Towards Worldwide Literacy: Technological Affordances, Economic Challenges, Affordable Technology

Gregory Aist

4215 Newell-Simon Hall
Robotics Institute
Carnegie Mellon University
5000 Forbes Ave.
Pittsburgh, Pennsylvania 15213 USA

Introduction

The ability to read is taken for granted by those of higher socioeconomic status in the developed world. However, many people from disadvantaged backgrounds in the developed world – and even more in the developing world – are unable to read or write well enough to thrive in today's technologically advanced global society; many cannot read at all. The International Adult Literacy Survey reports that "... while countries differ in the literacy attainment of their adult populations, none does so well that it can be said that it has no literacy problem" (Human Resources Development Canada, 1994ff). UNESCO (1999) reported 876 million illiterate people in the world as of 2000, with especially high rates of illiteracy in (for example) sub-Saharan Africa (39.7%) and southern Asia (45.8%).

In this paper, we explore how computer technology can meet the economic and social challenges of literacy learning in the developing world – not only learning in widely spoken languages such as Arabic, (Mandarin) Chinese, English, and Swahili, but also in a child's own native language – every language in the world.

Technological affordances

There is a long tradition of research on educational technology for literacy, from early work such as Atkinson and Hansen (1966) to a broad range today (for example, Vanderbilt 1996). Researchers have demonstrated that computer software for literacy instruction can improve various literacy skills, from phonological awareness and word identification (Barker and Torgesen 1995) to word comprehension (Aist et al. 2001) and passage comprehension (Mostow et al. in press). While certainly science is far from solving all problems related to computer-assisted literacy learning, we can say that research is available that sheds light on how computers can be used to help children learn to read (e.g. NRP 2000 chapter 6, "Computer Technology and Reading Instruction").

The present author has worked on technology for helping children learn to read for over five years as a member of Carnegie Mellon's Project LISTEN. Project LISTEN (since 1990 or thereabouts) has been developing, improving, and evaluating a Reading Tutor that adapts automated speech recognition to listen to children read aloud, and helps them learn to read. To date Project LISTEN has focused principally on first language learning, in English, in U.S. elementary schools. Table 1 presents some highlights of results from Project LISTEN's research, focusing on those with which the present author has been directly involved.

The Reading Tutor can not only take alternating turns, but also allow user interruption, produce backchannel feedback, and interrupt in response to a (perceived) student mistake.	Aist & Mostow AAAI CAHM 1997; Aist & Mostow CALL 1997; Aist ICSLP 1998.
Speech data collected during the course of Reading Tutor use can be employed for acoustic training – without the need for manual transcription.	Aist et al. 1998.
Users can write and narrate new stories for children to read with the Reading Tutor.	Mostow & Aist AAAI 1999; Mostow & Aist USPTO 1999.
Automated experiments let researchers test the effectiveness of educational interventions.	Mostow & Aist AAAI 1997; Aist & Mostow AAAI AMDLP 1998; Aist & Mostow ESCA 1999; Aist & Mostow ITS-AML 2000; Mostow et al. NAACL 2001.
Taking turns picking stories results in faster, better story choice than always letting the student choose what to read next.	Aist & Mostow STILL 2001 SC.
Adding automatically constructed factoids to text – that is, inserting a comparison of a word in the text to a synonym, hypernym ("*astronaut* is a kind of *traveler*"), or antonym– can help children learn the meaning of words better than text alone.	Aist AI-ED 2001.
Computer-assisted oral reading can help third graders learn vocabulary better than a classroom control.	Aist Ph.D. 2000; Aist et al. AI-ED 2001; Mostow et al. AI-ED 2001.

Table 1. Highlights of research on Project LISTEN's Reading Tutor.

Economic challenges

The world of science is developing solutions to the challenge of literacy learning – but will the developing world benefit from the science of literacy learning? Read this telling quote from the Atkinson and Hansen paper cited earlier (1966):

> "In September 1966 the CAI reading program will be used with approximately 100 first-grade children in Brentwood School, which is located in a racially mixed low-income area of East Palo Alto. Currently a building is being constructed on the school grounds to house the computer system and student terminals." (p. 21.)

Note the stark contrast between the resources available to communities that need the most help – whether East Palo Alto or the developing world – and the resources needed to truly transform their educational opportunities via educational technology – a new building and a large (in every sense of the word) computer, or even today expensive desktop computers. Following Moore's Law of exponential doubling in power every 18 months, computers have come an amazing distance since then in terms of computational power, physical size, and cost. For example, each student in Figure 1 is reading at a desktop computer far more advanced than any computer available at the time of Atkinson and Hansen's writing (Figure 1). Nonetheless, the cost of desktop or laptop computers to run current generations of software remains prohibitive for many communities.

