# Dialog Move Generation and Conversation Management in AutoTutor

**Natalie K. Person[1], Arthur C. Graesser[2], Derek Harter[2], Eric Mathews[1], and the Tutoring Research Group[2]**

[1]Department of Psychology, Rhodes College, 2000 N. Parkway, Memphis, TN 38112
person@rhodes.edu, matec@rhodes.edu
[2]Department of Psychology, University of Memphis, Memphis, TN 38152
a-graesser@memphis.edu, dharter@memphis.edu

## Abstract

AutoTutor is an automated computer literacy tutor that participates in a conversation with the student. AutoTutor simulates the discourse patterns and pedagogical dialog moves of human tutors. This paper describes how the Dialog Advancer Network (DAN) manages AutoTutor's conversations and how AutoTutor generates pedagogically effective dialog moves that are sensitive to the quality and nature of the learner's dialog contributions. Two versions of AutoTutor are discussed. AutoTutor-1 simulates the dialog moves of normal, untrained human tutors, whereas AutoTutor-2 simulates dialog moves that are motivated by more ideal tutoring strategies.

## Background

Human one-to-one tutoring is second to no other instructional method in yielding positive student learning gains. This particular claim has been supported in numerous research studies and is not particularly controversial. However, when the tutors of typical tutoring situations are considered, this claim becomes somewhat perplexing. Most tutors in school settings are older students, parent volunteers, or teachers' aides that possess some knowledge about particular topic domains and virtually no knowledge about expert tutoring techniques. Given their limited knowledge, it is somewhat impressive that these untrained tutors are responsible for the considerable learning gains that have been reported in the tutoring literature. Effect sizes ranging from .5 to 2.3 standard deviations have been reported for untrained tutors versus other comparable learning conditions (Bloom, 1984; Cohen, Kulik, & Kulik, 1982).

In order to identify the mechanisms that produce such positive learning gains, several members of the Tutoring Research Group (TRG) extensively analyzed a large corpus of tutoring interactions that occurred between untrained human tutors and students (Graesser & Person, 1994; Graesser, Person, & Magliano, 1995; Person & Graesser, 1999; Person, Graesser, Magliano, & Kreuz, 1994; Person, Kreuz, Zwaan, & Graesser, 1995). One reoccurring finding in many of our analyses is that untrained, human tutors rarely adhere to sophisticated or ideal tutoring models (e.g., Socratic tutoring, reciprocal teaching, anchored learning) that have been advocated by education and ITS researchers. Instead, untrained human tutors tend to rely on pedagogically effective dialog moves that are embedded within the conversational turns of the tutorial dialog. More specifically, human tutors generate dialog moves that are sensitive to the quality and quantity of the preceding student turn. The tutor dialog move categories that we identified in human tutoring sessions are provided below.

(1) <u>Positive immediate feedback</u>. "That's right" "Yeah"
(2) <u>Neutral immediate feedback</u>. "Okay" "Uh-huh"
(3) <u>Negative immediate feedback</u>. "Not quite" "No"
(4) <u>Pumping</u> for more information. "Uh-huh" "What else?"
(5) <u>Prompting</u> for specific information. "The primary memories of the CPU are ROM and _____?"
(6) <u>Hinting</u>. "What about the hard disk?"
(7) <u>Elaborating</u>. "CD-ROM is another storage medium."
(8) <u>Splicing</u> in the correct content after a student error.
(9) <u>Summarizing</u>. "So to recap," <succinct recap of answer to question>

After spending nearly a decade reading thousands of pages of tutoring transcripts and viewing hundreds of hours of videotaped tutoring sessions, we decided to it was time to put our knowledge to good use and build something. We built AutoTutor.

## What is AutoTutor?

AutoTutor is an animated pedagogical agent that engages in a conversation with the learner while simulating the dialog moves of untrained human tutors. AutoTutor is currently designed to help college students learn about topics that are typically covered in an introductory computer literacy course (e.g., hardware, operating systems, and the Internet). AutoTutor's architecture is comprised of five major modules: (1) an animated agent, (2) a curriculum script, (3) language analyzers, (4) latent semantic analysis (LSA), and (5) a dialog move generator.

