

Skill-Specific Spoken Dialogs in a Reading Tutor that Listens

Gregory Aist

Project LISTEN

LTI, CMU, 4910 Forbes Ave

Pittsburgh, PA 15213 USA

aist@cs.cmu.edu

ABSTRACT

Project LISTEN's Reading Tutor listens to children read aloud. A controlled study indicates that the Reading Tutor helps children's reading comprehension. However, the results for word attack (decoding) skills and word identification skills were not statistically better than in the control condition. Our thesis therefore proposes to develop skill-specific dialogs based on cognitive skill models and successful tutoring strategies. These dialogs will be dynamically assembled by the Reading Tutor and include text, speech, illustrations, and dialog parameters. We hypothesize that such dialogs will improve elementary students' reading abilities.

Keywords

Children, reading, speech recognition, spoken dialog, intelligent tutoring systems

THE LITERACY CHALLENGE

Learning to read is a critical yet challenging task. One-on-one instruction by trained human tutors can succeed for students whom large-group classroom instruction has failed [9], but is expensive and often unavailable.

THE TECHNOLOGY OPPORTUNITY

Advances in speech recognition and spoken dialog technology enable computer-based oral reading tutoring [4, 7, 8]. Intelligent tutoring systems based on cognitive principles have proven successful for other skills such as algebra and computer programming (e.g. [3]). By combining spoken dialogue and intelligent tutoring systems we hope to come closer to the substantial gains in reading skill observed for individual reading tutoring.

BACKGROUND: A READING TUTOR THAT LISTENS

Project LISTEN's Reading Tutor [5, 2] runs in Windows™ 95, 98 or NT 4.0 on a Pentium™, with a noise-canceling headset microphone and a standard

mouse. The Tutor incorporates materials adapted from *Weekly Reader* (a newsmagazine for children) and other sources. Roughly speaking, the Reading Tutor displays a sentence, listens to the child read it, and provides help in response to requests or on its own initiative based on student performance. The student can read a word, a sentence, or part of a sentence aloud. The student can click on a word or on *Help* for assistance. The student can click on *Back* to move to the previous sentence or *Go* to move to the next sentence (Figure 1). The student can click on *Story* to pick a different story, or on *Goodbye* to log out.

The Reading Tutor can provide help on a word (e.g. by speaking the word), provide help on a sentence (e.g. by reading it aloud), backchannel ("mm-hmm"), interrupt a student's mistake, provide tips on using the system, and navigate within the story (e.g. go on to the next sentence[1]).

Successive versions of the Reading Tutor have had extensive field trials in classroom, school computer lab, and public library settings [6].

THESIS WORK: SKILL-SPECIFIC DIALOGS

Current versions of the Reading Tutor use a completely literature-based set of materials and allow the student to

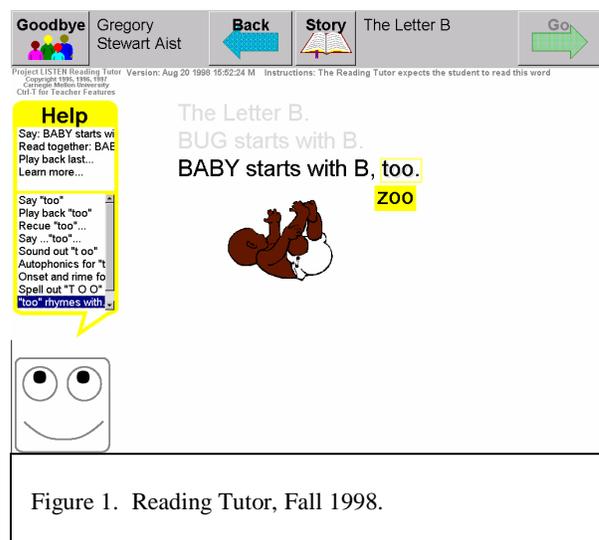


Figure 1. Reading Tutor, Fall 1998.

freely choose material to read. Recent preliminary results from a spring 1998 4-month-long study indicate that the Reading Tutor can be more effective than large-group classroom instruction at improving passage comprehension. However, the results for other important reading skills -- word attack (sounding out) skills and word identification skills -- were not as heartening. Our thesis work proposes extending the Reading Tutor to automatically construct instructional materials -- text, recordings, illustrations, and dialog parameters -- aimed at improving specific reading skills. These materials will be constructed by a combination of automatic generation -- for text and synthesized speech -- and automatic selection -- for text, pre-recorded audio, and illustrations. Dialog parameters here include whether the child may skip a sentence by clicking *Go*, what kinds of help (such as sounding out or rhyming hints) to supply, and whether the Reading Tutor should read the material first or let the child read first.

