

An Intelligent Tutoring System for Adult Literacy Enhancement

Sowmya Ramachandran, Richard Stottler
Stottler Henke Associates, Inc.
San Mateo, CA

sowmya@stottlerhenke.com, stottler@stottlerhenke.com

ABSTRACT

Highly advanced interactive computer-based training courses are being developed as technology plays an increasingly significant role in training warfighters. Job aids and other field operation guides also contribute to warfighter performance enhancement. But these programs are only as effective as the readiness of the warfighters to benefit from them. Lack of certain basic skills, like reading comprehension, can significantly hinder the effectiveness of training and performance support systems. Of what use is a job aid to someone who does not have the literacy skills to comprehend the information presented therein? This problem is magnified by increased reliance on computer-based training since such systems still rely quite heavily on text for conveying information. Having a minimum standard for recruitment often does not guarantee that recruits will possess adequate reading comprehension skills. An effective warfighter training program will include courses in basic reading skills.

This paper will discuss the issue of using technology to support training in basic reading comprehension skills, specifically in the military environment. We will describe a novel approach to this issue, sponsored by the Navy, that applies findings from research on the cognitive processes underlying reading comprehension to develop an Intelligent Tutoring System (ITS) for teaching reading comprehension skills. ITSs are computer-based tutors that make dynamic assessments of a student's unique skill levels and provide instruction that is adapted to individual skill levels. The assessments and adaptive lesson selection are entirely automated. The tutor is accompanied by an authoring tool that enables the customization of the content completely to suit an organization's needs. The tutor currently includes content related to the Navy.

AUTHOR BIOGRAPHIES

Dr. Sowmya Ramachandran is a research scientist at Stottler Henke Associates, Inc. Dr. Ramachandran received her Ph.D. in Computer Science from the University of Texas at Austin. She has a strong background in a wide variety of Artificial Intelligence techniques, including Intelligent Tutoring Systems, and Machine Learning. Her research interests include application of Artificial Intelligence techniques to Education Technology with a focus on addressing motivational, affective, and meta-cognitive issues. Dr. Ramachandran has headed several ITS development efforts, including one for adult literacy enhancement and one to teach Algebra to at-risk high-school students. She is currently developing a general-purpose authoring framework for rapid development of ITSs.

Richard Stottler co-founded Stottler Henke Associates, Inc., an artificial intelligence consulting firm in San Mateo, California in 1988 and has been the president of the company since then. He has been principal investigator on a large number of tactical decision-making intelligent tutoring system projects conducted by Stottler Henke, including projects for the Navy, Army, and Air Force. Currently, he is working on an intelligent tutoring prototype for the future combat system control vehicle, funded by the US Army STRICOM. He has a Masters degree in computer science from Stanford University.

An Intelligent Tutoring System for Adult Literacy Enhancement

Sowmya Ramachandran, Richard Stottler
Stottler Henke Associates, Inc.
San Mateo, CA

Sowmya@stottlerhenke.com, Stottler@stottlerhenke.com

THE PROBLEM OF LITERACY AND ITS SIGNIFICANCE

Highly advanced interactive computer-based training courses are being developed as technology plays an increasingly significant role in training warfighters. Job aids and other field operation guides also contribute to warfighter performance enhancement. But these programs are only as effective as the readiness of the warfighters to benefit from them. Lack of certain basic skills, like reading comprehension, can significantly hinder the effectiveness of training and performance support systems. Of what use is a job aid to someone who does not have the literacy skills to comprehend the information presented therein? This problem is magnified by increased reliance on computer-based training since such systems still rely quite heavily on text for conveying information. Having a minimum standard for recruitment often does not guarantee that recruits will possess adequate reading comprehension skills. A significant percentage of the Navy's enlisted population, for example, scores below the eighth grade level in reading, writing, and arithmetic. An effective warfighter training program will include courses in basic reading skills.

Literacy software tends to fall into one of two categories, both of which are inappropriate for the Navy's target group. Adult literacy software tends to target truly illiterate adults - those who cannot read at all. This aims too low. To enlist in the Navy requires an ability to read. Navy personnel have mastered the skill of decoding, at least at its most basic level. The other category of software attempts to enhance reading skills. Unfortunately, although the goals of such systems are appropriate, the content is not. It is typically geared toward children, and adults find the content uninteresting or insulting.

The content should be directed toward documents of importance to Navy personnel, such as technical manuals, general materials describing Navy procedures and policies, documents describing general citizenship, etc. Also appropriate would be material of general interest to adults. This might include life-skills (such as balancing a checkbook), descriptions of regions, countries, and cities with which the Navy is currently involved, news, parenting tips, etc. Tools that help Navy personnel create their own

content for reading practice and evaluation would also be very helpful.

Another important requirement is that Navy personnel should be able to remediate themselves through independent study. There are no reading instructors aboard ships where enlisted personnel spend a lot of their time. Therefore, the software should strive to largely replace the instructor. Most existing educational software is designed to work in concert with an instructor.