All of these modules have been discussed rather extensively in previous publications; and therefore, will only be mentioned briefly in this paper. (see Foltz, 1996; Graesser, Franklin, Wiemer-Hastings, & the TRG, 1998; Graesser, Wiemer-Hastings, Wiemer-Hastings, Harter, Person, & the TRG, in press; Hu, Graesser, & the TRG, 1998; Landauer & Dumais, 1997; McCauley, Gholson, Hu, Graesser, & the TRG, 1998; Olde, Hoeffner, Chipman, Graesser, & the TRG, 1999; Person, Graesser, Kreuz, Pomeroy, & the TRG, 2000; Person, Klettke, Link, Kreuz, & the TRG, 1999; Wiemer-Hastings, Graesser, Harter, & the TRG, 1998; Wiemer-Hastings, Wiemer-Hastings, & Graesser, 1999).

AutoTutor is controlled by Microsoft Agent. AutoTutor's images were first created in MetaCreations Poser 3 and then loaded into Microsoft Agent. He is a three-dimensional embodied agent that remains on the screen during the entire tutoring session. AutoTutor's dialog moves are synchronized with head movements, facial expressions, and hand gestures that serve both conversational and pedagogical functions (Person, Craig, Price, Hu, Gholson, Graesser, & the TRG, 2000; Person et al., 1999).

AutoTutor begins the tutoring session with a brief introduction and then asks the student a question from the curriculum script. A curriculum script is a loosely ordered set of skills, concepts, example problems, and question-answer units. Most human tutors follow a script-like macrostructure, but briefly deviate from the structure when the student manifests difficulties, misconceptions, and errors. The content of the curriculum script in tutoring (compared with classrooms) has more deep reasoning questions (e.g., why, how, what-if, what-if-not), more problems to solve, and more examples (Graesser et al., 1995; Person & Graesser, 1999). AutoTutor has a curriculum script that organizes the topics (i.e., the content) of the tutorial dialog. The content of the curriculum script is represented as word phrases, sentences, or paragraphs in a free text format. The script includes didactic descriptions, tutor-posed questions, example problems, figures, diagrams, and simple animations.

There are 36 tutoring topics in the curriculum script. Each tutoring topic contains the following:

- a focal question (or problem)
- a set of good answer aspects that are included in the ideal answer (each aspect is roughly a sentence of 10-20 words)
- a set of tutor dialog moves that express or elicit each ideal answer aspect (i.e., hints, prompts, and elaborations)
- a set of anticipated bad answers (i.e., bugs, misconceptions)
- corrections for each bad answer (i.e., splices)
- a set of basic noun-like concepts

- a set of anticipated student questions with corresponding answers
- a summary of the ideal answer or solution

For each topic in the curriculum script, didactic information is presented prior to a focal question (e.g., "How does the operating system of a typical computer process several jobs simultaneously with only one CPU?"). The answers to focal questions are lengthy in nature and contain several good answer aspects. There are approximately five good answer aspects for each focal question in the curriculum script. During the conversation for a particular topic, the student's contributions are constantly compared to each of the good answer aspects that correspond to the focal question. Latent semantic analysis is used to assess the match between student contributions and each good answer aspect. AutoTutor attempts to extract the desired information from the student by generating prompts and hints for each good answer aspect.

Students respond to AutoTutor by typing contributions on the keyboard and hitting the "Enter" key. A number of language analyzers operate on the words in the student's contribution. These analyzers include a word and punctuation segmenter, a syntactic class tagger, and a speech act classifier. The speech act classifier assigns the student's input into one of five speech act categories: Assertion, WH-question, Yes/No question, Frozen Expression, or Prompt Completion. These speech act categories enable AutoTutor to sustain mixed-initiative dialog as well as dictate the legal DAN pathways that AutoTutor may pursue. The DAN pathways will be discussed in the section on Conversation Management.

AutoTutor's knowledge about computer literacy is represented by Latent Semantic Analysis (LSA) (Foltz, 1996; Foltz, Britt, & Perfetti, 1996; Landauer & Dumais, 1997; Landauer, Foltz, & Laham, 1998). LSA is a statistical technique that measures the conceptual similarity of two text sources. LSA computes a geometric cosine (ranging from 0 to 1) that represents the conceptual similarity between the two text sources. In AutoTutor, LSA is used to assess the quality of student Assertions and to monitor other informative parameters such as Topic Coverage and Student Ability Level. Student Assertion Quality is measured by comparing each Assertion against two other computer literacy text sources, one that contains the good answer aspects for the topic being discussed and one that contains the anticipated bad answers. The higher of the two geometric cosines is considered the best conceptual match, and therefore, determines how AutoTutor responds to the student Assertion. For the domain of computer literacy, we have found our application of LSA to be quite accurate in evaluating the quality of learner Assertions (Graesser et al., in press; Wiemer-Hastings et al., 1999).

