

Delivering Hints in a Dialogue-Based Intelligent Tutoring System

Yujian Zhou¹, Reva Freedman², Michael Glass¹,
Joel A. Michael³, Allen A. Rovick³, Martha W. Evens¹

¹Department of CSAM
Illinois Inst. of Technology
10 W. 31st Street 236-SB
Chicago, IL 60616

²LRDC #819
University of Pittsburgh
3939 O'Hara Street
Pittsburgh, PA 15260

³Department of Physiology
Rush Medical College
1750 W. Harrison Street
Chicago, IL 60612

zhouyuj@charlie.iit.edu, freedrk+@pitt.edu, glass@steve.iit.edu,
{jmichael, arovick}@rush.edu, csevens@minna.cns.iit.edu
<http://www.csam.iit.edu/~circsim>

Abstract

Hinting is an important tutoring tactic in one-on-one tutoring, used when the tutor needs to respond to an unexpected answer from the student. To issue a follow-up hint that is pedagogically helpful and conversationally smooth, the tutor needs to suit the hinting strategy to the student's need while making the strategy fit the high level tutoring plan and the tutoring context. This paper describes a study of the hinting strategies in a corpus of human tutoring transcripts and the implementation of these strategies in a dialogue-based intelligent tutoring system, CIRCSIM-Tutor v. 2. We isolated a set of hinting strategies from human tutoring transcripts. We describe our analysis of these strategies and a model for choosing among them based on domain knowledge, the type of error made by the student, the focus of the tutor's question, and the conversational history. We have tested our model with two classes totaling 74 medical students. Use of this extended model of hinting increases the percentage of questions that students are able to answer for themselves rather than needing to be told.

Introduction

Hinting is a general and effective tutoring tactic in one-on-one tutoring when the student has trouble solving a problem or answering a question. In many student-oriented tutoring systems, the machine tutor will give hints when the student asks for help, e.g. Andes (Gertner et al., 1998). In this tutoring setup, the central issue of hinting is to help the student recall the related domain rules or facts that the student may have trouble with. In a system where the tutor has control over the conversation and asks the questions, hinting is also a good strategy to help the student find the expected answer when the student gives an unexpected one.

But in this tutoring setup, not only the student's possible weakness but also the tutor's plan and the tutoring context are important for issuing hints. Since there may be more than one pedagogical plan for tutoring a domain concept, the hinting strategy is closely related to the tutoring method or tutoring plan, although the detailed content of each hint is closely related to the domain concept. So how to issue a follow-up hint which is helpful to the student, coordinated with the tutoring plan, and coherent in the dialogue context is an important issue in this tutoring setup.

In this paper we will address how to deliver hints by considering all of these factors in CIRCSIM-Tutor, a conversational intelligent tutoring system (ITS) which uses free-text input and output as its dialogue interface. To do so, we will first study human tutors' hinting strategies. Then we describe our attempts to implement these strategies in CIRCSIM-Tutor in order to dynamically deliver versatile hints in the tutoring dialogue.

Another purpose of this work is to find out how far the tutoring system can go and how effective it will be by depending mainly on hinting to help students find the expected answer after they have given an unexpected one. Other approaches to the problem of unexpected answers have been proposed. For example, the model used by Freedman and Evens (1996) and Kim et al. (1998) is based on schemata, and allows unlimited nested plans and plan updating during execution. This is a more sophisticated model but it cannot be implemented in the current version of CIRCSIM-Tutor, as the current planner does not support unlimited nested plans. The work described here will allow us to gain experience in hinting.

Background

The CIRCSIM-Tutor Project

CIRCSIM-Tutor is an intelligent tutoring system designed to help medical students understand the negative feedback system that controls blood pressure. CIRCSIM-Tutor tutors by having students solve problems. The system presents the

This work was supported by the Cognitive Science Program, Office of Naval Research under Grant No. N00014-94-1-0338, to Illinois Institute of Technology. The content does not reflect the position or policy of the government and no official endorsement should be inferred.

Copyright © 1999, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

student with a description of a physiological change and asks for predictions about the effect of that change on seven important physiological parameters. Then it conducts a dialogue with the student to correct the errors in the predictions.

