Interactive Books

Final Report

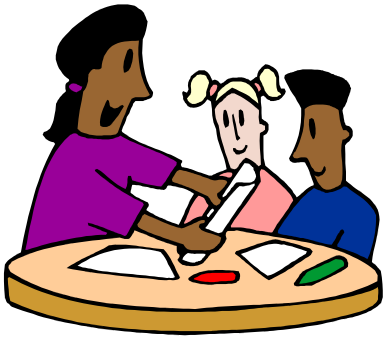

Tim Lloyd, Jon Roeth, Trevor Stone

Sponsor: Sarel van Vuuren, Center for Spoken Language Research

2 May, 2003

Table of Contents

Table of Figures

Overview of the Problem

The Center for Spoken Language Research (CSLR), in conjunction with the Colorado Literacy Tutor program, has developed a program called Interactive Books with the goal of teaching children to read. This program is designed to teach reading, reading comprehension, and knowledge application to K-12 students, as well as provide authoring tools for the interactive books. The Interactive Books provide an interface to a book where a student can interact with an animated agent that can read to the student, listen to the student reading the book, and ask questions of the student relating to his or her comprehension.

The Interactive Books program, in its current design, has several flaws that hinder its educational potential. In particular, the interface between the child and the agent has several issues.

- The current program highlights each word on the page as it is read by the student, but that is the limit of the program's feedback during read-aloud. Additional functionality is required, such as the agent suggesting a pronunciation for a word if the student hesitates and allowing for different methods of indicating the words being read and any problems with words.
- When a student clicks on a word in the Interactive Book, he or she can choose between several different scopes relating to that word, i.e., pronouncing the word, sentence, paragraph, or page, recording the student pronouncing the word, or comparing the student's pronunciation with the correct pronunciation. That interface is not very intuitive and has been redesigned.
- Additionally, when a student is reading to the Interactive Book, he or she must hold down the "Alt" key to activate the audio interface between the client computer and the server, and release the key to have the program stop recording/listening. This interface has been completely redesigned.

We believe that our redesign of the Interactive Books program to correct these design issues enhances the student's learning experience and makes the program a more effective reading tutor.

Literature Summary & Related Products Review

Designing With/For Children

Since the Interactive Books are designed specifically for children learning to read, consideration was given to receiving input during the design process from children in the target user group. However, for our project we decided to focus on designing with teachers *for* children. Druin et.al. (1998) concluded that control, social experience, and expressive tools were what children want in technology experiences. While expressive tools (being able to express sound, visuals, movement and physical appearance in technology) lend themselves to the content adding/editing portion of the Interactive Books interface that we did not work on,

integrating control (children are empowered when they feel in control of their environment) and social experiences (children want to work with other children) would have been helpful.

A consideration for testing with children is making them feel at ease and be able to get think-aloud data, but considering the nature of the think-aloud method, this may be problematic for use with children. An alternative is offered by Kahler (2000) called Constructive Interaction for Testing Collaborative Systems (CITeCS). The CITeCS method creates an environment where two users use two different workstations in close proximity to work together through a problem. The users each have different problems to solve, but the problems are related and the users can discuss their own tasks with the other as well as see the other's computer. Adapting this method to working with children, we had hoped to have two children use the same computer to use the Interactive Books program together. It was hoped that we would be able to hear the children discuss the interface and what they are doing thereby allowing us to collect think-aloud data, but reducing the stress and requirements on the children. However, we were not able to test our designs on students, for reasons explained later in this report.

Soliloquy Reading Assistant

One of the more desirable features we found in Soliloquy is its "How Am I Doing" button, which highlights each word that the student has read in one of three colors: green for "great," yellow for "okay," and red for "try again." The user clicks on a button labeled "Read & Record" to have the program begin listening to the user read the story. A "Stop" button stops the recording and listening. This setup is one possible solution to our need to replace the Alt key trigger to begin and end listening. A "Hear Me" button plays back what the user has read to the program, highlighting each sentence as it is read back. The program is very colorful, which we believe to be a good design for children. When the mouse rolls over a word, it is changed from black to blue to indicate that the user can click on the word to have it read by the program. While the program is listening to the user reading, only the current sentence is in black type, while the surrounding text is grayed out. As each word is read, it is grayed out by the program. If the user halts on a word, the program said the word or the first syllable of the word. If the user struggles longer, the program has the user re-start the current sentence. Finally, Soliloquy implements a volume meter in the top left corner of the screen as a yellow bar that grows or shrinks with the user's speaking volume.

Living Books - www.livingbooks.com

The Living Books examples do not offer much help in the way of design solutions. However, the program has a simple "Play" and "Stop" set of buttons to start and stop the audio reading of the text. We thought that the standard "Play" and "Stop" icons as found on any VCR or stereo component would be intuitive for most K-12 students. The Living Books also utilize small animated graphics, both as illustration for the text and as part of the page layout. Finally, we liked

the content of the example book, “The Silly Sea,” which was educational while being lighthearted and fun.

Corporate Comics - www.corporatecomics.com

To solve the problem of printing the words the animated agent is saying as words in addition to the audio output of the words, one idea was to use the dialogue representation that comic books use. This is explored in the example Flash animation at the Corporate Comics web site.

Summary of Initial Interviews

The biggest pieces we learned from our interviews is the need for flexibility. Students with different reading skills, from essentially none to expert, need to be able to use the system. Students, with teacher guidance, need to be able to focus on particular aspects of reading, such as pronunciation or comprehension. An extension of this need for flexibility is that students who can’t read well need to be able to use the system. This necessitates primarily visual or spoken interface elements.

Another thing our interviewees stressed was the need for the system to “just work” – otherwise, students will get bored or play around with ancillary parts of the system (which they’ll do even with a system that works right).

We also learned some techniques teachers currently use to teach kids to read. Pictures play a big role, since kids can learn what words are likely to appear. When a student struggles with a word, teachers may encourage them to sound it out or see what it might be from context. They may also suggest a pronunciation, but it’s important to let students try to figure it out on their own. Teachers also ask students to answer questions about what they’ve read. Positive feedback is important for students, and teachers keep track of what kinds of problems students have and use this knowledge to plan future readings.

Since the program is targeted at students at all grade and reading levels, it is difficult to put together an abstract user profile – very little is shared between students reading in kindergarten and students reading in high school. Our User Profiles section outlines relevant information for several types of students.

Key Design Priorities & Issues

Our project has a somewhat unusual set of design priorities due to the nature of our user group. Although the program is targeted at all grade levels from K-12, we decided to target our design at K-2 students, primarily because teaching students to read is the main focus of this program.

The interactive books feature we worked on was the read-aloud interface. In this part of the system, students are presented with text. They can read aloud to the system and learn how well they perform. They can have an animated agent read all or part of the text to them. They can learn the pronunciation of words they don’t recognize. They can also have the system test their comprehension with quiz questions, whether they have read aloud, read silently, or listened to the story. The interactive books should provide most if not all feedback mechanisms a reading teacher would provide.

At the most general level, the interactive books must work smoothly. While this is a priority for most systems, it is especially important in systems designed for kids, because they don't have the patience or the focus to try to learn the "proper way" to use a system. This means we need to make sure controls and information are as obvious and intuitive as possible from a child's perspective.

Flexible operation is another broad priority. Not only do different students read at different levels, students may have different reading needs and styles. Some students tend not to read aloud, but have excellent comprehension while reading silently. Other students may just want to look at the pictures or play with the animated agent. Different teachers may want students to focus on different reading tasks. One annoyance we found in Soliloquy was its requirement that students read each word in the story aloud before they took the comprehension quiz.

Another priority for students learning to read is for information to be available in multiple formats. For instance, if the system suggests a word pronunciation, the student should be able to hear the word, see the agent's lips move, see the word written, hear how he pronounced the word, etc. This information doesn't always need to be presented, or to be presented simultaneously, but it is all helpful for students learning to read.

Since many of our users will not know how to read yet, we need to rely on images and spoken words for most aspects of the user interface. Furthermore, we must carefully consider the extent of students' vocabularies when we do use written words in secondary capacities or for features used by better readers. Do students know what "paragraph" means?

The interactive books present a lot of information simultaneously, and one of our challenges will be developing ways for students to take in all of this information without being overloaded. For instance, while a student reads a text, the interface will show her progress. She will also see the story's illustrations, and may refer to them for clues about what words might be in the text. The animated agent may provide suggestions when the student is stuck on a word, so that talking head needs to be visible to a student reading.

Another set of issues emerge from the existing system with which we interact. The operation of the speech recognition and synthesis software and the animated agents is well outside the scope of our project. While these systems are robust, they restrict he feasibility of some interface solutions. We probably cannot, for instance, listen continuously to the student and try to detect when she starts reading at a random place in the text. We therefore needed a way to mark where and when reading begins. Since we are not experts in the existing systems which form the context for our project, we checked the technical feasibility of our solutions with folks from CSLR.

User Profiles

In our discussions with members of CSLR, several user types have emerged. They can be broadly divided into students at several reading levels, as well as teachers.

Teachers: While teachers won't often use the read-aloud functionality which is the focus of our project, they cannot be ignored as a user group. Teachers need to know how the system works so that they can help students who encounter troubles. Teachers also need to be sure that the interactive books are teaching students what the teachers

want them to. Furthermore, the teachers need to be able to assess student ability and progress. Finally, teachers will create interactive books, though the editing interface is out of the scope of our project; CSLR has plans for significant enhancements to this part of the system.