Figure 1. Children using Project LISTEN's Reading Tutor (http://www.cs.cmu.edu) during a summer reading clinic at an urban elementary school in Pittsburgh, Pennsylvania, U.S.A. Photo credit: Mary Beth Sklar.

There is also the challenge of building sustainable models of software development to bring educational technology within the reach of the developing world. Even in the developed world, educational software is a difficult economic proposition. As Soloway (1998) noted, "No one is making serious money selling educational software… It's basic economics: prices are high because the costs for developing good educational software are sincere; demand is low because the schools aren't spending their funds on software." (p. 11). Furthermore, that quote refers (more or less) to schools in the developed world. It is difficult to see how any public company in the developing world could begin – let alone sustain – a substantial development effort focusing on software for schools serving the children of the poorest of the poor, where $1 per day must provide for their needs and wants. This business and social challenge on top of the technical concerns discussed before compounds the problem – yet does not place a solution fully out of reach.

Affordable technology

Handheld devices are part of the solution to the economic challenge of the cost of computing. Such devices are not only less expensive than desktop or laptop computers, but require less infrastructure as well. Concerns of price might be partly addressed for now by transfer of previously used but still functional hardware from the "early adopters" in developed countries to the developing world. A handheld device such as the Palm Pilot shown in Figure 2, while over three years old and thus two generations obsolete, still functions well without interruption of service. A more complete solution might trade off computing power for price: the intentional manufacture of less than state-of-the-minute hardware – but at progressively decreasing prices close to those of inexpensive digital watches.

Figure 2. A three-year-old Palm Pilot: Technologically obsolete but still highly functional.

Even if computer hardware is available, the software challenge remains. How can we build educational technology for developing countries when there are over 6,500 languages in the world (Grimes et al. 2000)?

Towards first language literacy for every child

Currently, the author is in the process of formulating a social and technical proposal to address the following problem: effective software for first language literacy in any language, for any child, anywhere. While the author has been involved in educational software for literacy development for over five years, the present project is less than three months old and is as yet in the pre-proposal stage. We hope that participation in the Development by Design workshop will enable us to meet others interested in working in this area and further formulate a specific and feasible plan for research. For now, we first roughly sketch out steps toward a solution for this massive challenge, and then describe the first step in detail.

1. Software tools. Freely available software tools, focusing on core issues of natural language processing and instructional technology.
2. Local development. Cooperation between local teachers, researchers, companies, government bodies, and/or NGOs to build language-specific and country- or region-specific educational software for first language literacy.
3. Educational implementation. In order to be effective, educational software must be used in fruitful ways – used in classrooms, homes, and perhaps in educational toys.

We now describe the first step: software tools for literacy instruction.

Software tools for literacy instruction

Literacy instruction encompasses a wide range of skills in both reading and writing, and also requires the skills of speaking and listening (sometimes distinguished as oracy). Consider reading: Children must acquire a wide range of skills to ultimately comprehend text (NRP 2000, Snow et al. 1998). Phonemic awareness allows children to distinguish and manipulate individual sounds in spoken words. Knowledge of print conventions enables children to work with text as placed on a page – for English, left-to-right, top-to-bottom. Mastery of the alphabetic principle reveals that individual sounds are written with letters or letter patterns. Decoding skills codify how to turn printed letters into sounds. Increased fluency leads to faster and more automatic reading. Background knowledge increases text understanding. Vocabulary knowledge is critical for comprehension. Drawing inferences from text and integrating information from multiple sources finally allow the reader to make meaning from print. Furthermore, each of these processes may be more or less challenging in different languages. For example, English has a complicated mapping between letters and sounds, but relatively little inflectional morphology save for some remnants such as plurals; some Romance languages have simpler letter-sound relationships but more complex morphology. Software tools for literacy learning will eventually have to cover that entire space, but we can begin with a restricted ('core') set of skills – for example, we might choose to begin with decoding and vocabulary knowledge.

Rather than develop a single instructional system that works in a variety of languages, we propose to develop a set of language-independent tools for generating Web-based language instruction in any language. By identifying such language-independent components of literacy instruction software, we can amortize their development costs over multiple languages.