The current working version is a revised version of CIRCSIM-Tutor v. 2, developed in Lisp by Woo and others (Woo et al., 1991). It has seven modules: instructional planner, student modeler, input understander, text generator, screen manager, problem solver, and knowledge base. The instructional planner includes two levels of planners, a lesson planner and a discourse planner. The lesson planner generates lesson goals and decomposes them into discourse tasks. The discourse planner is responsible for controlling interactions between the tutor and the student. Most discourse tasks are executed as questions. When the student gives an unexpected answer to a question, the planner adds tasks to the agenda to complete the original goal by other means—giving the answer or giving a hint.

Earlier Work on Hints in CIRCSIM-Tutor

Hume et al. (1996) studied the use of hints by experienced tutors in the hope of formulating a strategy for using hints in an ITS. They observed that human tutors frequently use hints as a pedagogical tactic. However, the theory of hints in their framework is too broad, as they identify most tutoring moves as hints in some contexts. Furthermore, these hints were defined without reference to the structure required by the CIRCSIM-Tutor planner.

The original CIRCSIM-Tutor v. 2 produced very general hints to ask about physiological variables missed by the student. But it failed to tailor the content to the student's previous answer, as it issued only fixed hints such as "Think about the value of <the desired physiological variable>." To improve the hinting capability in CIRCSIM-Tutor, we started by adding hints, tailored to the student answer, which were given in response to a partially correct answer (this is the main situation in which the tutor gives hints in the original CIRCSIM-Tutor). This improved version was used by 24 students from Rush Medical College in April 1998. After reading the log files from this experiment, we found that the new hints were effective but there were other kinds of student answers that the system failed to respond to with follow-up hints.

For this reason we broadened the hinting ability in CIRCSIM-Tutor to include responses to other categories of student answers. This new version was used by 50 students from Rush Medical College in November 1998. We will discuss the experimental results later in this paper. This improved version, which also has a new input understander (Glass, 1999), is robust and useful enough that our expert tutors from Rush Medical College believe that it can be used without supervision by medical students as a regular part of the curriculum.

Interviewing Human Tutors

In an interview with our expert tutors they identified two rules for how they give hints:

First give evoking terms or synonyms.

Otherwise try to give an intermediate step.

These two rules indicate that human tutors are trying to help the students think actively (by giving more evocative language) and also trying to help them think along the right chain of causal relationships by giving a small step along the causal path.

These two rules cover many of the cases of the human tutors' usage of hints, but they are too general and abstract to actually implement hinting in CIRCSIM-Tutor. So we analyzed transcripts of human-to-human tutoring sessions conducted by the expert tutors in order to identify the types of hints used in different situations and identify more specific strategies that could actually be used to build an ITS.

Hinting Strategies in Human Tutoring

Hints occur in many different surface forms. To implement them in a principled and useful way in a real ITS, we need to identify the underlying principles in order to avoid just giving canned hints in each situation. So we want to isolate hinting strategies that are not dependent on specific domain facts or rules.

The following strategies are some examples of hinting strategies frequently used in our human tutoring transcripts.

Strategy: Give an Intermediate Causal Link

This is one of the rules indicated by our human tutors. It actually has three sub-rules, each related to a different tutoring plan. Suppose there are several causally related physiological variables A affects X affects B, where the tutor usually teaches the relationship between A and B, ignoring intermediate steps like X.

- If the tutor asked which variable is affected by a change in A, then mentioning the link from A to X can be an effective hint toward the answer B.
- If the tutor asked which variables cause a change in B, then mentioning the link from X to B can be an effective hint backward toward the answer A.
- If the tutor asked how A and B are related, then mentioning either of the relationships from A to X or from X to B can be a hint.

By giving hints like this, the tutor offers a small piece of information relating the variable in question to the desired answer. The pedagogical expectation is that the student will think along these lines and find the desired answer.