We currently have two versions of AutoTutor (i.e., AutoTutor-1 and AutoTutor-2). The versions differ in terms of the mechanisms that control their respective dialog move generators. AutoTutor-1 simulates the dialog moves of normal, untrained human tutors via production rules, whereas AutoTutor-2 uses production rules and particular discourse patterns to simulate more ideal tutoring strategies. The dialog move mechanisms for both AutoTutor versions are discussed in greater detail in Dialog Move Generation section of this paper.

For each of the 36 questions and problems, the student and AutoTutor collaboratively improve the quality of the student's contributions while participating in a conversation. Thus, AutoTutor is more than a mere information delivery system. Once AutoTutor is "convinced" that enough of the critical components for a particular question/ problem have been covered, the conversation proceeds to the next question or problem in the curriculum script.

## Conversation Management

### Pre-DAN

Our goal from the onset of the AutoTutor project has been to develop an agent that simulates the dialog moves of untrained tutors while participating in a conversation with the learner. Needless to say, we have encountered a number of problems along the way. After all, having knowledge about when AutoTutor should generate particular dialog moves does not guarantee that a smooth conversation will transpire. In fact, quite the opposite occurred in some of AutoTutor's initial tutoring sessions with students. Consider the following dialog excerpts in which AutoTutor and the students are discussing the parts and uses of a computer. AutoTutor's dialog moves are labeled in parentheses and the students' utterances that were directed to the human lab assistant are included in brackets.

**Excerpt 1**
1.1 AutoTutor (Prompt): See if you can get this. Computers can be used to analyze data by finding averages, standard deviations, or performing many other kinds of . . .?
1.2 Student to lab assistant: [Oh, just answer the question? Um. So, am I still supposed to answer that question (pointing to the main topic question that remains on the monitor)?]
1.3 Student types: i don't know
1.4 AutoTutor (Prompt Response): kinds of statistics.
1.5 Student to lab assistant: [Okay, this is really confusing. I don't understand what the question is supposed to be.]
1.6 Student types: i don't know

**Excerpt 2**
2.1 AutoTutor (Elaboration): Also, a computer has a monitor.
2.2 Student to lab assistant: [I just answer yes or no?]
2.3 Student types: yes

These two excerpts highlight several of AutoTutor's original conversational deficits. One glaring deficit involves the students' uncertainty about when it is their turn to respond to AutoTutor. Turn-taking is an integral feature of the conversational process. To facilitate the turn-taking process in human-to-human conversations, speakers signal to listeners that they are relinquishing the floor (i.e., it is the listener's turn to say something) (Clark & Schaefer, 1987; Grice, 1975; Hobbs, 1979; McLaughlin, 1984; Nofsinger, 1991; Sacks, Schegloff, & Jefferson, 1978). However, human-to-computer conversations lack many of the subtle signals inherent to human conversations. When conversational agents like AutoTutor lack turn-taking signals, computer users (in our case, students) often do not know when or if they are supposed to respond. In conversations with AutoTutor, students were frequently confused after AutoTutor's Elaborations, Prompt Responses, and assertion-form Hints (some Hints were in question-form and were not problematic for students).

Another obvious deficit is that AutoTutor's dialog moves are not well adapted to the students' turns. For example, in Excerpt 1, AutoTutor's dialog moves are clearly not sensitive to the content of the student's turns. Participants engaged in human-to-human conversations, however, are able to adapt each conversational turn so that it relates in some way to the turn of the previous speaker. This micro-adaptation process is somewhat problematic for AutoTutor because the content of AutoTutor's dialog moves is predetermined. That is, AutoTutor doesn't generate the content of his dialog moves on the fly but rather selects each dialog move from a scripted set of moves that is related to the tutoring topic being discussed. Hence, we recognized early on that AutoTutor needed a mechanism that would allow him to make quasi-customized dialog moves given his limited number of dialog move options.