The computer skills of teachers vary wildly. Some teachers are expert computer users while others have trouble with basic tasks. The teachers in our user tests are familiar with CSLR's reading tutor program, which shares some features and elements (such as the animated agents) with the interactive books. We can assume that all elementary school teachers have training and experience teaching students to read. Furthermore, teachers usually have a good idea how well each of their students can read and what they need to work on. Some teachers will be experts in helping students with reading difficulties. Secondary school teachers will have particular domain foci and expertise. Their focus will be on concept mastery, rather than reading proficiency.

"Ordinary" Students: This user group includes students who have few if any problems reading. Students within a classroom tend to read at different levels, and each student has a different set of areas in which he excels or is challenged. CSLR's goal is for interactive books to be used by students in every K-12 grade level, and that they should be presented with material appropriate to their skills. Students at the kindergarten level are likely to be familiar with letters and sound, but reading whole words is a challenge. Students a few grades later are able to read most words and follow somewhat complex narratives, but may have systematic or sporadic trouble with some words or concepts. Students in secondary school can read well, and are learning to analyze text, learn specific concepts, and induce facts about the material.

Students vary in their ease and comfort in reading aloud. Students may read quietly or loudly. They might read faster than the recognizer can follow along or they may pause between many words. Students may have an accent or odd patterns of speech. Some may not want to read aloud, or may not want to read at all.

Even among "ordinary" readers, attention and focus may be a problem. Some students are able to read by themselves for a long period of time while other students may easily stop reading and start messing around with parts of the system, playing with other kids, etc. Some students need a lot of feedback about their performance while others are able to read without much assistance. We will assume that students in this user group are able to produce and understand spoken language reasonably well, though there may be significant issues like accent or shyness.

Students have a wide variety of computer background. Some use computers at home every day while for others, interactive books may be their first encounter with a computer.

ESL Students: This user group includes students whose native language is not English. They may or may not be able to read in their native language. When teaching an ESL student, teachers make sure the kid understands the important concepts in her native language, often using pictures, and then connects that concept to its English counterpart. ESL students often mix English words with non-English, or apply grammatical rules from their native language to English words.

Special-Needs Students: Many students have fundamental difficulties learning to read due to a variety of physical, neurological, or psychological problems they face. Hearing-impaired students, for instance, have poor feedback on their own pronunciation, so it can be hard for them to match example locution. Deaf students also have trouble "sounding out" words they don't recognize. Other students may have trouble with visual recognition of letters and words or may easily forget what they have read so far.

It is difficult, if not impossible, to characterize all special-needs students. We are not experts in reading challenges, and neither we nor CSLR are approaching this project by producing a list of reading disabilities and targeting them. In our discussion with Scott, it seemed like what special-needs students need is techniques that are already provided to "ordinary" students by the interactive books system. Special-needs students just need more focus on one technique or heavy practice on certain aspects of reading.

Task Statements

Task 1: Jimbo Jones's kindergarten teacher, Ms. Krabappel, wants him to listen to a story and follow along with the words as they're being read. Jimbo is just learning to read, but doesn't have a very large vocabulary, so Ms. Krabappel wants him to learn a few new words while listening to the story. She wants him to answer a few questions about these words after listening to the story so that she knows how well he has learned them.

Task 2: Ralph Wiggum is a second grade student who has trouble pronouncing several words when he reads them aloud. Ms. Krabappel wants him to read a story aloud, and then have him hear the words he mispronounced in his own voice along with the correct pronunciation. She wants him to practice these words in this manner until he gets them right.

Task 3: Pedro Chespirito is an ESL student in third grade. He speaks and reads Spanish at a third-grade level, but only reads English at a kindergarten level, and often swaps Spanish words for English when speaking. Ms. Krabappel wants Pedro to practice reading a story with lots of pictures to help him identify concepts and match them with the words he is reading. Ms. Krabappel wants to know which specific words he has trouble reading aloud.

Task 4: Lisa Simpson is a second grader who reads at a ninth grade level. Her teacher, Ms. Hoover, wants her to read *The Pit and the Pendulum* and write about Poe's use of imagery and metaphor. Ms. Hoover wants to know roughly how good the essay is without reading it, but she would also like to read a copy.

Task 5: Mr. Hawking, a high school teacher, wants his 11^{th}-grade student Carl Sagan to read the relativity tutorial Mr. Hawking put together. Mr. Hawking wants Carl to answer Mr. Hawking's prepared test questions such as "Which of the following is not an inertial reference frame: a) A train traveling at 50 mph, b) The moon, c) A rocket ship blasting off from Earth" and "Selma travels around the Earth repeatedly at .9c. Patty

measures that one year has gone by, according to her clock. How much time has elapsed on Selma's clock?"

Key Issues for User Testing

One of the primary problems in using a think-aloud method with children is that the method is too stressful for children, and children often are not capable of thinking aloud about the task which they are performing. A method to solve this problem involves observing two children working together, and recording their dialogue and interaction between themselves and the computer.

As we found during this project, scheduling interviews with teachers was difficult at best and impossible at worst. We were unable to test our designs with any BVSD teachers due to several scheduling issues and unforeseen circumstances at the target schools. This problem naturally extended to testing with students, although this was not a concern given the level of fidelity of our final prototype.

The Interactive Books client-server interface is currently optimized for a 100BaseT network, and performs best on this type of network with high-end computers. This technological requirement is significant for testing with children, as they do not tolerate faults in the system, and will move on to another activity upon encountering what they perceive to be a problem with the program.

We have also run into problems with the preferred Interactive Book, "The Backyard Zoo." The speech recognition begins to falter past the fourth page, as does the "Next" button. We had hoped that this issue would be fixed before testing with students so that we would have had at least one complete story to test with the students. Ideally, having a selection of Interactive Books would be preferable to allow the students to read at their desired reading level and feel in control of their reading experience. As mentioned above, however, we did not test with students and in fact did not require a fully functional Interactive Book program for a significant portion of our design process.

Existing Design

Figure 1 shows the existing design as of the start of this project, created by CSLR designers including Nattawut Ngampatipatpong and Jariya Tuantranont.

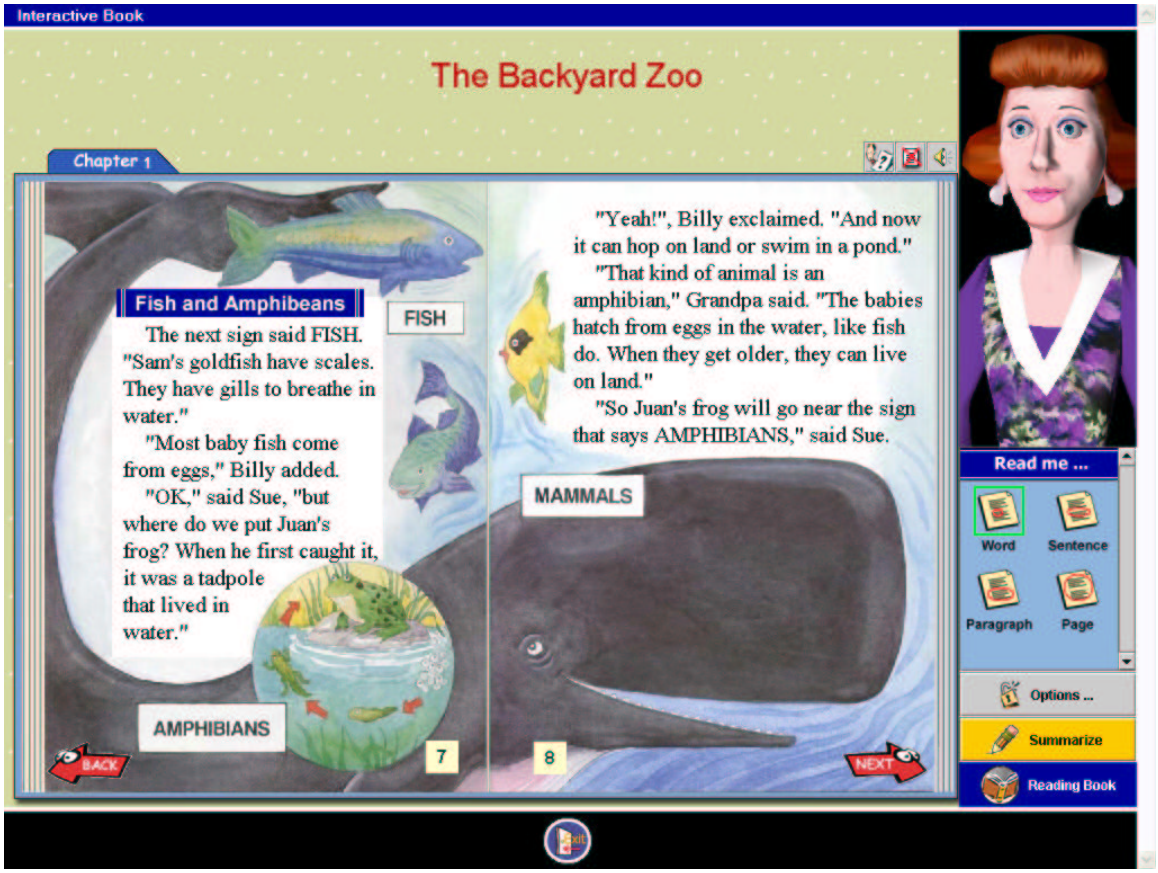

Figure 1. Existing design.

Heuristic Evaluation of Existing Design

Below we have listed several results of our heuristic evaluation of the usability of the existing Interactive Books program. The particular heuristic for each result is noted in italics.