Strategy: Refer to an Anatomy Object

Although our tutors prefer to focus on the physiology, they

occasionally point to an anatomy object to help the student concentrate on the right part of the cardiovascular system if the student can not answer a question. This kind of hint is especially useful when the student has trouble finding the first variable affected by a physiological change. For example:

T: What is the first variable affected by the alpha agonist?

S: I don't know.

T: Where in the CV [cardiovascular] system are alpha receptors found?

Strategy: Point Out the Laws of Physics Involved

Although our domain is physiology, it is occasionally useful for the tutor to point to some physics rules to help the student visualize why the causal relation should be the way it is. For example:

T: When MAP [mean arterial pressure] increases, what will that affect?

S: (incorrect answer)

T: When MAP increases, it's harder for the ventricle to pump blood.

Strategy: Give Evoking Terms or Synonyms

While most of the time our tutors use a specific set of physiology terms in order to encourage students to use the same terms, they sometimes choose more evocative phrases. For example, in certain contexts they often use "afterload" as a synonym for "mean arterial pressure," evoking images of the pressure the heart is pumping against. This strategy is used mostly when the tutor is tutoring the causal relationship from mean arterial pressure to stroke volume after an incorrect student answer.

Strategy: Linguistic Hint

Since our human tutors use natural language (just as our tutoring system does), they sometimes give subtle linguistic hints which include very little domain information. These hints are intended to help the student to think more actively.

A typical example occurs when the tutor is expecting several parameters from the student and the student gives only some of them. The tutor may simply reply with "And?" to indicate that more is expected.

Other Strategies

The above strategies are the most frequently used. There are also some other strategies that are used infrequently or are used only in special tutoring situations. These include pointing out the function of a drug, using capital letters to indicate a core variable, giving a definition, pointing out the problem-solving context, and referring to an equation.

Implementing Hinting in CIRCSIM-Tutor

Now that we have analyzed the hinting strategies of human tutors, we will discuss our implementation of these strategies in a running intelligent tutoring system—CIRCSIM-Tutor. Although there may be some deeper cognitive reasoning behind the human tutors' hinting strategy, we do not model such reasoning in the tutoring system. First, we lack a sufficiently comprehensive cognitive theory of hinting. Second, from a practical point of view, our analysis of human tutors' hinting strategies demonstrates that we can generate precise hints without invoking such a theory. The following algorithms describe our simulation of human tutors' hinting behavior.

Factors Determining the Hinting Strategies

There are several factors that may affect the choice of a specific hint: tutoring topic, tutoring context, tutoring history, student's answer, and so on. From the interview with human tutors and the study of their tutoring transcripts we find several factors to be particularly relevant.

First, to be pedagogically useful, a hint has to be related to the tutoring topic and be useful in helping the student find the expected answer. So the tutoring topic is important.

Second, the student's answer is important since hints are intended to help the student figure out the expected answer from what he or she has already said.

Third, the specific question used by the tutor, which is a reflection of the high level tutorial plan, is important because there may be several questions available for tutoring the same concept. Different kinds of tutor questions may indicate a different conversational context or focus.

Finally, the tutoring history is also important, especially for the second or third hint in a row. The tutor needs to base further hints on earlier ones for two reasons. From a discourse point of view, it increases the coherence of the conversation. From a pedagogical point of view, it makes the tutoring logic stand out more clearly.

A Classification Model for Student Answers

We added a classification module to categorize student answers. Below are the categories that we use to classify students' answers:

1. Correct
2. Partially correct answer, i.e. some part of the answer is correct and the rest is incorrect
3. Near miss answer, which is pedagogically useful but not the desired answer (Glass, 1997)
4. "I don't know" answer
5. "Grain of truth" answer, where the student gives an incorrect answer, but also indicates a partially correct understanding of the problem (Woolf, 1984)

6. Misconception, a common confusion or piece of false knowledge about the concept being tutored
7. Other incorrect answers
8. Mixed answers, i.e. a combination of answers from other categories

These categories, which were abstracted from our analysis of human tutoring transcripts, are one of the features used to decide which hint to give. The more information the tutor can find in the student's answer, the more specific the hint can be. Although the categorization is based on the domain model, it is important to recognize that it is largely a pragmatic categorization, i.e. a correct answer is one which our human tutors do not feel the need to correct or augment.