The need for this software extends far beyond the U.S. Navy. Literacy is one of the most fundamental requirements for succeeding in today's world. Yet, 50 million Americans are not functionally literate. With the growth in availability of enormous amounts of free information via the Internet and the Web, the plight of these people can only worsen. They will be unable to take full advantage of this new resource and could fall further behind in managing their lives as access to the Web becomes an accepted norm.

To address these problems, the Navy has funded research on the development of an Intelligent Tutoring Systems for teaching reading comprehension skills with customizable content. An Intelligent Tutoring System (ITS) is tutoring software that can tailor instruction to an individual student's requirements, and can provide intelligent instructional support similar to human instructors. Such systems would address Navy's need for reading tutors that do not require the presence of human instructors. The ability to easily customize content would enable the Navy to maintain a tutor that is relevant to the trainees' jobs. This paper describes one ITS which was developed as a part of this initiative. We will describe ReadOn, an ITS that can model a student's reading abilities and provide customized instruction. We will also describe an accompanying authoring tool that allows non-programmers to expand the set of reading material available to the tutor.

THE COGNITIVE UNDERPINNINGS OF READING COMPREHENSION

The Cognitive Processes of Reading

Researchers into the processes of reading have developed a componential theory of reading which specifies a set of functionally defined information-processing components that, in interaction with one another, accomplish the more complex task of text comprehension. According to this theory, several processes are active in parallel. Low-level processes output information which is used as input to higher-level processes.

Differences in reading skills are attributed primarily to the degree to which the various processes are automated. Because they are parallel in execution, the processes compete for cognitive resources. Non-automatic processes require substantially more cognitive resources than automatic ones. Improvements in the efficiency of one allow more cognitive processing resources to be allocated to the others, thus improving comprehension. For example, a student who has poor vocabulary will spend a substantial amount of cognitive resources on identifying words. Once the student becomes familiar with more words, word identification becomes more of an automatic process, thus freeing the cognitive resources for higher-level processing of information. Thus, a student model should attempt to model each of these processes separately, at least when there is a reading problem, to support assessment, diagnosis, and appropriate remediation. Furthermore, since repeated practice leads to improved automaticity of skills, our approach is designed to provide extensive guided practice in reading comprehension skills.

Just and Carpenter (1987) have identified the various process components that contribute to reading comprehension. Examples of such components are 1. perceptual processes to decode words, 2. processes to access the meaning of words, 3. processes to assign contextual meaning and roles to words in phrases, and 4. processes to construct stories and mental representations of the text. These processes are highly inter-dependent. For example, mental representations of what has been read so far helps with the identification of the meaning of a following word. Based on their theories, we have developed the reading comprehension skill hierarchy model shown in Figure 1. Under this representation, the skill at the tail of an arrow supports the skill at the head of an arrow. Note that this is an approximate representation since it does not capture the fact that higher-level skills also support lower level skills (e.g. an overall understanding of the context is required sometimes to disambiguate the meaning of a word). The skills shown in the network are described below.

Familiar Topics: This reflects the background knowledge possessed by the reader, and includes an estimate of his/her world knowledge.

Knowledge of integrative cues: This reflects the reader's knowledge of the roles of textual cues that indicates relationships between sentences. Examples of such cues include words like *because*, *since*, *however*, and *nevertheless*.




Figure 1: Part of the hierarchy of reading comprehension skills

Contextual Reasoning: This is the ability to infer word meanings from the context.

Following Causal Inferences: This reflects the ability to follow causal inferences presented in the text.

Following Logical Inferences: This reflects the readers' ability to follow logical inferences presented in the text.

Following Temporal Inferences: This reflects the readers' ability to follow and make inferences from temporal sequences presented in the text.

Syntactic Processing: This refers to syntactic processing skills.

Semantic Processing: This represents the student's skill at processing a sentence semantically.

Constructing Referential Representation: This refers to the conversion of text into internal knowledge structures (propositions). Referred to by some researchers as "building referential representations", this is aided by skills like *anaphora resolution*, and *resolving co-references*.

Anaphora Resolution: This is the ability to match pronouns with their references.

Co-reference Resolution: This is a more general skill than anaphora resolution. It represents the ability to process references to the same object or concept by different names. For example, read the following sentences.

Angela was proud of her son's accomplishments. The boy was an honor student.

Here, the words "her son" and "The boy" refer to the same person.

Word Lexical Identification: This is the ability to decode a printed word, and identify its meaning.

Local Comprehension: This models the ability to understand a short piece of text about the length of a paragraph.

CURRENT COMPUTER BASED APPROACHES TO READING INSTRUCTION

Many current computer-based training systems for teaching reading comprehension skills require various degrees of intervention from human instructors. Some require instructors to assign lessons to students based on their performance on prior lessons. Some involve exercises that need to be graded by human instructors. Such approaches are appropriate for organizations that have reading instructors on hand to carry out these essential tasks.