- Kids don't know what word/sentence/paragraph/page distinction is. For some kids books, the distinction is ill-defined anyway. *Match between system and the real world.*
- Top-right buttons – what do they do? We discovered turning on/off captioning by accidentally clicking one. They seem more like debugging buttons than final design buttons. The kids will click on them, because they'll click on anything, so keep debugging off the screen. *Recognition rather than recall, Aesthetic and minimalist design.*
- The entire bottom row of the screen is used for a single exit button. *Aesthetic and minimalist design.*
- The top of the screen isn't used for anything (well, almost). *Aesthetic and minimalist design.*

- Currently, you must hold down the alt key to have the program listen to you read. This is clearly not a production solution. *Recognition rather than recall, Consistency and standards, Flexibility and efficiency of use.*
- What does "Reading Book" mean? It switches between reading and editing, which shouldn't be available to kids. *Visibility of system status.*
- User shouldn't be able to waste time trying to look at the monkey agent's butt. (Don't let user set agent/display options or fiddle with agent orientation.) *User control and freedom.*
- Next/Previous page buttons are intuitive and consistent. *Consistency and standards.*
- Words only highlighted on recognition – no indication of how well the word was pronounced. *Visibility of system status.*
- You must turn the page to take the quiz and you must take the quiz to turn the page. If the student has already taken this quiz, and just wants to reread the chapter, she shouldn't have to take the quiz again. Furthermore, you may want to take the quiz and, if you perform poorly, reread the page. *Flexibility and efficiency of use.*
- First page doesn't appear when you load a book, instead a "Click Here to Begin." *Match between system and the real world.*
- Agent lets student struggle indefinitely with a word until the student takes action to click on the word. *Help users recognize, diagnose, and recover from errors.*
- Book author should specify whether or not to include a Summarize button? Some books won't be integrated with summary. *User control and freedom.*
- All of the buttons are targeted to an already-literate user, defeating the purpose of the program. *Recognition rather than recall, Flexibility and efficiency of use.*
- Looks like a real picture book. Good! *Match between system and the real world.*
- Agent is fairly large and holds a position similar to a teacher holding a book and reading to the class. Good. *Match between system and the real world.*

The most important issues found in our heuristic evaluation were the word/sentence/paragraph/page distinction and the lack of a good way to start the "Listen to Me" function. Less important issues were the various unnecessary buttons, use of screen real estate, agent re-orientation, and the link between turning the page and taking the quiz. The good features included the book-like visual design and the size and position of the agent.

Low Fidelity Mockups

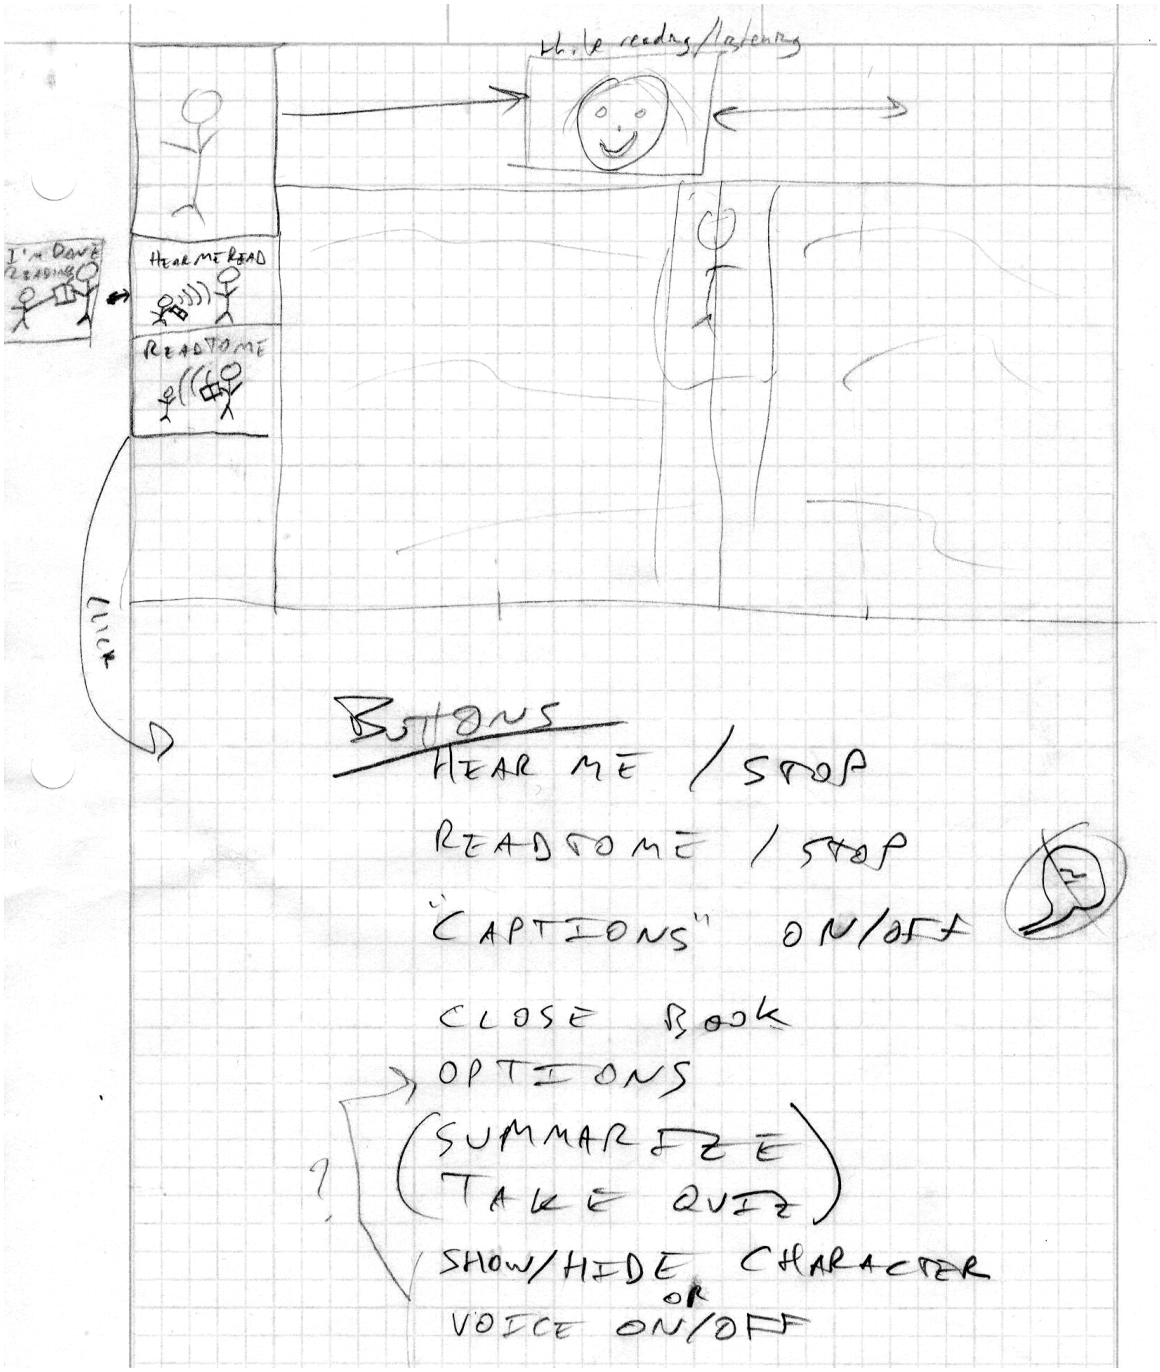

Figure 2. Initial sketch.

Figure 2 shows the initial sketch used as a starting point for the below mockups. Some simple, graphical buttons are drawn on the left, and various positions for the interactive agent are sketched out. An early idea was to have the agent move along the top of the book while reading or listening to the user.

Figure 3. Trevor's low-fidelity mockup.

Figure 4. Tim's low-fidelity mockup.

Figure 3 and Figure 4 demonstrate two different possible layouts. Trevor's mockup shows how an illustrated set of pages would be broken up by the agent and buttons placed between them. Tim's mockup illustrates only displaying the bust of the agent, and has two cartoon-style buttons for the "Read to me" and "Hear me read" functions.

High Fidelity Mockup

Below we show various screenshots of our high-fidelity prototype. It is important to note that while the above low-fidelity mockups were non-functional while visually pleasing, the high-fidelity prototype is functional and interactive, but does not incorporate all of the visual features shown later in our final design. This was primarily due to the restrictions of the Perl/Tk widget toolbox, which was chosen for its simplicity and power for rapid prototyping. The purpose of the high-fidelity prototype was to test out the functionality of our design with our users, while showing them the low-fidelity mockups for feedback on the visual layout.

The general concept behind our button images was to depict a representation of the computer, a representation of the student, and an arrow or "sound wave" indicating directionality. For the purposes of the user tests, we purposely did not attach labels to the buttons, but both experts thought that simple labels would be helpful for advanced users as well as the learning to read users in order to give them yet another opportunity to learn words, especially with a graphic association. We think that this underlying concept should be implemented, and we leave the task of devising the exact representations to CSLR.

Figure 5 through Figure 9 below demonstrate the various positions of the agent and buttons with respect to the text. The "text bubble" displaying what the agent would speak is shown at the top in Figure 5 through Figure 7, and beside the agent in Figure 8 and Figure 9. These screenshots were taken in the "Listen to me" function of the book, as can be seen by the stop sign replacing the second button in Figure 5. The text is colored green, yellow, or red where the user has read to the agent, reflecting how well the user spoke the text. Figure 5 also demonstrates highlighting of the text ahead of where the user is reading.

