the stark difference in effectiveness between his tutoring and the other two tutors as an indication that it was the Goal Level manipulation and not just a “warm body” effect (i.e., that students just preferred working with a human tutor) that led to the significant main effect for the Goal Level manipulation.

Because of larger differences in standard deviation within sections on the test than overall, the differences between conditions were less clear within individual test sections. On the conceptual part of the test, there was a significant main effect for Goal Level manipulation but not Goal Orientation manipulation, and no interaction effect. Again the order was $PS < S < NPSG$. Using a Bonferroni post hoc analysis, we determined that both S and NPSG were significantly better than PS ($p < .05$), whereas the difference between NPSG and S was only a trend ($p=.16$). On the objective part of the test there was no main effect either for Goal Level manipulation or Goal Orientation manipulation. However there was a marginal crossover interaction $F(2,83) = 2.98, p=.06$. The crossover interaction was between the P and PS conditions where PS was better in the Performance orientation condition (PO), but S was better in the Learning Oriented Condition (LO).

We then evaluated student performance on the Free Exploration assessment. There we found no main effect for Goal Level manipulation or Goal Orientation manipulation overall, nor an interaction. However, we found a significant difference in effectiveness between tutors within the NPSG condition using a binomial logistic regression ($p < .005$). For Tutor 1, 100% of his students were able to successfully complete the Free Exploration portion of the assignment. For Tutor 2, only 36% of his students were able to complete it. For Tutor 3, whose data was thrown out of the learning gains analysis, 0% of his students were able to complete the free exploration portion of the lab. 58% of PS students and 63% of S students were able to complete it. Obviously, Tutor 1, as the best performing representative of the NPSG condition, was significantly more effective than the other tutors as well as the other Goal Level manipulations on this assessment.

Overall, we found significant Goal Level manipulation effects, with NPSG being the clearest win across the three outcome measures, especially Tutor 1, as displayed in Table 1. However, in contrast to findings in McNeil & Alibali (2000), we found very little evidence of any Goal Orientation effect.

4. Conclusions and Current Directions

The results of our empirical investigation offer strong support that a tutorial dialogue system based on the idea of Negotiable Problem Solving Goals for support in an exploratory learning environment is a promising new direction for the tutorial dialogue community. One common pattern that we have observed is that students start out with the idea that more sophisticated designs will be more efficient. Thus, students have a tendency to be drawn towards the more advanced portions of the design space before they are ready to fully understand how to use that sophistication to an efficiency advantage. When our tutors observe this behavior, they encourage students to keep it simple and direct them back to more basic design explorations until students demonstrate a solid understanding at that basic level. This high level structuring provides many of the advantages of previously explored problem solving conditions. Because of it, students are not hampered by their preconceptions that would have led them to spend their time in explorations that would

have been devoid of educational value. Yet, students in the NPSG condition take more initiative than in the S or PS conditions because they still have a hand in deciding how they will spend their exploratory time. We are currently conducting an in-depth corpus analysis to gain deeper insights into what lead to the differences in effectiveness between Tutors 1, 2, and 3 within the NPSG condition. We plan to use that analysis as the foundation for the CycleTalk tutorial dialogue system, which we are developing [14].

Acknowledgements

This project is supported by ONR Cognitive and Neural Sciences Division, Grant number N000140410107.

References

- [1] Aleven V., Koedinger, K. R., & Popescu, O. (2003). A Tutorial Dialogue System to Support Self-Explanation: Evaluation and Open Questions. *Proceedings of the 11th International Conference on Artificial Intelligence in Education, AI-ED 2003*.
- [2] Charney, D.H., & Reder, L.M. (1986). Designing tutorials for computer users: Effects of the form and spacing of practice on skill learning. *Human Computer Interaction, 2*, 297-317.
- [3] Dweck, C. S., & Leggett, E. L. (1988). A social-cognitive approach to motivation and personality. *Psychological Review, 95*, 256-273.
- [4] Forbus, K. D. (1999). CyclePad: An Articulate Virtual Laboratory for Engineering Thermodynamics. *Artificial Intelligence 114(1-2)*: 297-347.
- [5] Graesser, A. C., Bowers, C. A., Hacker, D.J., & Person, N. K. (1998). An anatomy of naturalistic tutoring. In K. Hogan & M. Pressley (Eds.), *Scaffolding of instruction*. Brookline Books.
- [6] Graesser, A., VanLehn, K., the TRG, & the NLT. (2002). *Why2 Report: Evaluation of Why/Atlas, Why/AutoTutor, and Accomplished Human Tutors on Learning Gains for Qualitative Physics Problems and Explanations*, LRDC Tech Report, (2002) University of Pittsburgh.
- [7] Graesser, A. C., Jackson, G. T., Mathews, E. C., Mitchell, H. H., Olney, A., Ventura, M., Chipman, P., Franceschetti, D., Hu, X., Louwerse, M. M., Person, N. K., and the Tutoring Research Group, (2003). Why/AutoTutor: A Test of Learning Gains from a Physics Tutor with Natural Language Dialog. *Proceedings of the Cognitive Science Society*.
- [8] Klahr & Nigam (2004). The equivalence of learning paths in early science instruction: effects of direct instruction and discovery learning, *Psychological Science*, 2004.
- [9] Leutner, D. (1993). Guided discovery learning with computer-based simulation games: effects of adaptive and non-adaptive instructional support. *Learning and Instruction, 3*, 113-132.
- [10] Mayer, R. E. (2004). Should there be a three-strikes rule against pure discovery learning? The Case for Guided Methods of Instruction, *American Psychologist 59(1)*, pp 14-19.
- [11] McNeil, N. M., & Alibali, M. W. (2000). Learning mathematics from procedural instruction: Externally imposed goals influence what is learned. *Journal of Educational Psychology, 92*, 734-744.
- [12] Person, N., Bautista, L., Graesser, A., Mathews, E., & The Tutoring Research Group (2001). In J. D. Moore, C. L. Redfield, & W. L. Johnson (Eds.), *Artificial Intelligence in Education: AI-ED in the Wired and Wireless Future, Proceedings of AI-ED 2001* (pp. 286-293). Amsterdam, IOS Press.
- [13] Rosé, C. P., Jordan, P., Ringenberg, M., Siler, S., VanLehn, K., & Weinstein, A. (2001). Interactive Conceptual Tutoring in Atlas-Andes, In J. D. Moore, C. L. Redfield, & W. L. Johnson (Eds.), *Artificial Intelligence in Education: AI-ED in the Wired and Wireless Future, Proceedings of AI-ED 2001* (pp. 256-266). Amsterdam, IOS Press.
- [14] Rosé, C. P., Torrey, C., Aleven, V., Robinson, A., Wu, C. & Forbus, K. (2004). CycleTalk: Towards a Dialogue Agent that Guides Design with an Articulate Simulator, *Proceedings of the Intelligent Tutoring Systems Conference*.
- [15] Shute, Valerie J. and Glaser, Robert, (1990). "A Large Scale Evaluation Of An Intelligent Discovery World: Smithtown", *Interactive Learning Environments 1*, 51-77.
- [16] Tuovinen, J. E., & Sweller, J. (1999). A comparison of cognitive load associated with discovery learning and worked examples. *Journal of Educational Psychology, 91(2)*, 334-341.