Commercially available reading tutors often provide instruction in individual reading skills but lack opportunities to practice reading extended text. For instance, they typically have content modules for skills such as inference, vocabulary, identifying the main idea, etc., but not modules that let students practice these skills in the context of everyday reading. Plato™ (www.plato.com) is

an exception and does provide opportunities for such practice.

Finally, most of the commercially available training systems come with a pre-determined content which is of general interest to adults. However, organizations can optimize the return on investment from literacy training if the training involves job-related content. Not only will the employees be learning to read, they will also be learning useful job-related information. Customization of existing literacy CBTs involves contracting with the product developers to re-engineer the content.

AN INTELLIGENT TUTORING FRAMEWORK FOR DEVELOPING READING COMPREHENSION SKILLS

Intelligent Tutoring Systems

Intelligent Tutoring Systems contrast with most Computer Based Training (CBT) systems in that the latter can usually be described as automated text books. That is, most CBTs are developed by using the same approach as a corresponding textbook. In some domains, CBTs include multimedia materials such as video, audio, and other animation that textbooks cannot include. However, these don't really reflect differences in instructional methods. Other than allowing self-navigation, typical CBTs do not attempt to adapt or tailor the instruction to the individual. Additionally, most CBTs do not embody any particular instructional approach, theory, or philosophy, other than the instructional approach that happened to exist in the textbook on which the CBT system is based.

ITSs, on the other hand, adapt instruction much as a human tutor would and approach the benefits of one-on-one instruction. ITSs assess, model, diagnose, and automatically develop a tailored course of remedial instruction for individual students. To truly tailor instruction, a training system must create, develop, and maintain a model of the student, which ITSs do and most CBTs do not. This model is used as a basis for instruction method and content selection, diagnosis, remedial course formulation, re-testing, progress monitoring and reporting, all done automatically. Studies have shown that ITSs are highly effective. Anderson et. al., (1985) showed that students working with a LISP ITS acquired knowledge in about one-third to two-thirds the time that it took a control group of students to acquire the same amount of expertise (as measured by assessment tests). Nichols et. al, (1992) performed an evaluation of an ITS designed to teach avionics troubleshooting skills. They found that students could acquire, in 20 hours, skills comparable to those possessed by technicians with 4 years of experience.

Overview of the ITS

The central idea of ReadOn is to improve the reading skills of students through assessment, reading practice, and remediation. The process of assessment consists of monitoring a student's performance constantly and determining his strengths and weaknesses with respect to reading comprehension skills.

Reading practice consists of activities that exercise a broad range of reading skills. Students are given a passage to read, which is followed by a set of questions relating to the passage to test vocabulary and comprehension. This pairing of text and questions will henceforth be referred to as an exercise set. The reading material and questions cover a broad range of skills. The passages to be read are extracts from the kinds of reading material that a student is likely to encounter in his/her daily life. For the Navy, the content was extracted from the Basic Military Reading package that is part of required reading for every recruit. These passages describe various aspects of the Navy, including its history, culture, and rules and regulations. Thus, much of the reading material offered by the tutor is of direct relevance to a trainee.

In contrast to general reading practice, remediation consists of exercises that strengthen specific skills. These exercises are targeted toward illustrating specific concepts and skills, and are used to remediate a student's weak skills, as diagnosed by the ITS. An example of such a remedial exercise, shown in Figure 2, is one that targets the specific skill of anaphora resolution.

The approach of providing practice on a broad range of skills, resorting to remediation when such practice does not lead to improvement, is most beneficial to the student because, while training in specific skills is important, synthesizing those skills in order to fully comprehend reading material encountered in daily life is a student's ultimate goal. Thus, ReadOn constantly monitors a student in the act of everyday reading, and diagnoses and remediates reading problems within this context.

ReadOn maintains an explicit representation of the skills discussed in the previous section. Exercises include activities for assessing these skills. The student model maintains an estimate of the student's proficiency on each of the skills. ReadOn uses this model to present reading practice that is customized in the following ways:

1. **Selection of optimal challenge level:** Each exercise set provides an appropriate level of challenge to the student. The reading passages should not be easy as to bore a student, yet not so hard that the student is frustrated. Placement tests are used to initially assess the reading skills of a student. Like the exercise sets used for reading practice, a placement test consists of a short piece of text followed by a fixed set of questions that test a broad range of reading skills. The purpose of such tests is to determine the grade level of the material appropriate for the student. There are several well-documented methods of classifying text into grade levels (Flesch, 1974) that reflect the difficulty of reading material. The lower the grade level of