Figure 5. Agent on right.

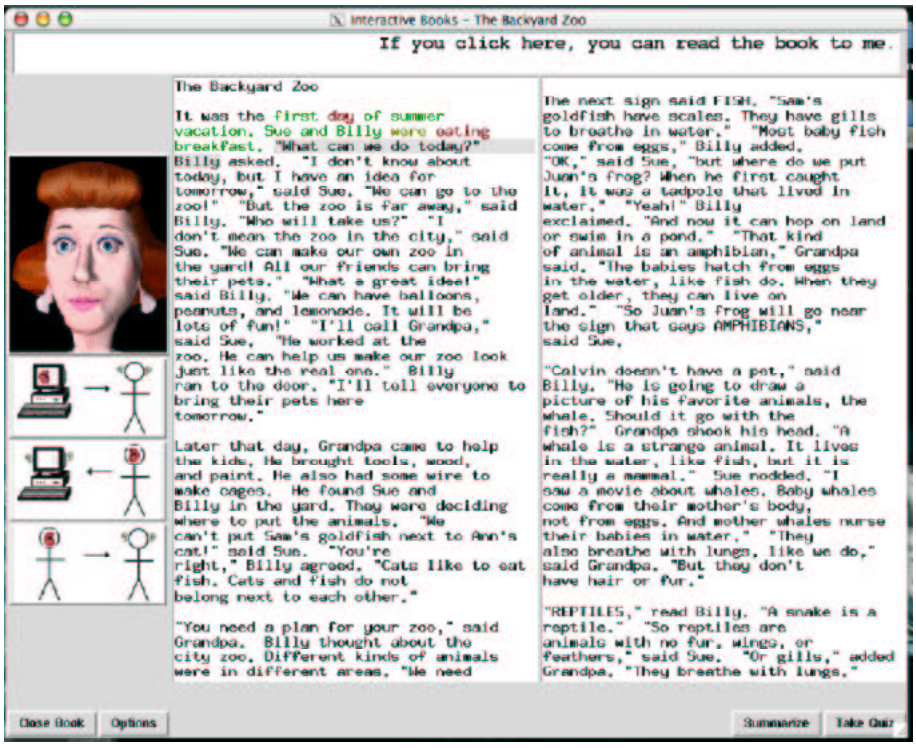

Figure 6. Agent on left.

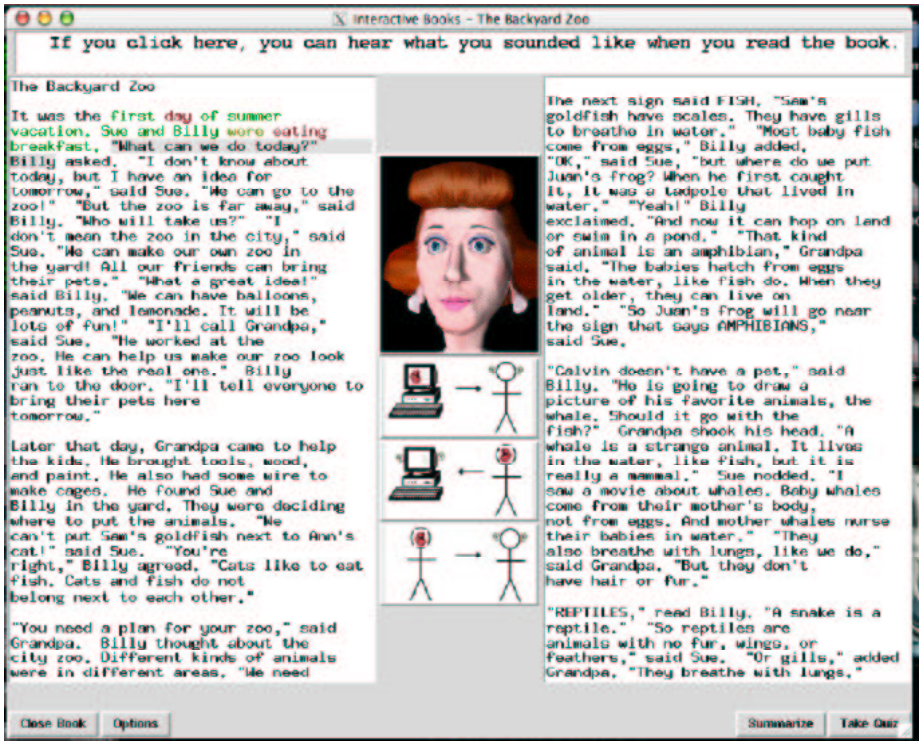

Figure 7. Agent between pages.

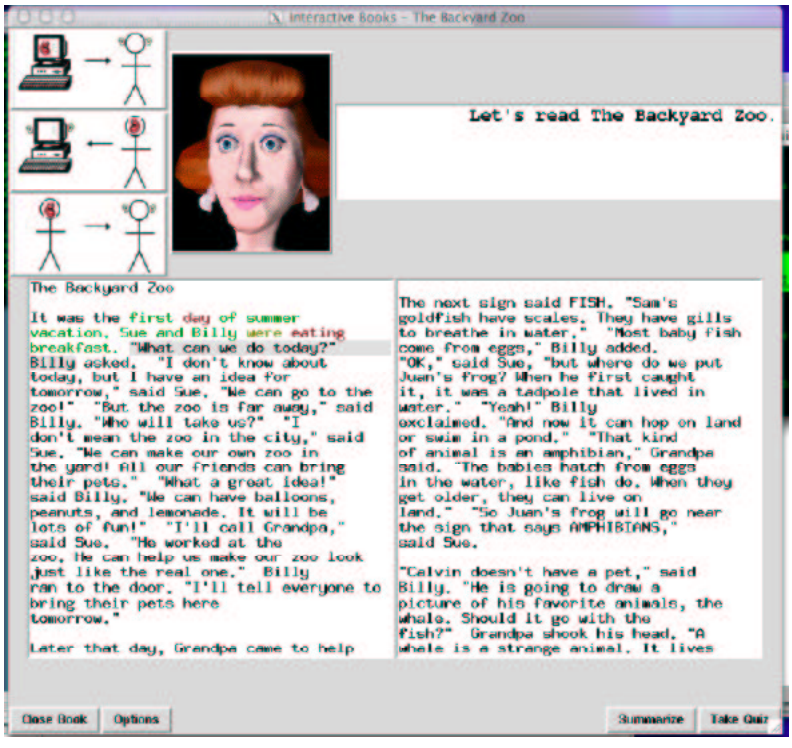

Figure 8. Agent above book.

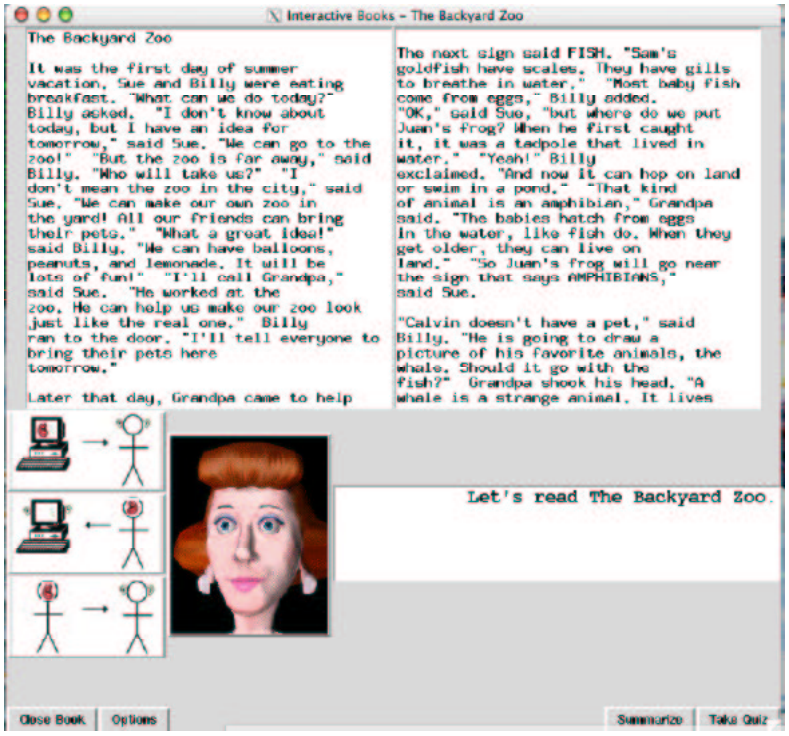

Figure 9. Agent below book.