Use of a Simple Quantitative Student Model

Hume et al. (1996) observed that human tutors maintain a rough assessment of the student's performance. They argued that when and how to give hints is based on that measurement of student performance. Although our definition of hint is narrower than theirs, we still feel that student performance is a good criterion for deciding when to deliver hints rather than giving the answer. Thus we added a student performance module to CIRCSIM-Tutor's original student model. It includes four levels of measurement: global assessment (total measurement of the student so far), procedure-level assessment (measurement for each problem the student is asked to solve), stage assessment (measurement for each of the three physiological stages in a problem), and the local assessment attached to each variable that has been tutored. The local assessment is updated after each tutoring interaction and other assessments are calculated from the local assessment. If the student's performance is too low, the tutor gives the answer instead of issuing a hint.

This assessment model is based on intuitive rules and is still being refined. From experiments with medical students, we have found that we need other history data along with the assessment of the student for deciding between giving a hint and giving the answer, especially when the student gives the same wrong answer twice in a row.

Identifying the Possible Hinting Strategies

From our analysis of human tutoring transcripts, we abstracted a set of hinting strategies, detailed below. We then built a hinting algorithm for each category of student answer. Each answer category is associated with a predefined list of strategies. Some of the algorithms are quite simple. For example, if the student gives a near miss answer, the tutor responds with a leading question that points to an intermediate link from the near miss to the correct answer. Some of the algorithms are more complex. For example, if the student's answer is incorrect, there are

several strategies available. If the tutor is tutoring a causal link in the forward direction, the hinting strategies mostly prefer to give evoking terms related to the variable already mentioned or to give an intermediate link in the forward direction.

Using Heuristic Rules to Rank the Strategies

If the tutor still has several strategies to choose among, the tutor ranks the possible hints using heuristic rules which attempt to generate more specific hints first. We consider hints in the following order:

1. Hints that are specifically related to the student's answer
2. Hints involving equations
3. Hints involving evocative synonyms
4. Hints involving a related anatomical object
5. Hints that give an intermediate logical step
6. Others

Locating Appropriate Content

The result of this procedure is a list of hint types only. To decide the details of the content in a hint, the tutor searches the domain knowledge base for each of the available hinting strategies. For the first hinting strategy, it looks to see if the knowledge base has an entry for the concept currently being tutored. If the domain knowledge has an entry, then the search terminates; if not, the algorithm tries the remaining hinting strategies in sequence. If no entry is available for any of the possible strategies, the tutor gives a default hint.

To support the hinting strategies that we identified from the human tutoring sessions, we are in the process of extending our domain knowledge base with additional evocative terms, related anatomical objects, and related physics rules.

Using Templates to Deliver Hints

We use hint templates determined by the content selection algorithms discussed above to deliver the hints. For example:

Like <the related variable in the student's answer>, the other determinant is also related to <the anatomy object>.

But what I am asking is <definition of the tutored variable>.

Do you remember an equation written as: <the variable been tutored> = ...?

The Model in Practice

Derivation of a Hint

The following example illustrates how the machine tutor

determines the follow-up hint step by step after the student gives an unexpected answer. Suppose the recent dialogue is:

T: What determines CVP [central venous pressure]?

S: I don't know.

Here the category of the student's answer is "I don't know." If the tutor decides to give hints, the hint algorithm related to the "I don't know" answer will be evoked. Since there are several strategies available, the algorithm will first check the tutoring plan. Here it is trying to tutor a causal relation backward. So the possible hinting strategies reduce to:

- Find equations related to the dependent variable.
- Point to a feature of an anatomy object related to the dependent variable.
- Give an intermediate step backward from the dependent variable.

Using the heuristic rules above, this list is the final list after ranking the preferences. Then the tutor checks the domain knowledge base and finds that the second strategy has suitable content available. Finally the tutor will find the related hint template and deliver the hint as:

T: Remember the CVC [central venous compartment] is very compliant. So what determines CVP?