For example, a low-tech template for a set of Web pages (or a Web page generator) aimed at vocabulary instruction for a particular word might specify a sequence of interactions:

1. Present to the student a sentence containing a word
2. Present a definition of the word

3. Present a multiple-choice question containing a possible meaning of the word.

The software localizer would then provide several pieces of information – without the need for actual software development: (a) a target word, in spoken, text, and/or graphical form; (b) a sentence containing that word, along with a suitable recording; (c) the definition, in spoken, text, and/or graphical form; (d) correct answers and distractors.

Software to collect student responses and present student averages may also yield language-independent components. For example, much of the computation involved in scoring student answers and aggregating scores is language-independent.

Natural language processing tools such as parsers, part-of-speech taggers, or morphological analyzers may help support language learning software. For an easy example, being able to identify rare words in a language may help a system choose words to explain to a reader (Aist unpublished data). However, such tools are often freely available only for research use, and may not be sufficiently well designed that those outside the field of natural language processing can easily use them. What is needed instead is a system – or toolbox, or system-generating language – that is freely available and probably public domain, and that enables courseware developers to plug in "data" – knowledge about a particular language, texts written in that language, recordings of such texts – and thus construct educational software in that particular language.

The challenge is great, but the rewards are tremendous: native language literacy, assisted by computer software, for any child in any language in the world.

Acknowledgements

This work was supported in part by the National Science Foundation under Grant Nos. REC-9720348 and REC-9979894, 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 authors and do not necessarily reflect the views of the National Science Foundation or the official policies, either expressed or implied, of the sponsors or of the United States Government.

References

Aist, G. S., Mostow, J., Tobin, B., Burkhead, P., Corbett, A., Cuneo, A., Junker, B., and Sklar, M. B. Computer-assisted oral reading helps third graders learn vocabulary better than a classroom control – about as well as one-on-one human-assisted oral reading. In J. D. Moore, C. L. Redfield, and W. L. Johnson (Eds.), *Artificial Intelligence in Education: AI-ED in the Wired and Wireless Future*, pp. 267-277. Amsterdam: IOS Press. Presented at the Tenth Artificial Intelligence in Education (AI-ED) Conference,San Antonio, Texas, May 2001.

Aist, Gregory S. and Mostow, Jack. 2001. Faster, Better Task Choice in a Reading Tutor that Listens. To appear in Philippe DeCloque and Melissa Holland (Editors), *Speech Technology for Language Learning*. The Netherlands: Swets & Zeitlinger Publishers (in revision).

Aist, G. AI-ED 2001. Factoids: Automatically constructing and administering vocabulary assistance and assessment. *Proceedings of the 10th International Conference on Artificial*

Intelligence in Education. San Antonio, Texas, May 19-23, 2001.

Aist, G., Mostow, J., Tobin, B., Burkhead, P., Corbett, A., Cuneo, A., Junker, B., and Sklar, M. B. AI-ED 2001. Computer-assisted oral reading helps children learn vocabulary better than a classroom control - and even just as well as one-on-one human-assisted oral reading. *Proceedings of the 10th International Conference on Artificial Intelligence in Education*. San Antonio, Texas, May 19-23, 2001.

Aist, G. 2000. Helping Children Learn Vocabulary during Computer-Assisted Oral Reading. Ph.D. dissertation, Language Technologies Institute, School of Computer Science, Carnegie Mellon University. http://www.cs.cmu.edu/~aist/Aist-PhD-dissertation.html

Aist, G. and Mostow, J. ITS-AML 2000. Using automated within-subject invisible experiments to test the effectiveness of automated vocabulary assistance. Workshop on Applying Machine Learning to ITS Design/Construction. Montreal, June 19, 2000.

Aist, G. and Mostow, J. 1999. Measuring the Effects of Backchanneling in Computerized Oral Reading Tutoring. *Proceedings of the ESCA Workshop on Prosody and Dialog*. Eindhoven, The Netherlands, September 1999.

Aist, G. 1998. Expanding a time-sensitive conversational architecture for turn-taking to handle content-driven interruption. International Conference on Spoken Language Processing (ICSLP) 1998, Sydney, Australia, Nov. 30 - Dec. 4, 1998. Paper 928.

G. Aist, P. Chan, X. Huang, L. Jiang, R. Kennedy, D. Latimer, J. Mostow, and C. Yeung. 1998. How effective is unsupervised data collection for children's speech recognition? International Conference on Spoken Language Processing (ICSLP) 1998, Sydney, Australia, Nov. 30 - Dec. 4, 1998. Paper 929.

Aist, G. S., and Mostow, J. AAAI-CAHM 1997. A time to be silent and a time to speak: Time-sensitive communicative actions in a reading tutor that listens. AAAI Fall Symposium on Communicative Actions in Humans and Machines. Boston, MA, November, 1997.