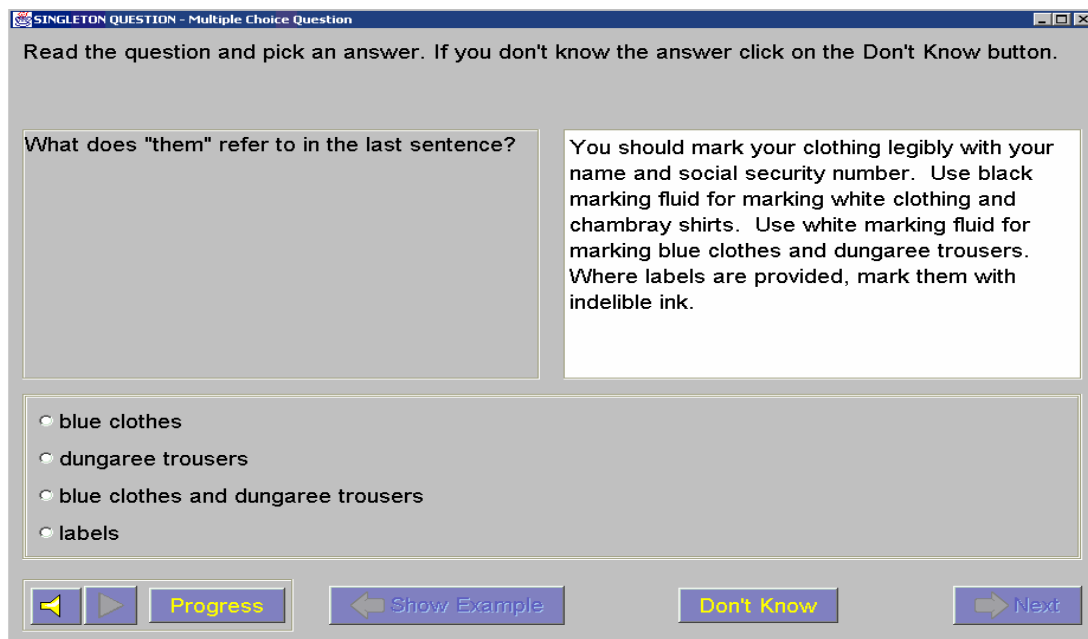


Figure 2: An example of a remedial exercise

some reading material, the easier it is to read and comprehend. For each person, there is a certain grade level below which he/she can read easily without any assistance. This is his/her *independent reading level*. Grade levels at which a student can read with assistance is called his/her *instructional level*. Anything more difficult than this would simply frustrate the student. The tutor uses the placement tests to determine the initial instructional level of the student and to form an initial student model. Furthermore, the tutor continues to update the student's instructional level throughout the course of instruction.

2. Targeted remediation: Remedial exercises are targeted to remediate specific skills in which a student is deficient. This information about weak skills is determined through observation of the student's performance during reading practice.

3. Dynamic selection of questions: The ITS is designed to improve a student's weak skills by building upon his strengths at simpler skills. This is realized through dynamic sequencing of questions that the student is required to answer as a part of reading practice exercises. The student is first presented with questions that involve lower-level skills. The questions gradually increase in complexity to tap into more complex skills that build upon lower-level skills. Such sequencing of questions helps students establish a basis of comprehension, and build upon it progressively. In this way, students build towards questions that they could not have answered had they been presented first. Consider the passage and related questions shown in Figure 2. A student may not be able to answer Q4 directly, but may arrive at an answer to the question when presented with questions Q1, Q2, and Q3 first. ReadOn uses the student model to determine an appropriate sequence of questions for a given student.

4. Socratic tutoring: When a student answers incorrectly, rather than providing some feedback or giving another try, ReadOn asks follow-on questions in a Socratic style to lead the student to the correct answer. Such type of follow-on questioning is intended to help students understand how more complex comprehension skills can be handled step-by-step. Referring to the example in Figure 3, consider a scenario where a student is presented with question Q4 (without being presented with Questions Q1 through Q3) and is unable to answer it. In this case, the tutor would present Q1 through Q2 as follow-on questions and prompt the student to use this information to answer question Q4.

England and most of Europe were monarchies, which means that they were ruled by a king or queen. These monarchies ruled by "divine right" - the belief that God had given them the right to rule. The power of government came from God to the monarch. The people existed to serve the monarch and the government. As a representative from God, the monarch's job was to take care of the people. The people were responsible to God and the monarch. In short, the power of government came from God and resided in the monarch, and the people served the government.

Q1. Who ruled over the people in a monarchy?

Q2. What was the role of the people in a monarchy?

Q3. Who gave a monarch the right to rule?

Q4. Is it true that everyone was considered equal in a monarchy?