We focus on three reading skills:

Word Attack (that is, decoding skills) Children must be able to turn a printed word into sound.

Word Comprehension Children must learn not just the sound of words but the meaning of words as well.

Passage Comprehension Making meaning out of print requires not just understanding single words, but constructing meaning for an entire story.

For each of these skills, we will investigate:

1. How to model the skill as a set of smaller components or rules: letter-to-sound rules for word attack, morpheme-to-meaning rules for word comprehension, and lexical entries for passage comprehension.
2. Which rules to use as case studies for developing tutorial dialogs, such as the letter-to-sound rule 'ph' → /f/ or the morpheme-to-meaning rule 'in-' → "not".
3. How to decide that extra help is needed.
4. How to construct an appropriate dialog to train the specific rules selected for the case studies.
5. When to present the training materials: at the beginning or end of a session or story, or when the student encounters difficulty?
6. How to adapt the Reading Tutor's normal reading assistance to the generated material.

We will conduct experiments to test how well the resulting dialogs work at training the target rules.

ACKNOWLEDGMENTS

This research is supported in part by the National Science Foundation (NSF) under Grants No. IRI-9505156 and CDA-9616546 and by the author's NSF Graduate Fellowship and Harvey Fellowship. Any opinions,

findings, conclusions, or recommendations expressed in this publication are those of the author(s) and do not necessarily reflect the views of NSF or the official policies, either expressed or implied, of the sponsors or of the United States Government. We also thank our committee: Jack Mostow, advisor (Carnegie Mellon), Alex Rudnicky (Carnegie Mellon), Al Corbett (Carnegie Mellon), and Charles Perfetti (University of Pittsburgh).

REFERENCES

1. Aist, G. S. 1997. Challenges for a Mixed Initiative Spoken Dialog System for Oral Reading Tutoring. AAAI 1997 Spring Symposium on Computational Models for Mixed Initiative Interaction. AAAI Technical Report SS-97-04.
2. Aist, G. S., and Mostow, J. 1997. Adapting Human Tutorial Interventions for a Reading Tutor that Listens: Using Continuous Speech Recognition in Interactive Educational Multimedia. In Proceedings of CALL 97: Theory and Practice of Multimedia in Computer Assisted Language Learning. Exeter, UK.
3. Anderson, J. R., Corbett, A. T., Koedinger, K., & Pelletier, R. (1995). Cognitive tutors: Lessons learned. *The Journal of Learning Sciences*, 4,167-207. http://act.psy.cmu.edu/ACT/papers/Lessons_Learned-abs.html
4. Bernstein, J., and Ritschev, D. 1991. A voice interactive language instruction system. In Proceedings of the Second European Conference on Speech Communication and Technology (EUROSPEECH 91), Genova, Italy, vol. 2, pp. 981-984.
5. Mostow, J., and Aist, G. S. 1997. The Sounds of Silence: Towards Automatic Evaluation of Student Learning in a Reading Tutor that Listens. In Proceedings of the 1997 National Conference on Artificial Intelligence (AAAI 97), pages 355-361.
6. Mostow, J., and Aist, G. S. 1997. When Speech Input is Not an Afterthought: A Reading Tutor that Listens. Workshop on Perceptual User Interfaces, Banff, Alberta, Canada, October 1997.
7. Mostow, J., Roth, S., Hauptmann, A. G., and Kane, M. 1994. A Prototype Reading Coach that Listens. In Proceedings of the Twelfth National Conference on Artificial Intelligence (AAAI 94).
8. Russell, M., Brown, C., Skilling, A., Series, R., Wallace, J., Bohnam, B., and Barker, P. 1996. Applications of automatic speech recognition to speech and language development in young children. ICSLP 96, Philadelphia.
9. Wasik, B., and Slavin, R. 1993. Preventing really reading failure with one-to-one tutoring: A review of five programs. *Reading Research Quarterly* 28(2): 179-200.