If the student still can not get the correct answer, the tutor could issue a further hint giving an intermediate step between CVC and CVP. But if the student gives a near miss answer, e.g. CBV (central blood volume) instead, the tutor could use the near miss as a basis for issuing a follow-up hint instead. That hint might be expressed as a question pointing to an intermediate link between the near miss and the desired answer:

S: How about CBV?

T: What determines CBV?

Influence of the Tutoring Question

The wording or intent of the tutor's questions is not an issue in systems where the student asks the questions. But the opposite is true in CIRCSIM-Tutor, particularly when the question indicates the direction the tutor is following along a causal link. When the tutor is teaching the relationship between A and B, different questions can address the same causal link in different directions. In response to an incorrect answer, each question might benefit from a different hint. In Dialogue 1 below, the tutor tries to teach about the link from cardiac output to central venous pressure, working backward from CVP. In Dialogue 2, the tutor works forward from CO. Although the student gives an incorrect response in both cases, the resulting hint is different.

Dialogue 1:

T: What determines CVP?

S: (*Incorrect answer.*)

T: Remember the CVC is very compliant. (*Looks backwards from CVP to the volume of blood in the CVC.*)

Dialogue 2:

T: When Cardiac Output decreased, how would that affect the value of Central Venous Pressure?

S: Decrease. (*Incorrect answer.*)

T: No. When CO decreases, think about what will happen to the volume of blood in the central venous compartment. (*Gives a forward hint from CO to the volume of blood in the CVC.*)

Influence of Tutoring History

The tutor keeps a tutoring history to avoid repeating the same hint and to make sure that the hints do not return to a causal relationship that was already tutored if there are several hints in a row.

T: What are the determinants of Stroke Volume?

S: It is RAP [right atrial pressure]. (*Near miss.*)

T: And what determines RAP? (*Hint based on the near miss.*)

S: EDP [end diastolic pressure]. (*A near miss, but further from the correct answer than RAP.*)

T: EDP comes from RAP. Think again, what parameter in the prediction table determines RAP? (*Point out the relation between RAP and EDP, then give further hint based on RAP.*)

If we did not keep a dialogue history, the tutor could ask "And what determines EDP?" In that case the tutor would be following the student's lead away from the right answer.

Sample Output

Although our model of hinting is based on a number of simple algorithms, it can generate a variety of hints that allow the students to answer for themselves rather than being told. Below is an example produced by CIRCSIM-Tutor using our model.

T: What are the determinants of Stroke Volume?

S: Heart Rate and Cardiac Output. (*"Grain of truth" answer, i.e. wrong but pedagogically useful.*)

T: Well, it is true that $CO = SV * HR$. But what I was asking is what determines how much blood is ejected from the heart each time it beats (the SV)? (*Hinting strategy: give definition.*)

S: Preload. (*Near miss.*)

T: Right. And what represents preload in the prediction table? (*Hinting strategy: point to the intermediate link from the near miss. Note that the tutor gave a positive acknowledgment ("right") to the near miss, showing that there is not always a one-to-one relationship between student evaluation and the correct discourse move.*)

S: Central Venous Pressure. (*Correct, but still missing one determinant for the original question, so we*

categorize it as a partial answer.)

T: Right. Central Venous Pressure determines SV. The other determinant of SV is neurally controlled. What is the other determinant of SV? (*Hinting strategy: point to the domain type of the missing variable.*)

S: It is Inotropic State. (*Correct.*)

T: Right. Inotropic State is the other determinant of SV. SV is determined by Central Venous Pressure and Inotropic State. (*The goal of the original question has been satisfied without telling the student the correct answer. The tutor can move to the next item on the agenda now.*)

Comparison to Related Work

There are several tutoring systems that use hints as a tutoring tactic. Andes (Gertner et al., 1998) generates individual hints. It uses a Bayesian-network based student model to tailor its follow-up hints to the student's knowledge, and delivers them by using an associated sequence of hint templates for each goal and fact in its knowledge base. The Lisp tutor (Anderson et al., 1995) also generates hints from a sequence of hint templates. It uses model tracing techniques to detect that the student is not following the correct solution path. Sherlock II (Lesgold et al., 1992) generates a paragraph after the conclusion of the tutoring session.