Figure 3: Example Text

Student Modeling

The tutor uses Bayesian inference to maintain a student model. The network of reading skills, shown in Figure 1, is mapped onto a Bayesian network. Bayesian networks [Pearl, 1988] provide a formalism for representing probabilistic knowledge. In general, Bayesian network is a directed acyclic graph, whose nodes represent to random variables. The links in the network represent dependencies between variables, such that two variables are assumed to be independent of each other if there is no undirected path between them, or if any of their common ancestors are instantiated. Associated with each node is a conditional probability table (CPT), which gives the probability of each value of the variable given each possible combination of values of its parent nodes.

Figure 4 shows a fragment of the Bayesian network representing the corresponding skills from the skill hierarchy. Referring to Figure 1, the skill "*Constructing Referential Representation*" can be thought of as a composite skill with sub-skills "*Co-reference resolution*", "*Familiar Topics*", and a component of constructing referential representation that is independent of the previous two sub-skills. The node in the student model labeled "*Constructing Referential Representation Prime*", in fact,

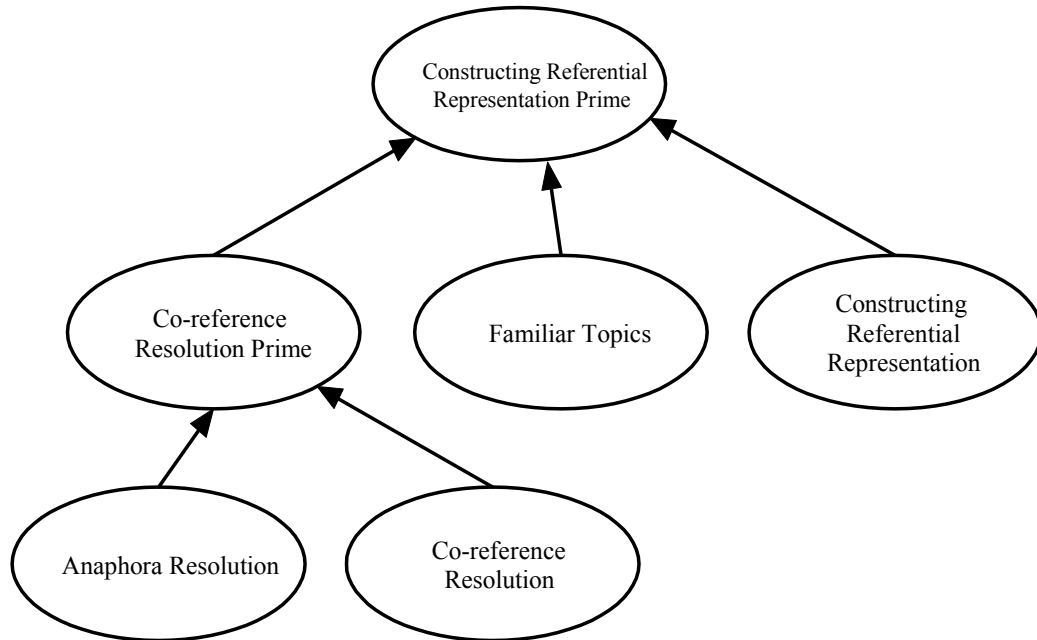


Figure 4: Part of the Bayesian network used for student modeling

represents this component skill. For such skills, the “prime” node is used as a representation of the skills strength/weakness. When a student responds to a question associated with these skills, it is the “prime” nodes that are updated with the resulting evidence.

The probability of each node represents the probability that the student is strong in that skill. Each link from a parent node to a child node has a parameter that represents the probability that the parent node is strong given that only the child node is strong. These parameters are initialized so that for a parent node with links to n children, each link is initialized to the value $1/n$. These link values are compiled into CPTs for each node. The prior probability of each node is set to 0.5.

When a student responds to a question, the tutor determines the likelihood that the response reflects the strength or weakness of the associated skill. This likelihood is a heuristic value that takes into account the likelihood of getting the correct answer by chance and the number of other skills associated with the question. This likelihood is assigned to the corresponding skill node as evidence of its strength or weakness for this case which is then propagated throughout the network using a Bayesian inference algorithm to arrive at an aggregate estimate of the student’s strengths and weaknesses.

Overview of the Authoring Tool

An important feature of ReadOn is the authoring tool that allows organizations to completely customize the content to suit their needs. The version of the tutor developed for the Navy incorporates reading passages extracted from the Basic Military Reading documents which are required reading for new recruits. Thus the content is highly customized to the needs of the Navy so that trainees using this tutor can learn material relevant to their jobs while improving their reading skills. Other organizations that do not find this material suitable can replace the content entirely on their own using the authoring tool. No programming skills are necessary to create new content.