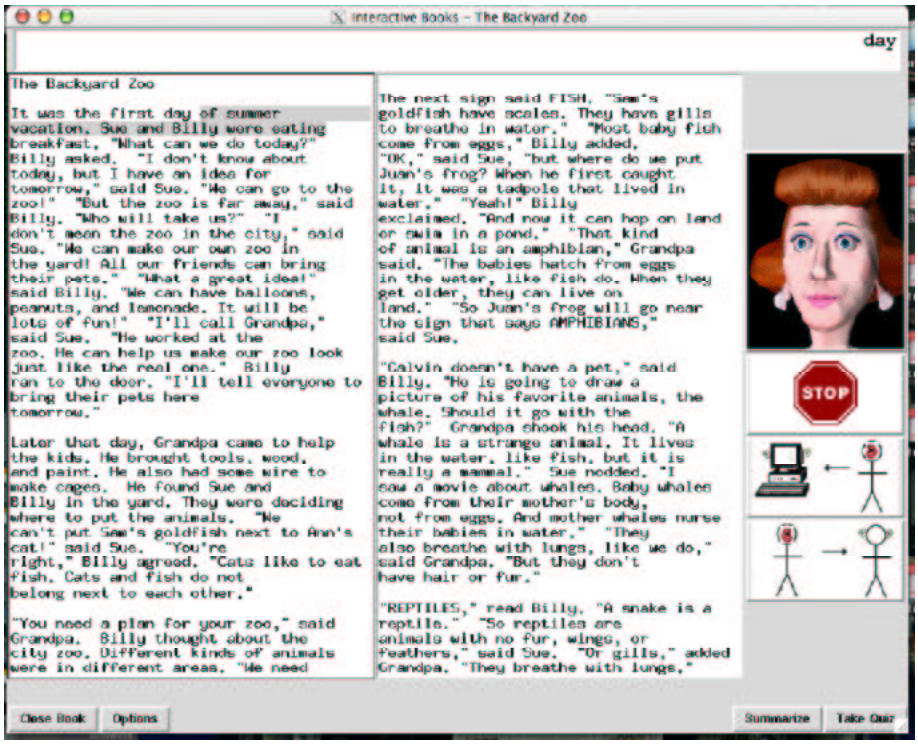

Figure 10. Agent reading to the user.

Figure 10 above demonstrates the agent reading to the user. The agent is speaking the word "day," and the next few words in the text are highlighted in gray. The "read to me"

button has been replaced by a stop sign, which is used to stop the agent reading to the user.

User Testing

Tasks Used in User Testing

- You are a kindergartener and your teacher has told you to have the tutor read to you and try to follow along with the text as the tutor reads. You haven't used the IB before, but "The Backyard Zoo" is open to the first page when you come to the computer.

- You are a first-grader who can read, but not very well. Using the IB, read "The Backyard Zoo" to the tutor. At the end of each sentence, stop and listen to how you read that sentence paying attention to the words the tutor says you didn't do well on and rereading the sentence until you get most of the words correct. The most important thing is getting the words right and learning how to pronounce new words, not getting through the story.

Participant Profiles

Due to external circumstances, we were unable to meet with any of the teachers in the target schools. We instead met with two CSLR education specialists who both spend time working with the teachers and students on the currently-deployed CSLR reading tutors and other reading initiatives.

Participant 1 has a background in linguistics and ESL, as well as database administration and design, and has been working at CSLR for over 3 years. Participant 2 is a speech/language pathologist who serves as a special education consultant for the Boulder Valley School District and is directly involved with the design and implementation of the CSLR reading tutors in the target schools.

Summary of User Testing

Participant 1:

One of the most important things that Participant 1 said is that the students are much more technologically savvy and interested in the program than the teachers are. Some fifth-grade students are currently using the Interactive Books' authoring tools to create reports about science experiments, and understand how to create the books an order of magnitude better than their teachers. She did say that Boulder Valley K-12 students are statistically more likely to have computer experience than K-12 students in general, but overall young children will learn how to use computer programs at an astonishing rate. Also, children will play with the interface for some time, clicking on all of the available buttons, before asking for help. She also noted that children like to have an obvious exit button for when they become frustrated with the program.

Overall, Participant 1 liked the buttons that we chose for the high-fidelity mockup. Having the animated agent provide some sort of introductory instructions when

the book is loaded would be very beneficial. She stated that generic, graphical, internationally recognized buttons are very important. The buttons could be separated from the agent. She liked that clicking on a single word would cause the agent to speak that word, and that we had removed the "word," "sentence," "paragraph," "page" buttons for the agent reading to the student. If the student wants to hear a sentence or paragraph, he or she will click the button to start reading, and then click "Stop" when the agent has read enough. She also suggested that authors may want the agent in a specific place for different books, i.e. in the center of the screen for hearing-impaired students to lip read the agent.

Participant 2:

Participant 2 understood the buttons fairly well before he was informed about the agent's "closed caption" speech bubble that was explaining the purpose of each button. He did note, however, that the icons for the buttons should be more descriptive, and should include text labels. The text labels should be included both for students able to read and to enhance the learning experience of younger students. While he was reading to the computer, he noted that the highlighting of the current sentence was beneficial, as well as being able to start reading anywhere. Being able to click on a word and have it read aloud was good, although he also wanted to be able to highlight text and have it read aloud. Participant 2 also noted that the fifth-grade students creating content in the target schools are five times better at using the software than their teachers.

Placing the animated agent in the center of the screen, according to Participant 2, best matches the real-world scenario of a teacher holding a book in front of himself or herself while reading to the class, although he did mention that the target students are used to the agent on the right of the screen. The placement should not be selectable by the student, since the students will spend time playing with the agent as they did with a previous version that allowed rotating and scaling the agent. He also mentioned that a few students rotated the agent 180° so that it was not facing the student while he or she read to the computer. When offered the option to minimize or hide the agent, however, Participant 2 declined, stating that as the animated agent becomes more expressive and reactive to the students, it will be more desirable to have it on screen.

Recommendations Resulting from User Testing

One thing that seemed important from user testing was a need for flexibility in the user interface. Our test users both expressed the importance of letting an interactive book's author or the teacher decide the layout of a book. For instance, the agent's position in the middle of the screen is valuable for hearing-impaired students so that they can easily see lip movements. However, advanced students might find a centrally-positioned agent too distracting, so an author might create a single "page" of text with the agent small and in the upper-right hand corner of the screen. At the same time, though, our experts thought that most teachers wouldn't spend time to configure the books, so the focus for layout would be on book authors.

Our experts said that students will play with anything that's on the screen, and stressed that no button should be on the screen unless kids were supposed to click it, because they will.

Kids tend not to ask for help – they'll play with it until they're either bored or frustrated and will then leave and do something else. Furthermore, if the program frustrates them and they can't get out of it, they won't use it in the future. Therefore, a failsafe method is needed to ensure that kids can exit the program at any time. Furthermore, this underscores the need for intuitive interfaces – kids won't ask the teacher for help, and may not listen to instruction. Finally, in a typical classroom environment, a teacher may not be able to give tech help even if the student asks.

Our design of the "Read to Me" interface was well-received. Users found clicking on a word to hear it intuitive. They also quickly discovered the way to get the agent to read the book to them. The previous design's method of selecting reading scope was universally despised by our users and they welcomed the new design.

Similarly, the new "Listen to Me" interface received praise. Our users found it intuitive and clear to locate and use. They praised our use of the universally recognized stop sign, which we borrowed from Soliloquy, although we had buttons change to the stop picture rather than having an always-present stop sign.

When selecting the final icon pictures, they should be as simple and universally-recognizable as possible. During design we struggled with how to convey a verb (such as read) along with its subject and object. The buttons we had in place for user testing were too complex, though our adult users could easily determine the meaning of the images.

Although we were unable to test the fine-grained coloration of recognized words, the test users liked the general pattern of coloring recognized words and highlighting the next segment of words to be read by the student. They suggested that when the agent reads to the student, highlighting goes to the end of the sentence. As CSLR continues development, they should explore using different colors to indicate the quality of the student's pronunciation.

The prototype's limitations also elicited some interesting comments from the experts. Because the speech recognition was being manually emulated by Trevor Stone pasting pre-created data into the speech recognizer while a tester was reading aloud, it was nowhere near as fast or accurate as the computer's recognizer. One of the experts pointed this out by saying that if it were that slow or problematic for children it would not be well received. Instead of having the agent/computer actually saying audibly what it would normally be saying, the prototype had a text bubble to show what the agent would be saying. The text bubble would probably be part of the finished interface, but the primary output would be audible and not visual for the agent's spoken part. Also, without the normally rich graphics and pictures that a children's book or the full-featured Interactive Books have, the interface looks like a drab journal article instead of an actively colored child's book.

Final Design

Figure 11. Final design.

The following are the decisions we made about the issues that we encountered and in some cases, just represent what we thought were intelligent defaults.

A guiding principle was simplicity throughout the interface in eliminating unnecessary and/or confusing functionality and keeping the interactions and views as simple as possible. As a result, we removed the word, sentence, paragraph, and page buttons which were used to specify to the agent how much to read to the student. We chose rather to have a simple button to turn on the agent reading, allow the student to choose where to have the agent start, and then let the student tell the agent when to stop. When not in the agent-reading mode, simply clicking on a word will have the agent read just that word to the student. This works even when the student is reading to the agent.

The existing interface required the student to hold down the ALT key to indicate to the agent to start "listening" or recording and release the key to indicate the student was done reading. We realize that this method was probably intended as a stop-gap measure until a better method of initiating the "listen" interface was designed. We provided a simple mouse button to tell the agent when the student was going to start reading. To give credit where it's due, we did like the similar button interface in Soliloquy. Removing the keyboard as a part of the basic interface makes the interaction between user and program much simpler. The only keyboard use now occurs when the student takes a quiz or writes a summary.

As the student reads to the agent, the text is colored based on the reading quality and the text just ahead of the reader is highlighted. The text coloring is green, yellow, and red depending on the quality of the word pronunciation where green is a perfect

match and red is no match. The red coloring is also used when words are skipped. The text highlighting is done in gray to not be too distracting but give greater emphasis and readability to the words being read. The rate of highlighting ahead of the reader is based on how fast the student is reading as determined by the rate of word recognition by the computer. For example, a slow reader sounding out each word and generally progressing carefully through the text would have only a few words ahead highlighted, while a more advanced reader speaking words in phrases would have many more words highlighted ahead.