### Post-DAN

In order to rectify many of AutoTutor's turn-taking and micro-adaptation problems, we created the Dialog Advancer Network (DAN) (Person, Bautista, Kreuz, Graesser, & the TRG, in press; Person, Graesser, & the TRG, in press). The DAN contains AutoTutor's dialog move options (i.e., 78 dialog pathways) for any given student turn category. Example DAN pathways are provided in Figure 1. The DAN has improved AutoTutor's micro-adaptation capabilities by providing customized pathways that are tailored to particular student turn

categories. For example, if a student wants AutoTutor to repeat the last dialog move, the DAN contains a Frozen Expression pathway that allows AutoTutor to adapt to the student's request and respond appropriately. A DAN pathway may include one or a combination of the following components: (1) discourse markers (e.g., "Okay" or "Moving on"), (2) AutoTutor dialog moves (e.g., Positive Feedback, Pump, or Elaboration), (3) answers to WH- or Yes/No questions, or (4) canned expressions (e.g., "That's a good question, but I can't answer that right now").

The DAN also solved practically all of AutoTutor's previous turn-taking problems. Most of the turn-taking confusion was eliminated by the Advancer States that occur in many of the DAN pathways. Advancer States are designed to advance the conversational dialog and to disambiguate who has the floor in the conversation. For example, prior to implementing the DAN, the conversation often stopped after AutoTutor delivered an Elaboration, Hint, or Prompt Response dialog move because the student did not know what to do. In the current version of AutoTutor, students are no longer confused about whether they should respond because each of these dialog moves is followed by an Advancer State that requires AutoTutor to keep the floor. For example, AutoTutor keeps the floor after an Elaboration by articulating a predetermined discourse marker (e.g., "Moving on") and selecting another dialog move (see Figure 1).

Two conversations that are managed by the DAN are provided in Excerpts 3 and 4. These excerpts illustrate how the Advancer States advance the dialog and indicate to the student that AutoTutor is keeping the floor in the conversation. In Excerpt 3, AutoTutor and the student are discussing the advantages of RISC processors; in Excerpt 4 they are discussing how typed information gets passed from the keyboard to the hard disk.

### Excerpt 3
3.1 AutoTutor (Prompt): Now, what kind of architecture allows a RISC processor to execute multiple instructions simultaneously?
3.2 Student types: simple
3.3 AutoTutor (Prompt Response): A superscalar architecture.
3.4 AutoTutor (Dialog Advancer): Okay,
3.5 AutoTutor (Elaborate): Also, RISC processors have fewer bottlenecks that complicate processing.
3.6 AutoTutor (Dialog Advancer) Okay,
3.7 AutoTutor (Summary): Now to wrap up, RISC processors are faster and more efficient than CISC processors (AutoTutor continues to summarize this topic).

### Excerpt 4
4.1 AutoTutor (Prompt): Well, where is most of the information you type in temporarily stored?
4.2 Student types: RAM
4.3 AutoTutor (Positive Feedback): Right!
4.4 AutoTutor (Prompt Response): In RAM,
4.5 AutoTutor (Dialog Advancer): Moving on,
4.6 AutoTutor (Elaborate): The CPU handles one instruction at a time.
4.7 AutoTutor (Dialog Advancer): Okay,
4.8 AutoTutor (Summary): Let's review, after you enter information, it is sent to the CPU. The CPU carries out the instructions on the data. (AutoTutor continues to summarize this topic).
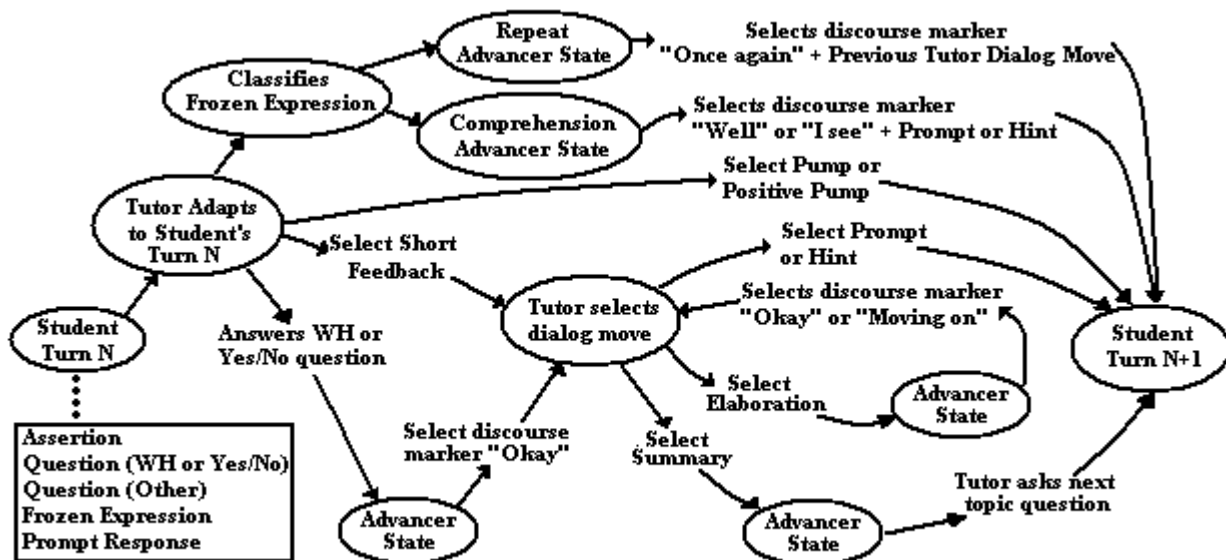