Figure 5 shows the high-level architecture of the authoring tool. The authoring tool facilitates the following:

1. **Entering the text that will serve as material for reading practice:** Authors can enter the passage for integrated lessons, the questions associated with the passages, and the remedial lessons including skill explanations, skill examples, and practice exercises.
2. **Categorizing the text according to its topics:** The authoring tool provides a way to categorize the integrated lessons into topics.
3. **Determining the reading level of the text:** The authoring tool includes two reading tests to automatically

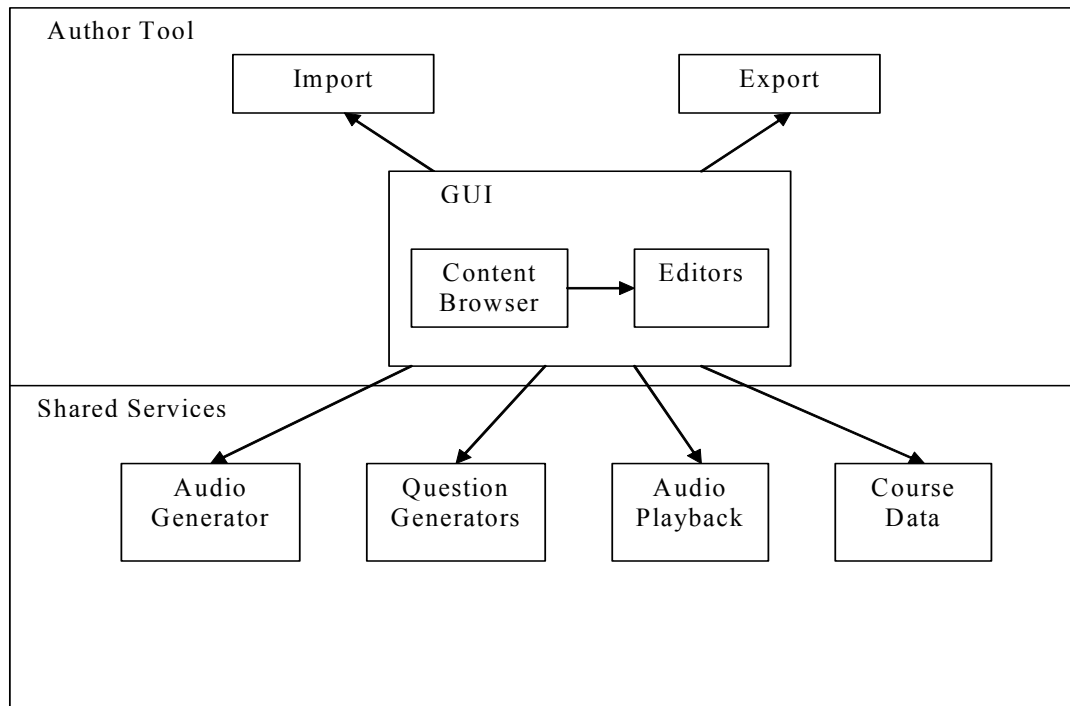


Figure 5: ReadOn Authoring Tool architecture

determine the grade level of reading passages: the Flesch test and the SMOG test (Flesch, 1974; <http://www.med.utah.edu/pated/authors/readability.html>).

4. Indexing questions according to the skills that they test: The authoring tool provides a way for authors to write questions and index them according to the skills that they test. Each question can be associated with multiple skills.

5. Generating questions to test student's comprehension of the text: The authoring tool includes a tool for automatically generating vocabulary questions. This uses data from the WordNet (Fellbaum and Miller, 1998) database to generate question templates that the author can modify and include as a question in an integrated skill lesson. Given a passage, the question generator identifies the difficult words in the passage, and generates a list of sample sentence fragments that illustrate the different senses of the word. The sentence fragments can be used to generate vocabulary questions that require the student to distinguish between the senses of a word.

We will now describe each of the components of the Authoring Tool architecture (Figure 5).

Content Browser

This is the main interface for the Authoring Tool application. All current course content is accessible from this component and is available for editing. New course content can be developed using the Editors associated with

the Content Browser or can be imported using the Import component. The Export component can be accessed to store course content in a text form that can conveniently be imported by someone using the Authoring Tool at another location.

Editors

The editors allow content authors to create new course content. Currently available editors are the Multiple Choice Question, Fill in the Blank Question, Hot Text Question, and Passage editors. Figure 5 shows the Passage Editor. These editors are invoked by the Content Browser to edit or create the associated type of course content.

Import

The authoring tool can import plain text files that are formatted with markup tags. This can be used to facilitate course authoring, since importing the content takes only a couple of seconds. Often it may be easier to develop the content in a textual format for import at a later time. We have found this feature useful for adding content developed by content developers not inside our organization. The Import component parses structured text files and incorporates the resulting content into the Course Data.