Aist, G. S., and Mostow, J. CALL 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.

Atkinson, Richard C., and Duncan N. Hansen. 1966. Computer-assisted instruction in initial reading: The Stanford project. Reading Research Quarterly 2(1): 5-25.

Barker, Theodore Allen, and Torgesen, Joseph K. 1995. An evaluation of computer-assisted instruction in phonological awareness with below average readers. *Journal of Educational Computing Research* **13**(1), pp. 89-103.

Donahue, P. L., Voelkl, K. E., Campbell, J. R., and Mazzeo, J. 1999. *NAEP 1998 Reading Report Card for the Nation and the States.* National Center for Education Statistics, Washington, DC. http://nces.ed.gov/nationsreportcard/pubs/main1998/1999500.shtml

Grimes, Barbara, ed. 2000. Ethonologue: Languages of the World. Summer Institute of Linguistics.
http://www.sil.org/ethnologue/

Human Resources Development Canada, 1994 and following. International Adult Literacy Survey.
http://www.nald.ca/nls/ials/introduc.htm

Jack Mostow, Greg Aist, Juliet Bey, Paul Burkhead, Andrew Cuneo, Susan Rossbach, Brian Tobin, Joe Valeri, and Sara Wilson. 2001. A hands-on demonstration of Project LISTEN's Reading Tutor and its embedded experiments. Demonstration at Language Technologies 2001: 2nd Meeting of the North American Chapter of the Association for Computational Linguistics. Pittsburgh, Pennsylvania, June 2-7, 2001.

Mostow, J., Aist, G. S., Burkhead, P., Corbett, A., Cuneo, A., Eitelman, S., Huang, C., Junker, B., Platz, C., Sklar, M. B., and Tobin, B. A controlled evaluation of computer- versus human-assisted oral reading. In J. D. Moore, C. L. Redfield, and W. L. Johnson (Eds.), Artificial Intelligence in Education: AI-ED in the Wired and Wireless Future, pp. 586-588. Amsterdam: IOS Press. Presented at the Tenth Artificial Intelligence in Education (AI-ED) Conference,San Antonio, Texas, May 2001.

Mostow, J. and Aist, G. 1999. Authoring New Material in a Reading Tutor that Listens. Proceedings of the Sixteenth National Conference on Artificial Intelligence (AAAI-99), Orlando, FL, July 1999, pp. 918-919. In the refereed Intelligent Systems Demonstration track. Also presented at 37th Annual Meeting of the Association for Computational Linguistics (ACL'99), College Park, MD, June, 1999.

Mostow, J. and Aist, G. 1999. Reading and Pronunciation Tutor. United States Patent No. 5,920,838. Filed June 2, 1997; issued July 6, 1999. US Patent and Trademark Office. http://www.uspto.gov/.

Aist, G., and Mostow, J. 1998. Estimating the Effectiveness of Conversational Behaviors in a Reading Tutor that Listens. AAAI Spring Symposium on Applying Machine Learning to Discourse Processing, Stanford, CA, March 1998.

Mostow, J., and Aist, G. 1997. The Sounds of Silence: Towards Automated Evaluation of Student Learning in a Reading Tutor that Listens. In Proceedings of the Fourteenth National Conference on Artificial Intelligence (AAAI-97). American Association for Artificial Intelligence, Providence, RI, July, 1997. Pages 355-361.

Jack Mostow, Gregory S. Aist, Cathy Huang, Brian Junker, Rebecca Kennedy, Hua Lan, DeWitt Latimer IV, Rollanda O'Connor, Regina Tassone, Brian Tobin, and Adam Wierman. 4-Month Evaluation of a Learner-controlled Reading Tutor that Listens. To appear in Philippe DeCloque and Melissa Holland (Editors), *Speech Technology for Language Learning*. The Netherlands: Swets & Zeitlinger Publishers (in press).

National Reading Panel. 2000. Teaching Children to Read.
http://www.nichd.nih.gov/publications/nrppubskey.cfm

Soloway, E. 1998. No one is making money in educational software. *Communications of the ACM* 41(2): 11-15.

UNESCO. 1999. UNESCO Annual Statistical Yearbook 1999. Online at http://unescostat.unesco.org/en/stats/stats0.htm

Vanderbilt (The Cognition and Technology Group at Vanderbilt). 1996. A multimedia literacy series. Communications of the ACM 39(8), 106-109. Now a commercial product from Little Planet Publishing, Nashville TN.