Figure 1. Dialog Advancer Network (DAN)

Since the implementation of the DAN, AutoTutor's interactions with students have improved considerably. The numerous pathways within the DAN have refined AutoTutor's micro-adaptation skills and the DAN Advancer States have eradicated much of the turn-taking confusion. Although the DAN is a relatively new feature of AutoTutor, it has already proven to be quite instrumental in helping us improve AutoTutor's overall effectiveness as a tutor and as a conversational partner (Person, Bautista, et al., in press).

## Dialog Move Generation

As previously mentioned, there are two versions of AutoTutor. AutoTutor-1 simulates the dialog moves of normal, untrained human tutors, whereas AutoTutor-2 simulates dialog moves that are motivated by more sophisticated, ideal tutoring strategies. Our analyses of human tutoring sessions revealed that normal, untrained tutors do not use most of the ideal tutoring strategies that have been identified in education and the intelligent tutoring system enterprise. These strategies include the Socratic method (Collins, 1985), modeling-scaffolding-fading (Collins, Brown, & Newman, 1989), reciprocal training (Palinscar & Brown, 1984), anchored situated learning (Bransford, Goldman, & Vye, 1991), error identification and correction (Anderson, Corbett, Koedinger, & Pelletier, 1995; van Lehn, 1990; Lesgold, Lajoie, Bunzo, & Eggan, 1992), frontier learning, building on prerequisites (Gagne′, 1977), and sophisticated motivational techniques (Lepper, Woolverton, Mumme, & Gurtner, 1991). Detailed discourse analyses have been performed on small samples of accomplished tutors in an attempt to identify sophisticated tutoring strategies (Fox, 1993; Hume, Michael, Rovick, & Evens, 1996; Merrill, Reiser, Ranney, & Trafton, 1992; Moore, 1995; Putnam, 1987). However, we discovered that the vast majority of these sophisticated tutoring strategies were virtually nonexistent in the untrained tutoring sessions that we videotaped and analyzed (Graesser et al., 1995; Person & Graesser, 1999). Tutors clearly need to be trained how to use the sophisticated tutoring skills because they do not routinely emerge in naturalistic tutoring with untrained tutors.

**AutoTutor-1.** The dialog moves in AutoTutor-1 are generated by 15 fuzzy production rules (see Kosko, 1992) that primarily exploit data provided by the LSA module (see Person, Graesser et al., 2000). AutoTutor-1's production rules are tuned to the following LSA parameters: (a) Student Assertion Quality, (b) Student Ability Level, and (c) Topic Coverage. Each production rule specifies the LSA parameter values for which a particular dialog move should be generated. For example, consider the following dialog move rules:

(1) IF [Student Assertion match with good answer text = HIGH or VERY HIGH] THEN [select POSITIVE FEEDBACK dialog move]

(2) IF [Student Ability = MEDIUM or HIGH & Student Assertion match with good answer text = LOW] THEN [select HINT dialog move]

In Rule (1) AutoTutor will provide Positive Feedback (e.g., "Right") in response to a high quality student Assertion, whereas in Rule (2) AutoTutor will generate a Hint to bring the relatively high ability student back on track (e.g., "What about the size of the programs you need to run?"). The dialog move generator currently controls 12 dialog moves: Pump, Hint, Splice, Prompt, Prompt Response, Elaboration, Summary, and five forms of immediate short-feedback (positive, positive-neutral, neutral, negative-neutral, and negative).