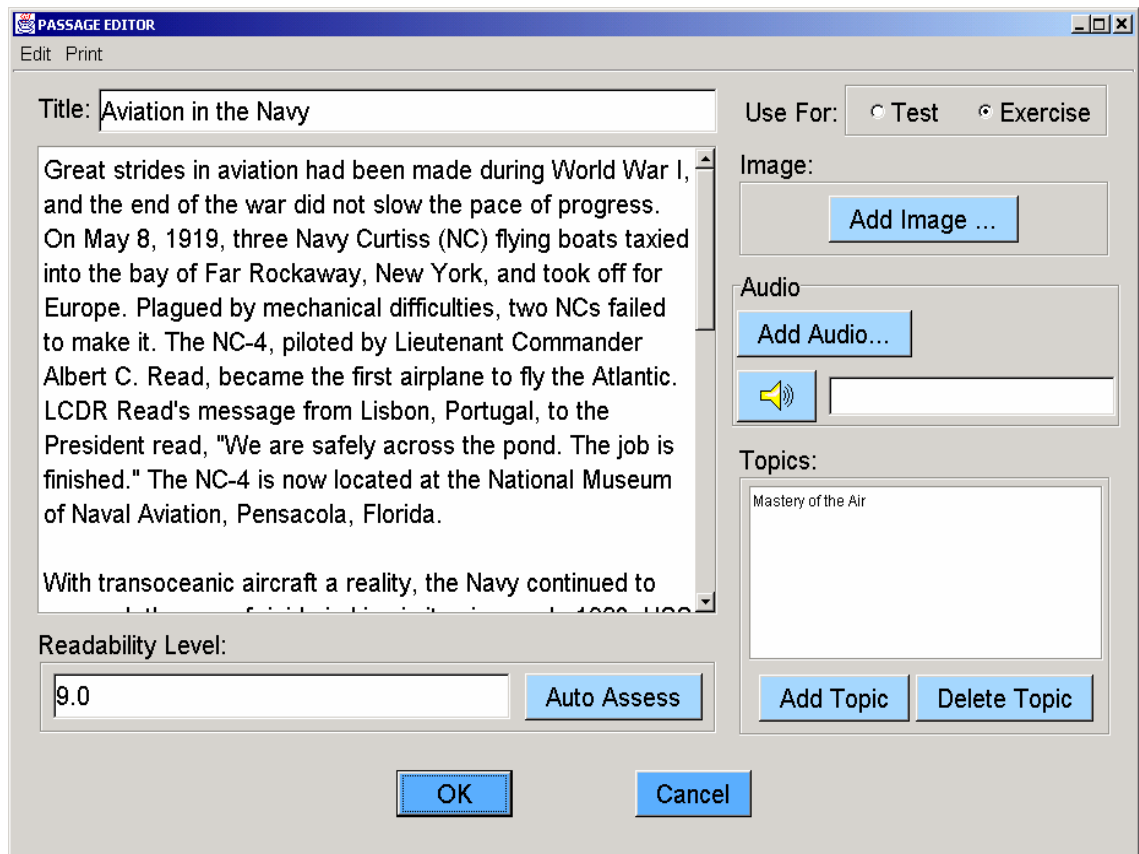


Figure 6: Passage Editor

Export

The Authoring Tool allows an author to export all or any subset of a course's content to text files. These text files can be imported into the same or a different Authoring Tool at a later time using the import feature. This is convenient for exchanging content between content developers, and useful when the content of a course is developed by multiple authors. Each person can export the content that they developed to a text file and all of these exported text files can be imported to create a unified set of courseware. The Export component formats current course content into a structured text format that can be subsequently imported using the Import content.

Audio Generator

The Audio Generator component converts text to speech. It is used by various editors to generate audio files that are associated with course content. The system makes use of the Fonix Text-to-speech system to generate speech, but will operate without speech generation capabilities if Fonix™ is not present (www.fonix.com). Additionally, another text-to-speech generation system could be swapped in with minimal change to the software.

Question Generators

Question Generator components analyze sections of text and generate questions based on its content. Currently, the system has vocabulary and anaphora resolution question generators that aid the content author in the creation of questions pertaining to these skills. The vocabulary question generator searches the text and compares each word to a database of words. The words that have the fewest "senses," or potential uses, are selected as possible vocabulary words. The anaphora resolution question generator parses the text into parts of speech and then finds noun phrases starting with the word "that." These phrases are extracted and used to form simple questions of the form, "What does the word 'that' refer to in the following sentence?"

Audio Playback

This component is used by multiple applications, and plays sound files that are associated with the course content. This component does not generate speech, as that is the domain of the Audio Generator component.

Course Data

Course Data is created using the Authoring Tool and consists of passages, skills, the questions and multimedia files that are associated with them. Content is split into four

instructional types: Pre-Test, Passage Lessons, Specific Skill Lessons, and Skill Examples. The Pre-Test and Passage Lesson types are organized by level. Within each level there are an arbitrary number of reading passages, each with associated questions. These questions are additionally indexed by the skills that they address. Skill Lessons and Skill Examples are organized by skill and contain questions that pertain to a single skill.