In CIRCSIM-Tutor, we use heuristic rules to choose a hinting strategy based on the category of the student's answer, the tutorial plan, and the tutoring history. We then decide the content by searching the domain knowledge base to instantiate the strategy. So our hints are focused on both the student's needs and the current tutorial plan. By considering the current tutorial plan, the tutor can make sure the hints are coordinated with the tutorial plan and ensure conversational coherence while at the same time tailoring the content of the hint to the student's needs.

Merrill et al. (1992) compared the effectiveness of human tutors and intelligent tutoring systems. Their study indicated that a major reason that human tutors are more effective is that they let the students do most of the work in overcoming impasses, while at the same time providing as much assistance as necessary. Although in CIRCSIM-Tutor the tutor mainly leads the students in correcting the errors they have made in the prediction table, it is also important to let the student do as much as possible. By giving follow-up hints tailored to the student's answer rather than giving the correct answer, the tutor provides necessary guidance to the student while promoting a more active style of learning.

In CIRCSIM-Tutor sometimes the student model can recognize the specific confusion of the student through its categorization of the student's answer. In that case, the hint is specifically related to the student's knowledge state. But

even if the student model can not infer a deep understanding of the student's mental model, hinting is still more useful than just giving the answer for two reasons. In addition to giving the student a second chance to correct the error, the content of the hint may offer the student useful information for understanding the material.

Evaluation

An earlier version of CIRCSIM-Tutor which implemented a portion of the hinting model described above was used by 50 first-year students from Rush Medical College in November 1998. All of the students had already completed the regular lectures. They used the program for one hour. Twenty-four students worked in pairs at a computer and 26 students worked alone. We obtained a log file from each student or pair of students, giving a total of 38 log files.

The tables below describe our initial formative evaluation of this portion of the hinting model. In this experiment, CIRCSIM-Tutor asked approximately 1700 questions. In the course of tutoring 565 student errors, it generated 97 hints. Table 1 shows the effectiveness of hints for different student answer categories and Table 2 shows the effectiveness of each hinting strategy.

Category of answer	No. of hints	No. of correct answers	% of correct ans.
Partially correct	55	41	75%
Near miss	12	9	75%
Incorrect	14	14	100%
Mixed	16	14	88%

Table 1: Hints used by answer category

Category of hinting strategy	No. of hints	No. of correct ans.	% of correct ans.
Involving equations	2	2	100%
Evocative language	17	10	59%
Point to anatomical object	4	3	75%
Intermediate step	19	16	84%
Point to variable type	31	25	81%
Others	24	22	92%

Table 2: Effectiveness of hints by strategy

In evaluating this performance it must be noted that in this experiment CIRCSIM-Tutor did not have a hint to give in all situations. In particular, hints for the incorrect answer category tended to occur on questions which had only a few possible answers, such as yes/no questions. Additionally, we believe that these questions were among the easier ones. As a result these hints tended to produce

good results. Hints for the near miss and partially correct answers were more likely to come from questions with a larger range of possible responses.

We have now implemented most of the possible hinting strategies for each answer category, and we hope to evaluate these hints in a later experiment. We are also looking forward to comparing the learning results between students who use the system with hints and without. Additionally, we are in the process of analyzing experimental data that will allow us to do a detailed analysis of student learning by comparing pretest and posttest results.

Another possible method for evaluating hints would be to let our human tutors compare the hints generated by CIRCSIM-Tutor to what they would like to say in the same situation. This method was used during the initial development of the system. We believe that this method of evaluation is important since the goal of this project is to simulate human tutoring behavior as closely as possible. Currently one of our expert tutors is working with the latest version of CIRCSIM-Tutor with this goal in mind.

The students were also positive about the quality of the hints and explanations (1.90 on a scale from 1= definitely YES to 5 = definitely NO, computed from the experiment survey form).