Part of our design implementation was to allow the agent's position to be flexible in relation to the text as well as to be different sizes. A default suggested by the user testing for the agent was to put it in the center of the screen between the two pages of text. Another advantage of having the agent in the center was that it is more compliant with accessibility guidelines as this is a better placement for hearing-impaired users.

Due to the nature of our program, all of the basic program functions must be usable by a person who can't read at all. Our general solution to this general problem is twofold – clear icons and spoken agent commentary. We felt that the most important buttons to work on were the "Read to Me," "Listen to Me," and "Playback" buttons. The former causes the agent to read the story to the student. The second allows the student to read to the program. The final button lets the student hear how she sounded while reading the story.

The visual challenge here is expressing a verb with its subject and object. Our basic idea is to have a symbol of the computer on one side and of the student on the other with an arrow, sound wave, or other graphical representation of directionality. Based on comments from testers, a promising candidate is a student "smiley face" with ears when the student is listening and with a mouth when speaking while the computer picture is undecorated. (We played with giving the computer ears or a mouth, but these look kind of silly and don't convey much more than the computer alone.) Each button should also have a textual label, both for students who can already read and for students to learn important words like "Read" and "Listen." We left these labels out of our prototype so that we could test the intuitiveness of the icons.

To complement what we hope are intuitive icons, when the user moves the mouse over a button, the agent explains what the button does through audio speech and the speech bubble (described below). This allows for students to learn what buttons to push even if they don't grasp the intended message of the icons. This functionality must be implemented with care, however. The explanation should stop when the mouse moves away from the button so that accidental motion doesn't produce a message that must be listened to. Furthermore, it may be beneficial to only explain the buttons a few times so that kids don't get distracted playing around – getting the agent to speak and stop quickly. The agent will also speak an introduction when the book is opened, directing the user to the read and listen buttons.

Clicking on the agent, which kids will do as a matter of exploration, will cause the agent to speak one of several tips about the program, such as how to read to the agent. We support CSLR's decision to remove the ability to rotate and resize the agent, which distracted students from the task at hand.

In addition to arrows which take the student to the previous and next pages, there are four other buttons on the screen – "Close Book," "Options," "Summarize," and "Take

Quiz." We didn't design icons for these, in part because existing icons work fairly well for them. "Close Book" should be marked with a green "EXIT" sign, which is universally recognized. The presence of an "Options" button should be optional – some students shouldn't have it available, to prevent distraction. However, its presence may be useful for other students who need to, for instance, move the agent. The "Summarize" and "Take Quiz" buttons are present to allow greater control over evaluation than was provided by CSLR's initial design, which brought up the quiz when the student turned the page. By making these buttons available to students whenever they want to use them, students can read without the distraction of a quiz (perhaps they already took it and are rereading a favorite part), or they can take a quiz without continuing in the book. This also separates the cognitively distinct actions of "turning the page" and answering questions about content. Possible icons for the "Summarize" and "Take Quiz" buttons include a pencil, and students raising hands. We placed these less-important buttons in a bar across the bottom of the screen so that they're fairly innocuous.

Inspired by the Corporate Comics website, we decided to give the agent a "speech bubble." This is a white area near the agent's head (which should have a angled affix indicating speech as in a comic). Anything that the agent says to the student will also appear in the speech bubble. This feature confers several benefits. First, it allows students to associate words with speech through casual use of the program without even reading the book. Second, it allows students who have trouble hearing to receive the instructions cross-modally. This advantage also allows users to interact with the agent without wearing headphones. Finally, the Corporate Comics website showed us that people will read something in a speech bubble that they wouldn't read in conventionally presented text (employing the Rosie the Riveter principle).

When the student struggles with a word, the agent should (after a sufficient amount of student effort) suggest a pronunciation or say something like "Try sounding it out." Although technical limitations prevented us from including this in the prototype, it is a feature of human reading tutors and was supported by our test users.

The original design also included the book's title and chapter tabs at the top of the screen. Although our final design is missing these items, we suggest that they are included in the working product. The original design underutilized the screen real estate where the book title and chapter tabs were located and because the title is important and the chapter tabs help when navigating through the book, they should be included, but not with a lot of surrounding space. In addition, the current user's login name would be useful to display, probably near the book title, as an aid for administrative purposes.

Future Issues

Experts are divided on whether presenting two pages side-by-side is preferable to a single "page" per screen. We think that layout issues like this should be decided by book creators. The design of the editing interface is a future project.

While we kept needs of advanced readers in mind while designing the interface, our primary target was the K-2 group. Some of the interface features, such as spoken instructions, might annoy students who have been reading for years. A beginning reader/advanced reader modality may be desired.

Our test users suggested we separate the agent from the buttons on the screen. Perhaps the read buttons should be on the bottom bar, though the buttons currently there are less important than the reading buttons. They might also be placed next to the speech bubble at the top of the screen. Perhaps these buttons should be positioned by the book's creator too, though they should have a good default position, since it's unlikely that book creators will be user-interface experts.

Perhaps the "Playback" button shouldn't be as prominent as the other two reading buttons. Until the student has read to the computer, there is nothing to play back, so perhaps the button should be hidden. Once it appears, perhaps it should be a green "play" arrow.

A function available in Soliloquy that we liked but did not implement in our design was the ability to read a definition of a word. One possible design is that the student would click on a word, the agent would pronounce it, and then the student could click on a "Define" button that would bring up a window with information about that word. This feature would be especially helpful for domains such as science books.

Due to external circumstances, we were unable to test the mockup with anyone besides the two CSLR experts whom we interviewed. While this was very beneficial, as the experts had a more broad knowledge base than an elementary teacher, limiting our test group to two falls short of the recommended 5-7 testers for think-alouds. Further testing with education experts, teachers, and finally students is of the highest importance for the next phase of development.

References

Center for Spoken Language Research (2003). "NSF-ITR Reading Tutor Project." http://cslr.colorado.edu/beginweb/reading/reading.html and the associated Reading Tutor software.

Druin, A. et al. (1998). "Children as Our Technology Design Partners."

Kahler, H. (2000). "Constructive Interaction and Collaborative Work." *Interactions*, May/June2000.

Kintsch, W. (1998). Comprehension: A paradigm for cognition. NY: Cambridge University Press.

Kolman, Joe. "Corporate Comics." http://www.corporatecomics.com/

Reeves, B., & Nass, C. (1996). The Media Equation , Cambridge University Press, 1996.

Scaife, M. and Y. Rogers (1998). "Kids as Informants: Telling Us What We Didn't Know or Confirming What We Knew Already?"

Soliloquy Learning (2003). "Soliloquy Reading Assistant." http://www.reading-assistant.com/

Van Dijk, T. & Kintsch, W. (1983). Strategies of Discourse Comprehension. San Diego, CA: Academic Press.

Wessling Co. / Kid Classroom. "Living Books." http://www.livingbooks.com/

Appendix A. Initial Interview Notes

- First explain the scope of our project
  -
- What do you see as the goals of the Interactive Books?
  - For students to learn to read. To test comprehension. To establish fluency. Improve comprehension and pronunciation. Through simple questions and complex summaries. Also to recognize specific problems and suggest help.
  - The pedagogical goals are "Read to me," "Tell me what I said," "Break down into syllables," "provide help on particular words," and "Learn to read all together."
  - Also to design the next generation of interactive agent interfaces and be able to apply these to any interactive agent system, such as an information kiosk or a web search agent.
- Describe the students who will be using the Interactive Books.
  - K-12 students with varying uses. Low grades: kids who can't read at all – agent will read to them and ask them questions. Possibly have a "choose-your-own-adventure" sort of thing. All interactions must be visual or spoken – no reading or typing.
  - Early Readers: Student reads to the agent.
  - Higher grades: Give a complicated text; ask for a summary, give advice on the summary. This focuses on and tests comprehension.
  - The system needs to be flexible with student reading level. Some students in third grade read at a high school level while others still read at a kindergarten level.

S.S.
A reading teacher will follow along with a student, who is reading to the teacher. The teacher will annotate their own copy of the text with a "running record" of how well the kid is reading. This record is used for diagnosis for specific patterns that indicate specific problems. Diagnose discourse errors.

Some students read at a level lower than their current grade because they are behind for some reason. Others read at a low level because they have a disability of some sort.

We asked him if he could read something to us, but he told us we wouldn't learn much from the experience. He then described how his daughter reads. She's about six, and she first spends a lot of time looking at the pictures, which then gives her a leg up on guessing what the words will be.
An educator working with a kid who's stuck on a particular word would tell the student to look at the word again, see what the context of the word is, possibly start the word for the student, suggest the student look at the picture and ask the student questions about it.

What do you hear from teachers that are using the reading tutors? They're saying "that's exactly what I do with a magnetic board," moving parts of the words around. The teachers haven't used the tutors directly. Many students aren't being challenged, but it's

good practice. Kids are "fighting" to use the computer tutor; they do so in 20-minute chunks. Within a given class, there are a variety of reading levels.

Students using – There is no target. Both typical kids and special ed and ESL. Students currently using tutors (the test group) are first-graders, around 6-years-old. An issue with ESL is recognition with picture. The learning process is for the kid to figure out what the word is in her native language, to understand the concepts, and then use those concepts to bridge the distance to understanding in English.
A key for hearing-impaired students is hearing what they said.

Challenges for typical kids. If there's a glitch/bug, they'll shut it down or go away. Silent readers, who don't read aloud – how can you tell where they are, or if they're reading at all? Speech and volume issues, accented speech, speech disorders, etc. Key – system knows where student is and can ask questions, even if reading silently. Also, silent readers may want to hear a word pronounced. ESL – students should have a foundation in their language before starting another. Can ground English in either pictures or their own language. It's more important for them to put words in context, with already-familiar concepts. Hard to teach idioms. Hearing-impaired – impoverished language base. E.g. in ASL, one sign can stand for several similar entries in the thesaurus. Sounding-out when learning to read is difficult, especially vowels. Even deaf folks who make it to university typically only read at a sixth grade level. Semantic knowledge deprived. Use comparison to golden vowel.