During the tutorial conversation for each tutoring topic, AutoTutor must keep track of which good answer aspects have been covered along with which dialog moves have been previously generated. AutoTutor-1 uses the LSA Topic Coverage metric to track the extent to which each good answer aspect ($A_i$) for a topic has been covered in the tutorial conversation. That is, LSA computes the extent to which the various tutor and student turns cover the good answer aspects associated with a particular topic. The Topic Coverage metric varies from 0 to 1 and gets updated for each good answer aspect with each tutor and student turn. If some threshold (t) is met or exceeded, then the aspect $A_i$ is considered covered. AutoTutor also must decide which good answer aspect to cover next. In AutoTutor-1, the selection of the next good answer aspect to cover is determined by the zone of proximal development. AutoTutor-1 decides on the next aspect to cover by selecting the aspect that has the highest *sub*threshold coverage score. Therefore, AutoTutor-1 builds on the fringes of what is known (or has occurred) in the discourse space between the student and tutor. A topic is finished when all of the aspects have coverage values that meet or exceed the threshold t.

**AutoTutor-2.** We believe that the most effective computer tutor will be a hybrid between naturalistic tutorial dialog and ideal pedagogical strategies. AutoTutor-2 incorporates tutoring tactics that attempt to get the student to articulate the good answer aspect that is selected. AutoTutor-1 considers aspect ($A_i$) as covered if it is articulated by either the student *or* the tutor, whereas AutoTutor-2 counts only what the student says when evaluating coverage. Therefore, if aspect ($A_i$) is not articulated by the student, it is not considered as covered. This forces the student to articulate the explanations in their entirety, an extreme form of constructivism. In order to flesh out a particular aspect ($A_i$); AutoTutor-2 uses discourse patterns that organize dialog moves in terms of their progressive specificity. Hints are less specific than Prompts, and

Prompts are less specific than Elaborations. Thus, AutoTutor-2 cycles through a Hint-Prompt-Elaboration pattern until the student articulates the aspect ($A_i$). The other dialog moves (e.g., short feedbacks and summaries) are controlled by the fuzzy production rules that were described for AutoTutor-1.

AutoTutor-2 has two additional features for selecting the next good answer aspect to be covered. First, AutoTutor-2 enhances discourse <u>coherence</u> by selecting the next aspect ($A_i$) that is most similar to the previous aspect that was covered. Second, AutoTutor-2 selects <u>pivotal aspects</u> that have a high family resemblance to the remaining uncovered aspects; that is, AutoTutor-2 attempts to select an aspect that has the greatest content overlap with the remaining aspects to be covered. Whereas AutoTutor-1 capitalizes on the zone of proximal development exclusively, AutoTutor-2 also considers conversational coherence and pivotal aspects when selecting the next good answer aspect to cover.

## Evaluation of the DAN and the Dialog Move Generation Mechanisms

This paper provided an overview of AutoTutor's mechanisms for managing conversations and generating dialog moves. We are currently analyzing pre- and post-DAN tutoring transcripts to determine whether the DAN significantly improves AutoTutor's conversational capacities. In order to evaluate the two dialog move generation mechanisms, we are planning a study in which the pedagogical effectiveness of the tutor's dialog moves will be assessed for AutoTutor-1 and AutoTutor-2. In addition, student learning gains will be measured for the two versions of AutoTutor versus comparable learning conditions.

## References

Anderson, J. R., Corbett, A. T., Koedinger, K. R., & Pelletier, R. 1995. Cognitive tutors: Lessons learned. *The Journal of the Learning Sciences, 4*, 167-207.

Bloom, B. S. 1984. The 2 sigma problem: The search for methods of group instruction as effective as one-to-one tutoring. *Educational Researcher, 13,* 4-16.

Bransford, J. D., Goldman, S. R., & Vye, N. J. 1991. Making a difference in people's ability to think: Reflections on a decade of work and some hopes for the future. In R. J. Sternberg & L. Okagaki, eds. *Influences on children,* pp. 147-180. Hillsdale, NJ: Erlbaum.

