

A Reading Coach that Listens: (Edited) Video Transcript

Jack Mostow, Alex Hauptmann, Steven F. Roth, Matt Kane, Adam Swift, Lin Chase, Bob Weide

Project LISTEN, 215 Cyert Hall, Carnegie Mellon University Robotics Institute
4910 Forbes Avenue, Pittsburgh, PA 15213-3890
mostow@cs.cmu.edu

COPYRIGHT NOTICE

This publication and its companion paper and video are copyrighted:

- J. Mostow, S. Roth, A. Hauptmann, M. Kane, A. Swift, L. Chase, and B. Weide. (August 1994). A reading coach that listens: (edited) video transcript. *Proceedings of the Twelfth National Conference on Artificial Intelligence (AAAI94)*. Seattle, WA.
- J. Mostow, S. Roth, A. G. Hauptmann, and M. Kane. (August 1994). A Prototype Reading Coach that Listens. *Proceedings of the Twelfth National Conference on Artificial Intelligence (AAAI-94)*. Seattle, WA, American Association for Artificial Intelligence, Recipient of the AAAI-94 Outstanding Paper Award.
- J. Mostow, S. Roth, A. Hauptmann, M. Kane, A. Swift, L. Chase, and B. Weide. (August 1994). A Reading Coach that Listens (6-minute video). *Video Track of the Twelfth National Conference on Artificial Intelligence (AAAI94)*. Seattle, WA, American Association for Artificial Intelligence.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

A Reading Coach that Listens: (Edited) Video Transcript

Jack Mostow, Alex Hauptmann, Steven F. Roth, Matt Kane, Adam Swift, Lin Chase, Bob Weide

Project LISTEN, Carnegie Mellon Robotics Institute, 215 Cyert Hall, 4910 Forbes Avenue, Pittsburgh, PA 15213-3890

At Carnegie Mellon University, Project LISTEN¹ is taking a novel approach to the problem of illiteracy. We have developed a prototype automated reading coach that listens to a child read aloud, and helps when needed. The coach provides a combination of reading and listening, in which the child reads wherever possible, and the coach helps wherever necessary -- a bit like training wheels on a bicycle. Let's see how the automated coach responds to various things a child might do. The output of the automatic speech recognizer is displayed at the bottom of the screen.

Help when needed: The coach recues a misread word by rereading the words that lead up to it, just like the expert reading teachers whom the coach is modelled after. This context often helps the reader correct the word on the second try:

Text: The cow lives on the farm.
Reader: "The cow lives on the farm."
Text: She eats grass all day long.
Reader: "She eats good all day long."
Coach: SHE EATS
Reader: "grass"
Coach: GRASS. PLEASE CONTINUE.

Support comprehension: The coach is designed to emphasize comprehension and ignore minor mistakes, such as repeated words. However, the word "very" is important to the meaning of the sentence, so the coach asks the reader to reread it. The coach rereads the sentence to help the reader comprehend it:

Text: At night she is very tired.
Reader: "At night ... at night she is tired."
Coach: READ THIS WORD AGAIN. [flashes "very"]
Reader: "tired?"
Coach: VERY. AT NIGHT SHE IS VERY TIRED.

Maintain flow: When the reader gets stuck, the coach jumps in, enabling the reader to complete the sentence:

¹**Note:** For acknowledgements, further references, and technical details, please see (Mostow et al, 1994b).

This research was supported primarily by the National Science Foundation under Grant Number MDR-9154059 and by the Advanced Research Projects Agency, DoD, through DARPA Order 5167, monitored by the Air Force Avionics Laboratory under contract N00039-85-C-0163. Lin Chase is supported by a Howard Hughes Doctoral Fellowship from the Hughes Research Laboratory. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the sponsors or of the United States Government.

Text: Then she slowly comes home.

Reader: "Then she then she s ... s ... "

Coach: SLOWLY

Reader: "slowly comes home."

Minimize disruption: Since short function words like "to" and "be" do not usually affect comprehension, the coach refrains from interrupting the reader to correct this omission:

Text: I want to be milked, she says.

Reader: "I want milked she says"

We're having children try out this prototype coach to help us improve it. [Show children using coach].

Clicking for help: To get help with a word, the child can click on it. [Coach speaks word.] This feature is very useful, but children often don't realize when they need help. [Child misreads "democratic" as "dramatic," even after coach recues it.]

Tolerate recognizer inaccuracy: We're working to make the coach recognize children's speech more accurately, and behave reasonably even when the speech recognizer is inaccurate. [Coach misrecognizes "slowly," causing it to reread the sentence.]

In an earlier study we compared how well second graders read with and without similar assistance. Without assistance, they missed one word in eight, which is considered overly frustrating. With assistance, they missed fewer than one word in forty, enabling them to read and comprehend material more than six months beyond their independent reading level.

Children can't read to learn until they learn to read -- whether it's a science passage or anything else. We need to find out how the coach can help children learn over time, and explore how the coach can help in ways that human teachers cannot -- for example, by modifying the text dynamically, and by tapping into the motivational power of computers. [Child comments on coach.]

Project LISTEN builds on years of previous government-funded research in basic speech technology. It has the potential to pay back many times over for the cost of that research, since illiteracy costs the United States over 225 billion dollars every year (literacy).

Moreover, this work applies to several important areas in addition to children's reading instruction, including adult literacy, English as a second language, and foreign language learning. It opens the door to a new generation of intelligent tutoring systems that can listen to their students.

References

