

Interactive Event: A Better Reading Tutor that Listens

Jack Mostow, Greg Aist*, Juliet Bey, Wei Chen, Al Corbett, Weisi Duan, Nell Duke*, Minh Duong, Donna Gates, José P. González, Octavio Juarez, Martin Kantorzyk, Yuanpeng Li, Liu Liu, Margaret McKeown*, Christina Trotochaud*, Joe Valeri, Anders Weinstein, David Yen

Project LISTEN¹ (www.cs.cmu.edu/~listen), School of Computer Science
Carnegie Mellon University, Pittsburgh, PA 15213, USA (* and various other institutions)
mostow@cs.cmu.edu

Project LISTEN's Reading Tutor listens to children read aloud, and helps them learn to read, as illustrated on the Videos page of our website. This Interactive Event encompasses both this basic interaction and new extensions we are developing.

To accelerate fluency development, we are generating real-time visual feedback on children's oral reading expressiveness by mapping prosodic features such as timing, pitch, and intensity to graphical features such as position, shape, and color. To design more effective practice on individual words, we are conducting an experiment to investigate whether and how the amount of context in which the student practices a word – in a sentence, in a phrase, in a bigram, in isolation, or not at all – affects the time to read the word subsequently in connected text.

To accelerate vocabulary development, we are augmenting children's encounters of words in stories with additional instruction and encounters in multiple contexts required to acquire word meaning. To foster active processing required for successful learning, these encounters challenge the child to think about how words relate to context and to other words. We are developing automated methods to help generate effective contexts for learning word meaning, to generate useful challenges, to compute their answers, and to provide informative feedback to children's responses.

To teach explicit reading comprehension strategies, we are adapting expert human instruction into scripted scenarios for Reading Tutor dialogue. The strategies include activating background knowledge, visualizing, asking questions, and summarizing. We are working to automate the scripting process of generating comprehension instruction, for example by generating good questions about a story and scaffolding children to make up their own. As Chen, Mostow, and Aist's ITS2010 paper reports, we are attacking the problem of recognizing children's free-form spoken responses to tutor prompts by training them to respond more predictably, and by exploiting this predictability to improve speech recognition. This work aims to enable the Reading Tutor, and perhaps other tutors some day, to listen to children not just read but talk.

¹ The research reported here was supported by the Institute of Education Sciences, U.S. Department of Education, through Grants R305A080628, R305A080157, and R305B070458. The opinions expressed are those of the authors and do not necessarily represent the views of the Institute and the U.S. Department of Education. We thank the educators and students who help generate and analyze our data.