Going into the program…
Student interface isn't terrible. Wishes it could save which agent was used last. Wants the agent to be nearer to the text. Editing interface isn't friendly. Sarel jumps in.
Options are transient and cosmetic.
Loads the book. "Where's the book?" Just an intro screen.
What are the buttons? What does this do. (All of the buttons.) Kid can't read labels – "what's a paragraph?" Suggests that maybe on roll over, the agents says "Click here and I'll read a sentence to you."
Roll-overs were confusing/unnecessary.
Played around with moving the monkey. "Hey, I can look up his butt." Clicks around on a lot of buttons.
"A Great Idea" – huh? Is this the title? Can it be read aloud?
"What's this?" Pokes at the buttons to the left of the agent
Captions should be the same size.

- ☐ What do you hear from the teachers who are testing the tutors?
  - ☐ Tutors
  - ☐ Interactive Books
- ☐ How familiar are you with Soliloquy? [Run Soliloquy now]
  - ☐ Briefly point out what you like and dislike about the program.
- ☐ How familiar are you with the Interactive Books?
  - ☐ Explain it somewhat if they're not familiar

- ☐ Describe some of the techniques currently used to teach reading to kids in the target groups.
  - ☐ What are some of the problems with these techniques?
- ☐ (Choose random book from shelf) – Can you show us how you would read some sentences from this book with a student?

With Interactive Books running:

- ☐ Tell us about a student that would be using the Interactive Books.
- ☐ Do what that student would do if he/she were using the program for the first time, and think aloud as that student. [Start at "click here to start reading."]
- ☐ Point out what features of the program would make it easy or hard for that student to use it, from your point of view and the student's.

Appendix B. User Testing Notes

User test
April 18, 2003

Jon Roeth (JR)
Trevor Stone (TS)
Tim Lloyd (TL)
Sarel van Vuuren (SV)
Tester 1 (T1)
Tester 2 (T2)

JR: We'll get started and anytime you want to stop for any reason whatsoever is fine. First of all, from your perspective from what you've seen of this or the Interactive Books at the beginning and what Nattawat has been working on, what kind of tasks do you see teachers wanting students to do with these Interactive Books? What kinds of things do you see them wanting to do?

T1: Do you want the kids' opinions or the teachers' opinions?

JR: Well, both then.

T1: I think the teachers would like it to be sort of seamless so that they could set the child in front of it and give them a really quick "click here," get them started, and then they go – that that would be a really great thing. We've seen a lot of teachers don't have a lot of time. The other thing that we've been working on is we've had fifth graders building books themselves, so making those interfaces as simple and easy to use. Actually, the kids learned ten times faster than the teachers and have been incredibly more productive.

JR: The students are actually quicker at learning how to edit books?

T1: Yes, they edit their own books. They've recorded their voices, typed in text, they use digital cameras. The kids are far more technologically savvy. And they have more time and more interest in doing it themselves. So I see that it would be a really great thing for the teachers to be able to introduce this and have things for the children to do. Have books, have tutorials, have products for kids where they just want them to have a learning experience in that realm, but also be able to have them create their own. Things have been, but the kids seem to be able to do it now without much change in the interface.

JR: So you say them are fifth graders. Do they make these books for each other at their own grade level or younger?

T1: They're making books right now, they're documenting their science experiments. But I had one of our previous tools that we were using had a vocabulary wizard and one of the things that I did was that I taught a group of fifth graders to build vocabulary tutors for themselves and also to build them for, they had reading partners, they were mentors for a kindergarten buddy or first grade buddy that they read with. And they could build vocabulary tools for younger kids. So, it's just very empowering for the children especially older ones, but I could see a third grader doing some of this. I have a group that is hearing-impaired in another project that are building their own books and they are using, they're waiting patiently for the video part of it because they've been doing little minimal video clips of ASL so that they can communicate with each other.

TS: So what's something that a teacher might specifically ask a student to do if they wanted the student to use the Books?

T1: On here or just in general?

TS: Particularly with the Interactive Books, but if you have a particular insight with "in general" that would be good to hear.

T1: I think it has always been helpful for teachers to be able to ask a child to read aloud, that's where you get that automaticisity and that fluency that they are reading. And we all have those memories of sitting in school during reading group and reading aloud. And the good thing that the Interactive Books is that they can actually be doing it with a headphone and a headset and it can be recorded and the teachers can hear it and it eliminates that embarrassment or stumbling over words in front of a group of kids. I find that for a child who, the teacher really wants them to work on their reading skills and their fluency skills, that this would be excellent for them to use that would take away that social embarrassment part of it, that social interaction. So teachers asking them to read aloud and recording that and having that and just getting them to just sit and visually read once we get the visual tracking – that gives a lot of information as well. The question at the end is "did they understand what they read?"

JR: You said to have them record what they are saying. Is that to listen to themselves later or immediately following what they just read or for the teacher listen to what they are still having trouble with?

T1: I've done a lot of things with another program where the child is doing maybe one word or sentence at a time and then they get to hear themselves and we use that in ESL and for hearing-impaired for perfecting your intonation and pitch and getting those ranges. I think that it wouldn't be so much for listening to themselves, but to have that stored somewhere so that the teacher at a later time could sit and listen to see where the child is stumbling where the problems were, or if in fact the child is actually doing what they are supposed to be doing and not just chanting garbage, which kids'll do for the fun of it. What questions have you developed, I mean, what is your interface?

TS: The interface that we have is, you can actually look at it there on Tim's screen there. We focused on the interface that students are presented with when they are reading or when they are having the book read to them. Currently we weren't really able to get live speech interaction to work for this demo, so one of the things that we have is sort of a speech bubble that is up at the top and anything that gets put there is also spoken by the agent and so as you are playing with it, if you see something that is up there that is what gets spoken by the agent. And I have it set up so that, we just recorded some sentences and so forth so that we can simulate some student speech going on so you can see sort of how it looks when a student will actually speak.

T1: So what other questions do you have?

JR: We're going to actually have you play with it, but to do that we are going to use a technique that we call "think-aloud" where we ...

T1: Oh yea, I know.

JR: OK, you've used "think-alouds" before? You have? Is it something you're comfortable with? Or you don't want to do that?

T1: No, I make people do it.

JR: Well, we're going to make you do it.

T1: I do this big thing with cooperative learning which is a great tool that I love to use with my students, but when I had to do it myself, I didn't like it. But when you're standing up there as a teacher and you say "Alright, so work together..."

JR: Well, we'd also like as you do it as much as possible to put yourself in the perspective of either an ESL student or a first-grader, somebody whose going to have a little bit of a struggle with, a little bit of word struggle, you know they can speak and hear and everything OK, but reading, the word recognition and everything.

TS: And maybe doesn't have a lot of computer background.

T1: Oddly enough, very few kids don't have computer background. I think that will change as we go out of Boulder county, but even in the poorer parts of the county have pretty good computer knowledge. The kids are pretty computer savvy.

SV: Oh, yea, my two year old can play with a mouse.

T1: Oh yea, your son can play with a mouse already.

SV: It's no problem with the mouse, he can even play games with it. He hits the specific bubbles on the screen in his game and they pop. Very good with the mouse.

T1: It's amazing how well children grasp this. My grandchildren send me email, which is just astonishing to me at four and five that they understand that concept that they can sit at the computer and type in a message, hit a button, and send a message. I think it's pretty astonishing.

JR: So we have our student [T1] and your going to start reading The Backyard Zoo. So, you're sitting down at the computer. Take it away.

T1: So what is it you want me to do first? You want me to read the story?

JR: We want you to read to the tutor Ms. Gurney and as you read it, you should be able to follow your progress through the story and she should be able to follow your progress through the story and be able to tell you how well you're doing by highlighting the words as you go.

T1: Is there something I need to click to get started.

JR: We'll let you see if you can figure it out. This should be relatively stand-alone.

T1: This middle one is probably me talking to the computer, so it makes the most sense to me. So I'm going to click on this one.

JR: There's the speech that happens [that the agent is saying].

T1: There's speech that happens?

TS: Yea, here's our little speech bubble right here showing what the agent is saying.

SV: We'll have actual audio speech recognition.

T1: The Backyard Zoo. It was the first day of summer vacation. Sue and Billy were eating breakfast. What can we do today Billy asked. I don't know about today, but I have an idea about t-t-t-tomorrow said Sue. We can go to the zoo. But the zoo is far away said Billy. Who will take us? I don't mean the zoo in the city said Sue. We can make our own zoo in the yard. All our friends can bring their pets. What a great idea said Billy. We can have balloons, peanuts, and lemonade. It will be lots of fun. I'll call Grandpa, said Sue. He worked at the zoo. He can make our zoo look just like the real one. Billy ran to the door. I'll tell everyone to bring their pets here tomorrow. Is that enough reading?

TS: So what were you thinking as the highlighting and coloring was going on?

T1: There was a lag.

TS: There's going to be a disconnect between what I was doing and what you were doing.

T1: IT was that lag. I think that kids are used to seeing things happen fairly instantaneously. So for me while I was reading this, I would want this to be highlighting pretty much in tandem as I was speaking.

SV: Let me interject here for a moment. The actual books have the highlighting working pretty much, right? So that's how this would look, but this is a little disconnect here because he is having to do it manually.