Conclusions and Future Work

In this paper we addressed how to systematically deliver pedagogically sound and conversationally coherent hints in a dialogue-based ITS, CIRCSIM-Tutor. Our strategy involved categorizing student answers, and considering both tutoring plan and dialogue history. We first studied human tutoring transcripts to identify human tutors' hinting strategies and factors that might affect their choice of a hinting strategy. We then implemented these strategies in a real tutoring system as much as possible.

During the spring semester of 1999, CIRCSIM-Tutor will be installed as a standard program at Rush Medical College to be used by any student who wants. We plan to analyze the log files to see how effective the new hints are. We will also analyze the inappropriate hints and discuss with our expert tutors how to fix them.

It is worthwhile to note that we are also planning to replace CIRCSIM-Tutor v. 2 by a completely rewritten v. 3 based on the work of Freedman and Evens (1996). That project, currently in progress, will allow us to add more complex kinds of remediation since we will be able to use nested plans and delete agenda items that have become irrelevant. We are looking forward to identifying uses for these new features. However, since the hinting algorithms described here are based on an actual corpus of tutoring transcripts, they will remain pedagogically valid and we intend to re-implement them in the new system.

Since most of the strategies isolated from the tutoring

transcripts are not related to specific domain knowledge, we also expect them to generalize to other causal domains.

References

- Anderson, J., Corbett, A., Koedinger, K., and Pelletier, R. (1995). Cognitive Tutors: Lessons Learned. *Journal of the Learning Sciences* 4(2): 167–207.
- Freedman, R. and Evens, M. (1996). Generating and Revising Hierarchical Multi-turn Text Plans in an ITS. *Intelligent Tutoring Systems: Third International Conference (ITS '96)*, Montreal, 632–640. (Springer-Verlag Lecture Notes in Computer Science, 1086.) Berlin: Springer.
- Gertner, A., Conati, C., and VanLehn, K. (1998). Procedural Help in Andes: Generating Hints using a Bayesian Network Student Model. *Proceedings of the Fifteenth National Conference on Artificial Intelligence*, Madison, 106–111. Menlo Park: AAAI Press.
- Glass, M. (1997). Some Phenomena Handled by the CIRCSIM-Tutor Version 3 Input Understander. *Proceedings of the Tenth Florida Artificial Intelligence Research Symposium*, Daytona Beach, 21–25.
- Glass, M. (1999). Broadening Input Understanding in an Intelligent Tutoring System. Ph.D. diss., Dept. of CSAM, Illinois Institute of Technology.
- Hume, G., Michael, J., Rovick, A., and Evens, M. (1996). Hinting as a Tactic in One-on-One Tutoring. *Journal of the Learning Sciences* 5(1): 32–47.
- Kim, J., Freedman, R., and Evens, M. (1998). Responding to Unexpected Student Utterances in CIRCSIM-Tutor v. 3: Analysis of Transcripts. *Proceedings of the Eleventh Florida Artificial Intelligence Research Symposium (FLAIRS '98)*, Sanibel Island, 153–157. Menlo Park: AAAI Press.
- Lesgold, A., Katz, S., Greenberg, L., Hughes, E., and Eggan, G. (1992). Extensions of Intelligent Tutoring Paradigms to Support Collaborative Learning. In Dijkstra, S., Krammer, H., and van Merriënboer, J., eds., *Instructional Models in Computer-based Learning Environments*, 291–311. (NATO ASI Series, series F: Computer and System Sciences, 104.) Berlin: Springer.
- Merrill, D., Reiser, B., Ranney, M., and Trafton, J. (1992). Effective Tutoring Techniques: A Comparison of Human Tutors and Intelligent Tutoring Systems. *Journal of the Learning Sciences* 2(3): 277–305.
- Woo, C., Evens, M., Michael, J., Rovick, A. (1991). Dynamic Instructional Planning for an Intelligent Physiology Tutoring System. *Proceedings of the Fourth Annual IEEE Computer-Based Medical Systems Symposium*, Baltimore, 226–233. Los Alamitos: IEEE Computer Society Press.
- Wolf, B. (1984). Context-Dependent Planning in a Machine Tutor. Ph.D. diss., Dept. of Computer and Information Science, University of Massachusetts at Amherst. COINS Technical Report 84–21.