Clark, H. H. & Schaefer, E. F. 1987. Collaborating on contributions to conversations. *Language and Cognitive Processes, 2,* 19-41.

Cohen, P. A., Kulik, J. A., & Kulik, C. C. 1982. Educational outcomes of tutoring: A meta-analysis of findings. *American Educational Research Journal, 19,* 237-248.

Collins, A. 1985. Teaching reasoning skills. In S.F. Chipman, J.W. Segal, & R. Glaser, eds., *Thinking and learning skills,* vol. 2, pp 579-586. Hillsdale, NJ: Erlbaum.

Collins, A., Brown, J. S., & Newman, S. E. 1989. Cognitive apprenticeship: Teaching the craft of reading, writing, and mathematics. In L. B. Resnick, ed., *Knowing, learning, and instruction: Essays in honor of Robert Glaser,* pp. 453-494. Hillsdale, NJ: Erlbaum.

Foltz, P.W. 1996. Latent semantic analysis for text-based research. *Behavior Research Methods, Instruments, and Computers, 28,* 197-202.

Foltz, P. W., Britt, M. A., & Perfetti, C. A. 1996. Reasoning from multiple texts: An automatic analysis of readers' situation models. *Proceedings of the 18th Annual Conference of the Cognitive Science Society,* pp. 110-115. Mahwah, NJ: Erlbaum.

Fox, B. 1993. *The human tutorial dialog project.* Hillsdale, NJ: Erlbaum.

Gagné, R. M. 1977. *The conditions of learning* (3rd ed.). New York: Holdt, Rinehart, & Winston.

Graesser, A.C., Franklin, S., & Wiemer-Hastings, P. & the Tutoring Research Group 1998. Simulating smooth tutorial dialog with pedagogical value. *Proceedings of the American Association for Artificial Intelligence,* pp. 163-167. Menlo Park, CA: AAAI Press.

Graesser, A. C., & Person, N. K. 1994. Question asking during tutoring. *American Educational Research Journal, 31,* 104-137.

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

Graesser, A.C., Wiemer-Hastings, P., Wiemer-Hastings, K., Harter, D., Person, N., & the Tutoring Research Group in press. Using latent semantic analysis to evaluate the contributions of students in AutoTutor. *Interactive Learning Environments.*

Grice, H. P. 1975. Logic and conversation. In P. Cole & J. Morgan, eds., *Syntax and semantics, vol. 3: Speech acts,* pp. 41-58. New York: Academic Press.

Hobbs, J. R. 1979. Coherence and coreference. *Cognitive Science*, *3*, 67-90.

Hu, X., Graesser, A. C., & the Tutoring Research Group 1998. Using WordNet and latent semantic analysis to evaluate the conversational contributions of learners in the tutorial dialog. *Proceedings of the International Conference on Computers in Education*, vol. 2, pp. 337-341. Beijing, China: Springer.

Hume, G. D., Michael, J.A., Rovick, A., & Evens, M. W. 1996. Hinting as a tactic in one-on-one tutoring. *The Journal of the Learning Sciences, 5,* 23-47.

Kosko, B. 1992. *Neural networks and fuzzy systems.* New York: Prentice Hall.

Landauer, T. K., & Dumais, S. T. 1997. A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review.*

Landauer, T. K., Foltz, P. W., & Laham, D. 1998. An introduction to latent semantic analysis. *Discourse Processes, 25,* 259-284.

Lepper, M. R., Woolverton, M., Mumme, D.L., & Gurtner, J.L. 1991. Motivational techniques of expert human tutors: Lessons for the design of computer-based tutors. In S.P. Lajoie & S.J. Derry, eds., *Computers as cognitive tools,* pp. 75-105. Hillsdale, NJ: Erbaum.

Lesgold, A., Lajoie, S., Bunzo, M., & Eggan, G. 1992. SHERLOCK: A coached practice environment for an electronics troubleshooting job. In J. H. Larkin & R. W. Chabay, eds., *Computer-assisted instruction and intelligent tutoring systems,* pp. 201-238. Hillsdale, NJ: Erlbaum.

McCauley, L., Gholson, B., Hu, X., Graesser, A. C., & the Tutoring Research Group 1998. Delivering smooth tutorial dialog using a talking head. *Proceedings of the Workshop on Embodied Conversation Characters,* pp. 31-38. Tahoe City, CA: AAAI and ACM.

McLaughlin, M. L. 1984. *Conversation: How talk is organized.* Beverly Hills, CA: Sage.

Nofsinger, R. E. 1991. *Everyday Conversation.* Newbury Park, CA: Sage.

Merrill, D. C., Reiser, B. J., Ranney, M., & Trafton, J. G. 1992. Effective tutoring techniques: A comparison of human tutors and intelligent tutoring systems. *The Journal of the Learning Sciences, 2,* 277-305.

Moore, J.D. 1995. *Participating in explanatory dialogues.* Cambridge, MA: MIT Press.

Olde, B. A., Hoeffner, J., Chipman, P., Graesser, A. C., & the Tutoring Research Group 1999. A connectionist model for part of speech tagging. *Proceedings of the American Association for Artificial Intelligence,* pp. 172-176. Menlo Park, CA: AAAI Press.

Palinscar, A. S., & Brown, A. 1984. Reciprocal teaching of comprehension-fostering and comprehension-monitoring activities. *Cognition & Instruction, 1,* 117-175.

Person, N. K., Bautista, L., Kreuz, R. J., Graesser, A. C. & the Tutoring Research Group in press. The dialog advancer network: A conversation manager for AutoTutor. In the ITS 2000 *Proceedings of the Workshop on Modeling Human Teaching Tactics and Strategies.* Montreal, Canada.

Person, N. K., Craig, S., Price, P., Hu, X., Gholson, B., Graesser, A. C., & the Tutoring Research Group 2000. Incorporating Human-like Conversational Behaviors into AutoTutor. Paper submitted to The Fourth International Conference on Autonomous Agents, Workshop on Achieving Human-like Behavior in Interactive Animated Agents, Barcelona, Spain.

Person, N. K., & Graesser, A. C. 1999. Evolution of discourse in cross-age tutoring. In A.M. O'Donnell and A.

King, eds., *Cognitive perspectives on peer learning,* pp. 69-86. Mahwah, NJ: Erlbaum.

Person, N. K., Graesser, A. C., & the Tutoring Research Group in press. Designing AutoTutor to be an effective conversational partner. *Proceedings for the 4th International Conference of the Learning Sciences.* Ann Arbor, MI.

Person, N. K., Graesser, A. C., Kreuz, R. J., Pomeroy, V. & the Tutoring Research Group. 2000. *Simulating human tutor dialog moves in AutoTutor.* Submitted to *International Journal of Artificial Intelligence in Education.*

Person, N. K., Graesser, A. C., Magliano, J. P., & Kreuz, R. J. 1994. Inferring what the student knows in one-to-one tutoring: The role of student questions and answers. *Learning and Individual Differences, 6,* 205-29.

Person, N. K., Klettke, B., Link, K., Kreuz, R. J., & the Tutoring Research Group 1999. The integration of affective responses into AutoTutor. *Proceeding of the International Workshop on Affect in Interactions,* pp. 167-178. Siena, Italy.

Person, N. K., Kreuz, R. J., Zwaan, R., & Graesser, A. C. 1995. Pragmatics and pedagogy: Conversational rules and politeness strategies may inhibit effective tutoring. *Cognition and Instruction, 13,* 161-188.

Putnam, R. T. 1987. Structuring and adjusting content for students: A study of live and simulated tutoring of addition. *American Educational Research Journal, 24,* 13-48.

Sacks, H., Schegloff, E. A., & Jefferson, G. 1978. A simplest systematics for the organization of turn taking for conversation. In J Schenkein, ed., *Studies in the organization of conversational interaction.* New York: Academic Press.

van Lehn, K. 1990. *Mind bugs: The origins of procedural misconceptions.* Cambridge, MA: MIT Press.

Wiemer-Hastings, P., Graesser, A. C., Harter, D., & the Tutoring Research Group 1998. The foundations and architecture of AutoTutor. *Proceedings of the 4th International Conference on Intelligent Tutoring Systems,* pp. 334-343. Berlin, Germany: Springer-Verlag.

Wiemer-Hastings, P., Wiemer-Hastings, K., & Graesser, A. C. 1999. Improving an intelligent tutor's comprehension of students with Latent Semantic Analysis. *Artificial Intelligence in Education,* pp. 535-542. Amsterdam: IOS Press.