T1: OK, well that would be the only thing. The bubble here will not be spoken?

TS: No, it will be spoken and it will also be there as a bubble.

T1: Because if you have a child who can't read, if she's not chatting it, then they aren't going to get it. But I understood the figures, I could figure them out, but I'm not sure that first-graders would actually get it. They might. They might get it 10 times faster than I would, but to just sit here, it's not real clear what to click and what to do first. So an introduction page . . .

TL: So maybe as soon as the program comes up, Ms. Gurney says “Hello, welcome to the program, these are the buttons...”

T1: Yea, it only needs to be a sentence or so of instructions. That would have made it clearer for me. And if the child is still struggling and a certain amount of time goes by and they are not figuring out the buttons, if the buttons could be explained. But the buttons do make sense.

TS: We also have it set up so that if you move the mouse over one of the buttons, she says, “if you click here, I'll read to you...”

T1: That was my only other thing. If the buttons are explained, then the child would ... and the other time you would have a problem with that is if you have an ESL or non-English speaking.

JR: So do you think that children are more inclined to stop what they are doing and if they can't figure it our, ask a teacher or move the mouse around...

T1: They'll move the mouse around.

JR: OK, so you think that they will play with it.

T1: Yes, they'll play with it. My experience has been sitting back and observing them and getting them something. Unless you have a child who has pretty serious problems and has become completely dependent on the teacher or outside help. Most kids would rather play around with it and they're not afraid to play around with it and press buttons. One of the problems with it is that if you have buttons that you don't want them to touch.

We found that at one point they could move Ms. Gurney, one child realized that and then there was this flood of, it looked like the first thing that all these kids would do, independently, hit upon because they were moving and hitting all these buttons, they were turning her upside-down and looking for her underwear. Which was embarrassing. We disabled that. But they'll move the characters around and they'll just keep hitting buttons. One of the things that I don't see on this screen is an exit button somewhere that is pretty obvious other than the big STOP when they are in the middle of it. But children will recognize an EXIT sign.

JR: Like a big green EXIT?

T1: Yes, a big green EXIT. I was working with children who were getting frustrated with the tasks that they were given to do and they wanted to exit, they didn't want to call the teacher, they didn't want to play anymore, and it was just haunting them that this was stuck up on the screen – this task that they were not completing very well. And they wanted to be able to exit. I don't know if you've seen the part that we've added to the tutor – Nattawat and Jariya have added to the tutor. They can exit out and then a dinosaur comes and asks them, "Do you really want to go away?" And the head actually bobs up and says yes and no and they can read and the head movement they'll recognize if they can't read. But an out, just giving them an out they can take if they're really frustrated.

JR: Just closes down the whole thing.

T1: Right, because if they get frustrated, they won't come come back and do it again. They'll play with it once, but if they get frustrated, when it's their turn to come up again, they're not going to want to be there. It was too stressful.

JR: OK. Well, we didn't label it "exit," but there is actually a method to close this, can you find it?

T1: On this screen? Well, "close book" but I'm not sure kids.

JR: So make it more graphical?

T1: Yea, I think that the more graphical the interface, especially for younger kids. I mean, I keep going back to these international symbols for things.

TS: We recognize that we wanted to get more graphical and more finished buttons for "close book" and the others.

T1: But the more graphical and more generic, I mean, in symbols, this becomes better. That's why you can go to another country and find the exit even if you can't read the signs. You might not know where you are going, but you can get out of there.

JR: That's a good idea.

T1: So on this button here, the read to the character, when do they have the capacity to choose whether he reads or she reads?

JR: That's what those three buttons are. The one on the bottom is the one for playback of what the student read.

T1: So will this play back to the child?

JR: Yes.

T1: So they can hear themselves read and work on their pronunciation. So is it storing those .wav files somewhere?

JR: For now, it is stooring them just temporarily right here.

TS: This prototype obviously doesn't implement that.

T1: The playback becomes more and more important, especially for children who maybe have a speech impairment or impediment or if they can break it down that the reason; you could see how they are decoding where the blocks are by what it is that they are missing. Are they not getting particular blends? Phonemes? Or just not reading it correctly, but are making consistent mistakes. So to have that saved somewhere so that the teacher can go back and hear it or play it back for a parent at a parent-teacher conference.

JR: Yea, they can just save these files.

T1: So when the kids are done reading, they can just click on this button to see what they sounded like.

JR: That's the idea. Then the top two buttons, in fact all three, once you click on them say “STOP” and you'll be able to hit STOP once you are done doing that activity.

T1: I've had a number of children when I was doing speech data collection who wanted to hear themselves and of course all children say “I don't sound like that, “ but there are other kids who know they have a speech problem and don't want to hear themselves. And are very adamant about it and may not even want anyone, it's part of them not wanting to read aloud in a reading group. They're the last kid, whose palms sweat when the teacher asks them to read aloud.

SV: How would you have kids read two or three sentences from a paragraph?

TS: How would you want to have the computer read several sentences to you if you were the student?

T1: I'd highlight it, but that doesn't seem to do anything. So, if I start here at this sentence and click the computer to read, it should just keep reading until I tell it to stop. How do I get her to read only one word? If I am in the middle of this and I'm reading it and I get to this word that begins with a "z" and I have no idea what that is, but I want to keep going and I want her to read that to me. I would want to be able to click on the word and have her read it.

TS: That is in fact what happens.

SV: See click on "vacation."

TS: And she says "vacation."

T1: So she'll say that – OK. So one word at a time. She can do one word at a time? That makes sense to me.

SV: So now there's no sentences and paragraphs in the book given that sentences and paragraphs are very ambiguously defined especially in dialogues.

T1: Especially in dialogues. One of the issues I had when I was downloading text and building books and import text, was finding what was and wasn't paragraphs, especially in early readers, a paragraph there's one sentence on a page, maybe five words. And there were others where the structure of the book was such that each sentence was a paragraph in and of itself so it was frustrating to figure out where the paragraph marking would be for inputing the text. So I agree with Sarel to give the child the option, not to define what is a sentence and what is a paragraph , but just to read until they want to stop.

JR: That's what this one does.

T1: I'm not sure that a single sentence is very valuable. And if they want to hear a particular sentence again, they can just click on the sentence and have it read again. I think kids will figure that out on their own.

TS: The other thing we thought of was that if a kid was reading along and struggles with a word, they get a couple of rejections on it, then Ms. Gurney will say "vacation" as they're reading along.

T1: I think that is an excellent idea because that's what a teacher would do, that's what a coach would do. Or say "Do you want me to help you?" Some kids might be, I think that depends on the age of the user. An older child might want to be asked or keep trying on their own. A younger child might, if they're used to being read to by a parent or a teacher, they are used to being helped in a particular way. So that may be linked to the level of their reading.

SV: Based on what you are saying, there might need to be two modalities of user: one where as you are reading and come to a word you struggle with, the tutor interrupts you

and says the word, but for a more advanced reader that may be too disruptive. So for this user, the tutor does not interrupt, but allows the reader to go on. Once you are done, it says "You struggled with these seven words, let's try to do them again."

T1: I like the second modality especially for older readers, because even I read a lot, but I still come across words that I certainly don't know but can get from the context or can glean it from what I am reading, but would have no idea how to pronounce it. I think that with a more advanced reader, that you will just slow them down, particularly when the wording becomes more advanced. The further along they get and they are reading words that are longer and more difficult, they are just not going to be able to sound them out, although they know what it means. Are there capabilities in here to link a particular word with a dictionary word?

SV: Yes, we haven't done it yet because we've had difficulty finding a suitable dictionary, but yes, sooner or later, we need to do this.

T1: You also want them to be able to figure out the words from the context, but later on there are definitely going to be words that they won't be able to figure out. And it may be a word that is crucial to understanding the whole concept. But there are words that you just don't get, that you see over and over again.

TS: This is the default layout we have. We have also talked about having the agent up on the top and seeing what issues that raises.

T1: What about putting her on the bottom?

TS: That is only position we didn't do.

JR: Tell us why the top and bottom.

SV: And comment on the center one as well actually.

T1: For some reason it just seems to me that I would, partly because I am used to reading a computer screen, information is on the bottom where the buttons are too. For me it just seems intuitive and it also I would want what I am reading to be the most prominent feature and this is distracting. I would want the buttons down here. I don't know where I would put Ms. Gurney because I don't like her.

SV: What about having her not be visible, but popping up when she needed to say something?

T1: In working with younger children, I have noticed that they engage her in the work they are doing, she is secondary to the work they are doing, but they still go to her for instruction and click on her and talk to her. If I were reading, I wouldn't want her on the screen while I was reading aloud.

JR: What about a young child?

T1: I would think they would find her distracting, but I would move her away from the text in some way so that you have a block of text. I would move these on the bottom, but I would have Ms. Gurney up on top depending on the age of the child or maybe the child themselves could choose. I would have her in the upper right so she is out of the way of the reading of the text.. The upper-left would be the most distracting because that is where your eyes start when you are reading a line of text. So for me that is the reason why I would have her on this side instead of that side because that is the periphery. But I would separate her from these [buttons].

SV: Do you prefer one or two pages of text on the screen?

T1: I prefer one especially for younger readers because they're used to seeing that much text.

SV: Or what about one page of text with a page of pictures on the opposing page?

T1: I would go with that especially for younger kids with a descriptive picture and then the text. And obviously children are used to having the two-page. But when I am reading, I think that that depends on the level of the book. I am thinking of a project I am doing now with hearing-impaired children and they might want the agent central so it really depends on the population.