SUMMATIVE EVALUATION

We conducted a final evaluation towards the end of Phase II to study the effectiveness of the tutor in improving reading comprehension skills. We installed the tutor at the Naval College in San Diego, CA and recruited some volunteers to use the tutor for a few hours over a period of two weeks. Only a small number of students volunteered to participate in the study and we were unable to get sufficient time commitment from them to lead to any significant results. Although we could not draw any significant conclusions about the effectiveness of the tutor from the study, a survey at the end of the study indicated that the participants reacted favorably to the tutor. We obtained the following scores for the questions in the survey (Table 1). These results indicate a positive response to the tutor.

RELATED WORK

A research group at Carnegie Mellon University has developed a tutoring system for teaching children how to read (Mostow and Aist, 2001). The tutor provides students with passages for reading aloud, and helps them when needed by reading words out aloud. This software is innovative in providing highly accurate speech recognition that can recognize when a student is having trouble with a word and provide assistance. This tutor has been field-tested in schools and found to be effective. This work is complementary to the ITS reported here as it restricts itself to the issue of teaching print-decoding skills.

Intelligent Automation Inc. has also developed a system for teaching reading comprehension skills to adults (www.i-a-i.com) that presents students with passages to summarize, and critiques the summaries automatically using natural language processing techniques. However, this system does not attempt to assess or remediate the particular skill deficiencies that may get in the way of generating good summaries (due to lack of comprehension, for example). The reported system is more focused on diagnosing and remediating particular reading comprehension problems encountered by students. The two approaches again are complementary.

Table 1: Survey results

Questions	Weighted score
How easy or difficult was it for you to understand the tutor's interface? (Scale: 0 – very difficult, 5 - very easy)	3.7
How confusing or easy were the instructions given by the tutor? (Scale: 0 – extremely confusing, 5 – extremely clear)	4.2
How easy or difficult was it for you to understand the performance feedback that you received from your tutor? (Scale: 0 - Very hard, 5 – very easy)	3.9
How effective was the tutor in helping you understand the passages you had to read? (Scale: 0 – extremely ineffective, 5 – extremely ineffective)	4.2
Overall, what was your reaction to the tutor? (Scale: 0 – extremely unfavorable, 5 – extremely favorable)	3.7
In general, how useful or useless do you think this tutor will be for Navy personnel? (Scale: 0 – extremely useless, 5 – extremely useful)	4.1

LESSONS LEARNED

The summative evaluation study reinforced the impact of student motivation on learning outcome. Students join training and educational programs for a variety of reasons. Some do so out of an intrinsic motivation for self-improvement while others try to educate themselves out of a desire for career advancement. Yet others join training programs because they feel pressured to do so by their supervisors. Differences in motivation can have a significant impact on the success of a training system. We found that students with strong intrinsic motivation stayed engaged with ReadOn longer than those with weaker intrinsic motivation. To keep the attention of latter kinds of students, the training system would have to include other motivational techniques like games or explicit extrinsic rewards. The problem of tailoring instruction to different motivational levels has not received much attention in ITS research. We feel that addressing this problem is an important aspect of creating highly effective ITSs.

FUTURE WORK AND CONCLUSIONS

We have described an Intelligent Tutoring System to teaching reading comprehension skills. Such tutoring systems are important for providing tools for adaptive, self-paced instruction in reading in situations where sufficient reading instructors are not available to provide individualized attention to students. The associated authoring tool enables organizations to customize the content to suit their training need.

We are currently negotiating with community colleges and high-schools to further field test the ITS. There are several enhancements that can be made to the tutor. Foremost is the development of additional content. Another near-term extension would be an expansion to include print decoding skill. We would also like to expand the activities available to students to include more interactivity and more generative tasks. We could also expand the skill list to include meta-cognition and reference skills. Finally, as discussed above, addressing the issue of motivation will be an important step in the future.

REFERENCES

- Anderson, J. R., Boyle, C., and Reiser, B, (1985). Intelligent Tutoring Systems. *Science*, 228, 456-462.
- Barr, R., Blanchwicz, C. L. Z., and Sadow, M. W. (1995). *Reading Diagnosis For Teachers: An Instructional Approach*. Longman Publishers, USA.
- Fellbaum C., and Miller G. (1998). *WordNet*. MIT Press.
- Flesch R. F. (1974). *The Art Of Readable Writing ; With The Flesch Readability Formula*. Harper and Row.
- Just, M. A. and Carpenter, P. A. (1987), *The Psychology of Reading and Comprehension*. Allyn and Bacon, Inc.
- Mostow J. and Aist, G. (2001). Evaluating Tutors That Listen. *In Smart Machines in Education*, Eds. Kenneth D. Furbus and Paul J. Feltovich, MIT Press.
- Nichols, P., Pokorny, R., Jones, G., Gott, S. P., and Alley, W. E., (1992). Evaluation of an avionics troubleshooting tutoring system. *Technical Report, Human Resources Directorate, Brooks AFB, TX*.
- Pearl J. (1988). *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufman.