Development of an American Sign Language Game for Deaf Children

Valerie Henderson, Seungyon Lee, Helene Brashear, Harley Hamilton[1], Thad Starner, Steven Hamilton

College of Computing
Georgia Institute of Technology
Atlanta, GA 30332
{vlh, sylee, brashear, thad}@cc.gatech.edu

[1]Atlanta Area School for the Deaf
Clarkston, GA 30021
HHamilto@doe.k12.ga.us

ABSTRACT

We present a design for an interactive American Sign Language game geared for language development for deaf children. In addition to work on game design, we show how Wizard of Oz techniques can be used to facilitate our work on ASL recognition. We report on two Wizard of Oz studies which demonstrate our technique and maximize our iterative design process. We also detail specific implications to the design raised from working with deaf children and possible solutions.

Author Keywords

Deaf, children, ASL, Wizard of Oz method, computer games, language acquisition, gesture recognition.

ACM Classification Keywords

H.5.2 [Information interfaces and presentation]: User Interface - User-centered Design; K.8.0 [Personal Computing]: Games}

INTRODUCTION

A 1999 study estimated that 17% of children ages 2-7 and 37% of children ages 8-13 play computer games on any particular day [16]. These games can provide sensory clues essential to the game's action which can render it uninteresting, confusing, or even completely inaccessible for children with sensory impairments. Deaf children, whose native language is ASL, are further impeded by the fact that many educational games rely on English grammar skills or spoken audio files. For deaf children, neither of these methods enable the transfer of knowledge required to play the games.

Language Acquisition

Ninety percent of deaf children are born to hearing parents who do not know sign language [19]. Often these children's only exposure to language is from signing at school. Early childhood is a critical period for language acquisition, and exposure to language is key for linguistic development. Research has shows that this critical period applies not only to spoken language, but also to ASL acquisition [11, 14].

By two years of age, hearing children learning a spoken language are combining words in their expressive communication [22]. By one and a half years, deaf signing children of deaf parents are also combining signs to communicate. A third group, deaf children of hearing parents, develop language in the same sequence as the first two groups, however, at a much slower pace. The slower linguistic development of this third group has been attributed to incomplete language models and lack of daily interaction using a language [6, 17]. Studies have linked delayed language acquisition with delayed short term memory development [5].

Although many children sign at school and use ASL on a daily basis, many teachers at Atlanta Area School for the Deaf (AASD) report practice/repetition and phrase generation are particularly problematic for their students. While children may have the vocabulary necessary for conversation, they have trouble stringing the vocabulary together into complete phrases, leaving interpretation to the listener. Many teachers and aides have the context necessary to fill in the blanks when signing with the children. For example, the child may sign only ball, but the teacher fills in the context and realizes the child would like to play with the ball at recess. She can then respond appropriately and say, "No, recess is in the afternoon; you have to wait until then." Prompting the child for a more complete phrase leads to confusion as the signed request overwhelms the child. Often the child abruptly ends the conversation because of his confusion.

Current Products

To enhance the language instruction they receive at school, hearing children have a multitude of educational software products which are available both at home and at school. Interactive ASL software is very limited and usually concentrates on students ability to receive and comprehend language rather than their ability to generate language independently. Two examples are Con-SIGN-tration [15] in which the child plays a memory game, matching cards bearing ASL signs to cards with English words and Aesop's Fables: Four Fables [2] in which the child watches several of Aesop's Fables interpreted into sign and is then asked a series of comprehension questions in English following the stories. However, to our knowledge, no games currently on the market allow the child to communicate with the computer via their native language of ASL. Games that do prompt children to mimic signs have no measure of evaluation to help the child improve the clarity and correctness of their signs. This lack of repetition with feedback prevents the child from benefiting fully from the software.

Goals

Given the difficulty of language acquisition and the lack of opportunities for repetition, we propose a system that provides children with a chance to practice their signing. It should help them use their vocabulary to generate phrases which convey complete thoughts and ideas, and it should also provide deaf children with a fun and engaging game unique to them and appropriate for their language.

Solution

In this paper, we detail our methodology for creating a game targeting deaf children who might benefit from increased repetition and practice of ASL. We target children ages 6-8 who are early in their language development process. This game is comprised of two distinct components which span different research areas:

1. **Game interface**: a Human-Computer Interaction work involving design principles and iterative design practices.
2. **ASL recognition engine**: uses gesture recognition techniques to observe ASL and determine the correctness of signed phrases.

By using a Wizard of Oz (WOz) technique, we maximize progress in both of these areas. In this paper we detail our previous game development work, our previous gesture recognition work, and two preliminary WOz studies. We show how our modified WOz study format maintains the ability to use iterative, inclusive design practices on our game interface while also enabling data collection for our gesture recognition system and thus greatly reducing the time and effort required to develop robust ASL recognition.

We also detail problems of both the interface and the recognition engine which must be addressed with future research.

GAME DEVELOPMENT

A sample interaction involves Iris the cat, AASD's mascot. The child sits in front of a desktop system equipped with a camera for computer vision recognition of the child's ASL. The child clicks to wake Iris and make her pay attention. The child signs any of eight predetermined phrases and clicks again to tell Iris to execute the action. While the ASL recognition is currently simulated by a wizard, in the future, our recognition engine will match the child's signing to a mathematical model of an ASL phrase with a certain level of confidence. If the confidence level is high enough (indicating clear and correct signing), Iris will execute the child's command and return to her original position to await the next command. If the confidence level returned by the ASL system is too low, Iris looks puzzled (a thought bubble with a question-mark in it), and the child must re-sign the phrase.

Höysniemi, Hämäläinen, and Turkki [8] investigated WOz prototyping for fast paced games based on computer vision recognition of children's movements in a contextual situation. They showed that WOz techniques could be used effectively to collect data necessary for distinguishing several different types of motions (swimming, running, jumping, and trying to escape from spiders). However, the coarse, whole body motions they collected were quite different than the fine, complex motor movements of ASL. Additionally, our task introduces the concept of evaluation because ASL is a structured language complete with grammar, vocabulary, and other linguistic features. This correctness of the signs figures prominently into the progress of the mathematical models, while Höysniemi, et.al 's work simply collected a corpos of movements.

The prototype includes a live video feed of the child (Fig. 1b), the Attention Button (Fig. 1c) which segments the data for the ASL recognition system, the action scene including Iris (Fig. 1d), and the Action Buttons (Fig. 1e) which the child can click to see the correct phrases demonstrated by the tutor in the tutor window (Fig. 1a) The prototype is built with Macromedia Flash and incorporates previously recorded video of a teacher signing ASL phrases.

Figure 1. Screenshot of ASL Game Interface. a) Tutor Video b) Live Camera Feed c) Attention Button d) Animated Character and Environment e) Action Buttons

Since English is often not the native language of people who are born deaf, the game environment is text-free. In lieu of English, the interface uses only ASL video, icons and animated visual indicators cuing the user to perform the proper action.

When a child presses a pictorial button (Fig. 1e), a video of a tutor signing the correct ASL phrase is loaded and automatically plays. A series of attention drawing events help guide the child through the interaction. After the video plays, the attention button (Fig. 1c) flashes red until the child clicks it, ostensibly to wake Iris up. After the child clicks the attention button, a square around the live video feed of the child (Fig. 1b) flashes red, drawing the child's attention to the mirror-like feedback of himself as he signs (Fig. 1b). After the child is finished signing, he again clicks the attention button (Fig. 1c) to tell Iris to execute his command.

When testing the interface among researchers, we discovered the wizard had difficulty discerning when the child clicked the attention button and would key Iris's response too early. We solved this problem by adding auditory cues to the wizard's interface to indicate the child's mouse actions.

HARDWARE DEVELOPMENT

Sign language recognition is a growing research area in the field of gesture recognition. Research on sign language recognition has been done around the world, using American Sign Language [20, 21], Taiwanese Sign Language [10], Korean Sign Language [9], and Chinese Sign Language [3, 4]. Artificial Neural Networks (ANNs) [9, 13] and Hidden Markov Models (HMMs) [4, 18, 20, 23] are the two most prevalent pattern recognition techniques in use, with the former being popular before 1995 and most current systems using HMMs. These statistical recognition techniques have been successfully utilized in the speech recognition community, and are embraced by the activity and gesture recognition research communities for their abilities to generate models from limited and potentially noisy sensor data.

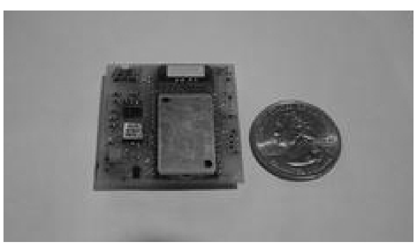

Figure 2 Wireless accelerometer boards

Our early sign language research demonstrated a continuous recognition system that performed at 98% accuracy with a 40 ASL sign vocabulary in a lab environment using HMMs [18]. This project has evolved to explore different sensor configurations and worked to increase flexibility and mobility of the system [1, 7, 12]. Our previous work demonstrates that the recognition accuracy of ASL is greatly increased by multiple modes of data. We have shown that accelerometer data complements computer vision data well and have engineered wireless, Bluetooth accelerometers. These accelerometers are about the size of a matchbox and worn on the wrist to provide [1]. See Figure 2.

The design of the current recognition system is progressing in parallel with the design of the WOz experiments and is still under development. Its development has been tightly integrated into the iterative development cycle, particularly in terms of system configuration, sensor design, and data input. The WOz studies have allowed us to iteratively test the acceptability and comfort of our sensors and system configuration and have provided valuable development feedback on sensor infrastructure needed for recognition.

Collecting data for use in statistical pattern recognition is laborious, time consuming, and tedious because a large number of samples must be collected and then labeled. We have collected ASL data in the past data by having a native signer click a mouse button, sign, click again, and repeat this process 30-40 times for each set of vocabulary. The data would then be categorized and labeled by hand both at the phrase level and the sign level. This data would then be ready for use with HMMs.

We use a "push-to-sign" mechanism to segment relevant data samples of signing from fidgeting and chatter. When using the game interface, the child clicks to alert Iris, signs, and then clicks to end the phrase. The signing is segmented by the start and stop clicks, labeled by the phrase selected for our data collection, and tagged as correct or incorrect. In this way, we will remove out-of-context and unscripted signing and ignore the child's meta comments such as "That's weird!" or "Why didn't it understand me?" This removing, segmenting and labeling greatly reduces the workload for the large-scale data collection needed for developing sign language recognition since most of the work is done concurrently with the data collection.

This push-to-sign mechanism is similar to those found in many speech recognition systems (called push-to-talk for speaking) and will allow our ASL recognition system to attempt recognition on only pertinent phrases of the children's sign. Our data consists of video of the user signing, and will, in the next iteration of testing, include accelerometer data from wrist mounted accelerometers. During the recognition phase, the video is processed, the hands are tracked, and the data is fed to the HMMs for classification.

ITERATIVE DESIGN

At the beginning of this project, we faced a dilemma: the primary developers are not deaf and do not have extensive experience with deaf children. Thus, an iterative design process with deaf children as our primary users figured prominently into project development. With the assistance of educational technology specialists and linguists at AASD, we developed a list of eight age appropriate phrases, listed in Table 1, columns 1 and 2.

For the game implemented in this phase, there are two ways Iris can respond to a correct ASL phrase: she can respond to commands by performing the action (chasing the butterfly in response to, "Go catch butterfly,") or she can respond to questions by nodding the affirmative and thinking of the answer (nodding and thinking of a hamburger when asked, "Are you hungry?"). Response to the phrases are listed in Table 1, columns 2 and 3

Glossed ASL	English Translation	Iris's Response
q(YOU LIKE MOUSE)	Do you like mice?	Nods and hearts fly from Iris to mouse
q(YOU HAPPY)	Do you feel happy?	Nods and throws confetti in air
q(YOU HUNGRY NOW)	Are you hungry now?	Nods and thinks of hamburger
YOU GO PLAY BALLOON	Go play with the balloon.	Plays with balloon
YOU MAKE FLOWERS GROW GO-ON	Go make the flowers grow.	Runs in garden and flowers appear
YOU GO CATCH BUTTERFLY	Go catch the butterfly.	Catches the butterfly
whq(WHO BEST FRIEND WHO)	Who is your best friend?	Thinks of duck
LOOK-THERE IRIS MOUSE OVER-THERE	Look, Iris! A mouse, over there!	Runs and catches mouse

Table 1. Glossed ASL Phrases, English Translations, and Iris's Response to Phrases

We developed a basic interface as described above. We then began to iterate on the prototype with small groups of children ages 9-11. While slightly older than our targeted age of 6-8, the older children were able to give us more detailed feedback about the prototype. Additionally, this left our primary subject pool (children ages 6-8) uncontaminated a future longitudinal study of ASL acquisition using the finalized game interface and integrated ASL recognition system.

Initial Trial

Three children participated in our initial trial at AASD. These children played the game as long as they wanted. Although we had a goal of twenty minutes of uninterrupted playing, during this trial we let the children self-direct and play as much (or as little) as they wished. The results for this trial are shown in Table 3, Subjects 1-3. The number of phrase attempts by each child is also included in this table. Because these are formative tests, we choose not to distinguish between correct and incorrect attempts. Indeed, many of the errors the children made could be attributed to mistakes the researchers made, both in interface design and understanding of ASL.

Subject	Gender	Time (min:sec)	Number Attempts	Self-Ended?
1	F	14:06	18	Yes
2	M	04:55	11	Yes
3	F	09:42	21	Yes
4	M	18:10	50	No
5	M	12:22	28	No

Table 2. Subjects' Test Data

Configuration

We had a facilitator, who was fluent in ASL and known to the child, stay with the child at all times. An ASL interpreter remained behind a partition with our wizard to assess the correctness of the child's signing. Due to our plans for a long term study, we needed to assess the viability of using different interpreters and having them adhere to the same standards of sign language grading.

During this test, we also needed to assess the children's tolerance and recall of our push-to-sign mechanism.

The WOz setup consisted of a single computer (Fig. 3d). The keyboard input was controlled by the wizard behind a partition. The mouse input was routed to the child (Fig. 3c). The interface output was split and routed to both the wizard and the child (Fig. 3b). The child was seated in a child's sized chair behind a small desk on which the mouse was placed. The facilitator sat behind the child, out of the line of vision, while the interpreter was behind the partition with the wizard. A camera recorded the child's signing and routed it directly to the interface for use as feedback when the child signed (Fig. 3a).

Results

We found that the children had no trouble with the push-to sign mechanism. At most, this mechanism was demonstrated twice by our facilitator and in some cases, the children understood it after only one demonstration. If the children forgot to push the attention button either before or (more often) after signing, they quickly realized what had happened and self-corrected.

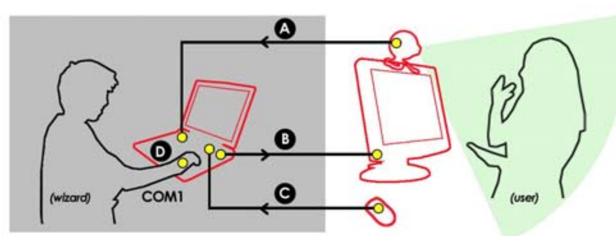

Figure 3. Initial system setup showing a) live camera feed b) interface output split between wizard and user c) child's mouse and d) the interface computer

However, the issue of interpreter variability proved to be problematic. Each interpreter graded the children differently. One interpreter was particularly strict due to her perception of the child's advanced linguistic status. Other interpreters were more lenient based on the desire to encourage the children's attempts. We found no clear criteria which could serve as metrics for allowing multiple interpreters and yet ensuring the accuracy of the phrases accepted as correct.

Because our facilitator was fluent in ASL and was more familiar with the layout of the game, for the third student, he became the interpreter. Because the student was deaf, the facilitator would call out yes or no to the wizard behind the partition. The wizard would then use the keyboard to direct Iris to the appropriate response. We discovered this method worked surprisingly well. The facilitator was often able to recognize critical mistakes in the child's signing before the child had completed the entire phrase. This gave our wizard time to make the cat's response to the child's signing appear seamless.

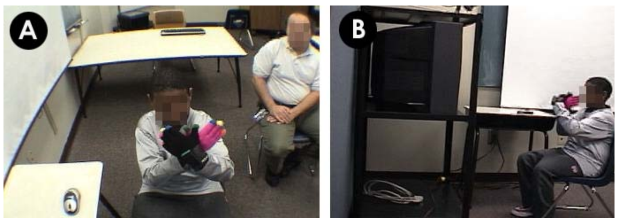

Figure 4. Wizard of Oz setup with participant and facilitator as seen from a) the ASL recognition camera and b) the sideview

We also discovered, in spite of the child's sized desk and chair, all the children had a tendency to trail their hands down to their waist (and thus behind the desk) toward the end of their signed phrases. While not presenting a problem during this phase of testing, this would prove problematic for our gesture recognition engine in the future. Before the second trial, we needed to rearrange the room furniture to accommodate the line-of-sight necessary for accurate and robust data collection.

After they indicated they were finished playing, the children were asked a few simple questions: 1) Did you like the game? 2) What was your favorite thing to make Iris do? and 3) What should the game have more of?

All the children responded that they liked the game, S1 most enthusiastically. The girls (S1 and S3) reported they liked the command, "Make the flowers grow." The other participant liked making Iris chase the mice. S2 and S3 asked for more actions in the games such as running or "run and do magic flowers!" S1 told us the game was cool, and she would use it to practice her signing. We noticed all the children preferred the action commands in contrast to the static questions.

Second Trial

Two children participated in this phase of the design prototype (S4 and S5 in Table 3).

Configuration

The physical setup of the room was as before. However, the child's sized desk was moved to the right of the children, giving our cameras a better line-of-sight to the child's signing area. Both of the children were right handed, although the desk position might have to be modified for left-handed children in future tests. See Figure 4.

Our facilitator was with the children at all times. However, he also acted as interpreter for both children's playing time, calling out responses to the wizard.

During the second iteration we also wanted to test the children's tolerance for devices which aid our gesture recognition system, specifically, lightweight, colored gloves and the small, wireless accelerometers described previously.

Tracking skin tones is particularly problematic for computer vision. For example, it is difficult to distinguish when the hands perform signs near the face or when hands cross and uncross as the skin tones are indistinguishable. During this test, we had the children wear small colored gloves. These provided a bright color easily identified by our computer vision algorithms. Both students wore a black glove on their left hand and a bright pink glove on their right. (See Figure 4a), but we needed to ascertain these would not hinder the child's signing.

Results

During the second iteration of the interface design, we replaced some of the tutor video with tutor video which included instructions, hoping to emulate a teacher's instruction to the student and reduce the child's interaction with the facilitator in the room. For example, during the first iteration, the tutor merely presented the phrase, signing, "Go catch the butterfly." In the second phase, the tutor signed, "You tell Iris, 'go catch the butterfly.'" Unexpectedly, this change introduced complications due to the spatial aspects of ASL.

ASL is a spatial language with rich directional expression. In ASL each signer has a signing space which is maintained in front of the signer. By setting up subsections of the signing space and indexing the subsections (by referencing them in conversation), signers indicate interactions between concepts, people, times, etc.

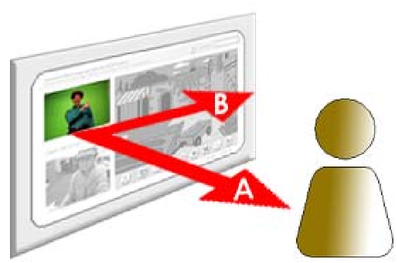

Figure 5. Interface and Tutor's Signing Space

For example, verbs can be directional and imply subject and object from the way they move through space. For example in the phrase, "Bob, tell Sue the answer," the sign for tell moves from Bob to Sue. If instead Sue was to tell Bob the answer the sign would move from Sue to Bob. During the second test, we used phrases such as, "You tell Iris, 'go catch the butterfly.'" We thought the children would understand that the imperative 'you tell Iris' was directed at them and the second half ('go catch the butterfly') would be the phrase needed to activate Iris. However, we discovered that both children included the directive 'you tell Iris' in their attempts. After reviewing the video clips that the children were shown, both we and the educators at AASD believe the interface layout and expression of the signing video were not spatially compatible.

We believe the children were treating the interface as an ASL construct rather than a 2D interface. Thus, when the tutor signs ‘you tell Iris,’ she should sign directly to the child (Figure 5a). When signing the command ‘go catch the butterfly,’ she should turn 90 degrees and sign toward the animated cat (Figure 5b). Thus, she would maintain proper ASL spatial grammar. Since the video was shot without this concern, we introduced confusion.

Unexpectedly, the change in the position of the desk and the colored gloves provided a significant marker for our gesture recognition. With the mouse placed on the desk and the child required to click a button before and after signing, the gesture of the hand on the desk begins and ends each signed phrase. This movement is very distinct and easier for our ASL recognition system to identify. Additionally, the light color of the desk provides a high contrast environment.

One child (S4) expressed no reservations about the gloves saying, “They’re fine; no problem!” when asked about wearing the gloves after playing the game. The other child said the gloves bothered him a little bit, but he would wear them if they were necessary to play the game. Neither child commented on the accelerometers on their wrists. During this trial, both children played until we asked them to stop (Table 3, Subjects 4-5).

Again, the children were asked several subjective questions about their experience playing the game. S4 was very enthusiastic and animated when asked if he liked the game. He also liked all of the actions he could make Iris do, but he wanted a race. S5 liked making Iris chase the mouse, but couldn’t think of anything else he’d like the game to add. Again, both children preferred the action commands to the static questions and thought bubble responses.

S4 also provided an interesting behavior study. He was very animated while signing. He signed, “That’s weird!” when Iris did not recognize his signing several times in a row. He would repeat the same phrase over and over until he got it right before moving on to another phrase. He began adding advanced sign language. For example, when he asked Iris, “Are you hungry?” Iris nodded and thought of a hamburger. The next time he attempted that phrase, he asked Iris, “Are you hungry for a hamburger?”

Iteration Three Setup

After the first two WOz tests, the test setting was revised to incorporate another machine dedicated to data processing for the ASL recognition (Fig. 6 COM2). The data machine receives input data through the live video feed (Fig. 6e) and accelerometers on the wrist (Fig. 6f).

As before, the live video is routed directly into the interface (Fig. 6a), the child controls the interface via the mouse (Fig. 6c), and the wizard controls Iris’s response via the keyboard (Fig. 6 COM1). The interface output is again routed to the child’s monitor (Fig. 6b). The correct or incorrect designation determined by the wizard and facilitator is also routed to the data machine for automatic labeling (Fig. 6g).

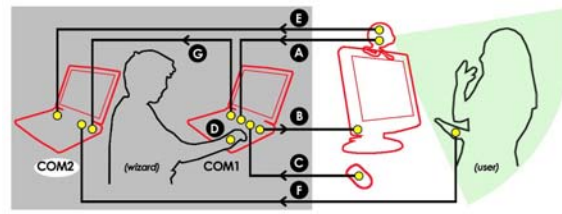

Figure 6. Proposed integration of ASL processing machine and interface machine

Based on our goal of 20 minutes continuous playing time and the children’s desire for a more competitive, action oriented game, we have revamped the prototype into a game with more levels and a goal oriented narrative flow. In the new prototype, “Kitten Escape!”, Iris’s multicolored kittens have been chased out of their basket by various animals such as spiders, snakes or alligators and have hidden somewhere around the backyard (e.g. behind the wagon, on the wall, etc.). Previously, the game involved directing Iris to do unrelated tasks. In the new version, the child tells Iris where to find each of her kittens and Iris returns them to the basket as the child signs correctly. The child will also tell Iris to chase away the predator before moving on to the next level. On each level, the kittens hide in different locations. Table 4 details the colors and locations that the child must sign.

Kitten Color	Locations
Blue	Behind wagon
Green	Under chair
White	In flowers
Orange	In bedroom
Black	On wall

Table 4. Vocabulary for "Kitten Escape!"

Progress Update

We have piloted this game and collected the corresponding sensor data. Based on preliminary analysis, it appears this new game design provides motivation for the children to play longer than in the previous designs. The modified WOz method allowed us to efficiently gather data, track collection progress, and rapidly assess data quality. This data has allowed us to move beyond infrastructure building and has given us “real world” data to incorporate into our current recognition system development cycle.

FUTURE WORK

Our current sign language recognition development is focused on generalizing models from user-dependence to user independence. We will continue our research in computer vision and machine learning algorithms to improve the recognition system. We have now collected a large data set with computer vision and accelerometer data. From this data we hope to build more descriptive and flexible models and begin in-house testing. Once satisfactory accuracy and performance have been achieved, we will begin user testing on an integrated system that uses the recognition technology. Important future questions for this integration include:

1) What are acceptable error rates for classification?

2) What are acceptable tolerances for signing errors?

2) How well does our WOz simulation reflect the integrated system performance without a human in the loop?

An area of future work for the game design is a help system translated into sign and accessible by the children. Currently, only an intro clip is played explaining how to use the push-to-sign mechanism and the basic plot of the game. This occurs in the small tutor window (Figure 1a). We have discovered that even with this, the children are still unsure of themselves and immediately ask the facilitator what to do. After he demonstrates, the children feel confident to attempt playing on their own. We would like to reduce this barrier and enable the children to play without an adult to guide them. We believe we need to introduce a full-screen introduction complete with a demonstration of a child playing the game for several phrases. However, the first step will be to ask children who have played and mastered the game, "How would you tell someone how to play this game?" We believe this will help us to identify the particularly confusing stages of the game and include more explanation of them.

CONCLUSIONS

Our method of using a Wizard of Oz study to facilitate research on both the visual interface and the ASL recognition engine has proved very successful. By modifying small aspects of our WOz setup such as the angle of the camera focused on the child or the location of the mouse and desk, we have developed a system capable of gathering and categorizing large amounts of sensor data quickly. This enables a more rapid project timeline. As gathering large amounts of data for ASL recognition algorithms is laborious and difficult task, our setup allowed us to leverage our valuable time with the children. By obtaining feedback about the interface and game design, we were able to improve the experience for the children while building a structure capable of furthering our recognition engine rapidly.

While applied to the very specific context of American Sign Language recognition, we feel that our method could provide benefits in more generalized research. This method allows collection of large amounts of data in a manner which frames the data in terms of an activity instead of rote, meaningless repetition. This method could be applied to speech recognition work, gesture recognition, and activity recognition.

We and the educators at AASD were very encouraged by the way in which the children refined and enunciated their signs after Iris did not understand them. All the children were willing to attempt phrases multiple times, making their signing more distinct and clear each time. We are also encouraged by the children's obvious enthusiasm for the game and their willingness to try new things.

We have presented a design for an interactive American Sign Language game geared for language development. In addition to work on game design, we have shown how WOz techniques can be used to facilitate our work on ASL recognition. We detail how we have adapted the technique to facilitate the collection and labeling of large amounts of data and to test a variety of sensors and algorithms.

ACKNOWLEDGMENTS

We would particularly like to thank the children, staff, and administration at Atlanta Area School for the Deaf for their enthusiasm and support of this project. This work is funded in part by NSF Career Grant #0093291 and the NIDRR Wireless RERC.

REFERENCES

1. H. Brashear, T. Starner, P. Lukowicz and H. Junker. Using Multiple Sensors for Mobile Sign Language Recognition. In *Proceedings of the Seventh IEEE International Symposium on Wearable Computers*, pp. 45-52, 2003.

2. Texas School for the Deaf and G. Pollard, Aesop: Four Fables, Texas School for the Deaf, Austin, TX, http://www.readingonline.org/electronic/project/post2.html, 1998.

3. G. Fang, W. Gao and D. Zhao. Large Vocabulary Sign Language Recognition Based on Hierarchical Decision Trees. In *Intl. Conf. on Multimodal Interfaces*, pp. 125-131, 2003.

4. W. Gao, G. Fang, D. Zhao and Y. Chen. Transition Movement Models for Large Vocabulary Continuous Sign Language Recognition (CSL). In *Proceedings. Sixth IEEE Intl. Conference on Automatic Face and Gesture Recognition*, pp. 553-558, 2004.

5. H. Hamilton and T. Holzman. Linguistic Encoding in Short-Term Memory as a Function of Stimulus Type. *Memory and Cognition*. 17 (5). pp. 541-50, 1989.

6. H. Hamilton and D. Lillo-Martin. Imitative Production of Verbs of Movement and Location: A Comparative Study. *Sign Language Studies*. 50. pp. 29-57, 1986.

7. J.L. Hernandez-Rebollar, N. Kyriakopoulos and R.W. Lindeman. A New Instrumented Approach for Translating American Sign Language into Sound and Text. In *Proceedings. Sixth IEEE Intl. Conference on Automatic Face and Gesture Recognition*, pp. 547-552, 2004.

8. J. Höysniemi, P. Hämäläinen and L. Turkki. Wizard of Oz Prototyping of Computer Vision Based Action

Games for Children. In *Proceedings of Interaction Design and Children*, pp. 27-34, 2004.

9. J.S. Kim, W. Jang and Z. Bien. A Dynamic Gesture Recognition System for the Korean Sign Language KSL. *IEEE Transactions on Systems, Man, and Cybernetics*. 26 (2). pp. 354-359, 1996.

10. R. Liang and M. Ouhyoung. A Real-time Continuous Gesture Recognition System for Sign Language. In *Third International Conference on Automatic Face and Gesture Recognition*, pp. 558-565, 1998.

11. R.I. Mayberry and E.B. Eichen. The Long-Lasting Advantage of Learning Sign Language in Childhood: Another Look at the Critical Period for Language Acquisition. *Journal of Memory and Language*. 30. pp. 486-498, 1991.

12. R.M. McGuire, J. Hernandez-Rebollar, T. Starner, V. Henderson, H. Brashear and D.S. Ross. Towards a One-Way American Sign Language Translator. In *Proceedings. Sixth IEEE Int.l Conference on Automatic Face and Gesture Recognition*, pp. 620-625, 2004.

13. K. Murakami and H. Taguchi. Gesture Recognition Using Recurrent Neural Networks. In *CHI '91 Conference Proceedings*, pp. 237-241, 1991.

14. E. L. Newport. Maturational Constraints on Language Learning. *Cognitive Science*. 14. pp. 11-28, 1990.

15. Institute for Disabilities Research and Training Inc. (IDRT), Con-SIGN-tration, http://www.idrt.com/ProductInfo.php?ID=32&u=1, 1999.

16. D. Roberts, U. Foehr, V. Rideout and M. Brodie, Kids and Media @ the New Millennium, Menlo Park, CA, 1999.

17. P. Spencer and A. Lederberg, Different Modes, Different Models: Communication and Language of Young Deaf Children and Their Mothers, in M. Romski, ed., *Communication and Language: Discoveries from Atypical Development*, Harvard University Press, 1997, pp. 203-30.

18. T. Starner and A. Pentland. Visual Recognition of American Sign Language Using Hidden Markov Models. In *Proc. of the Intl. Workshop on Automatic Face- and Gesture-Recognition*, pp. 1995.

19. Gallaudet University, Regional and National Summary Report of Data from the 1999-2000 Annual Survey of Deaf and Hard of Hearing Children and Youth., Washington, D. C., 2001.

20. C. Vogler and D. Metaxas. Adapting Hidden Markov Models for ASL recognition by using three-dimensional computer vision methods. In *Proceedings of the IEEE Intl. Conference on Systems, Man and Cybernetics*, pp. 156-161, 1997.

21. C. Vogler and D. Metaxas. ASL Recognition Based on a Coupling Between HMMs and 3D Motion Analysis. In *Proceedings of the IEEE Intl. Conference on Computer Vision*, pp. 363-369, 1998.

22. G. Wells, *Language Development in Preschool Years*, Harvard University Press, Cambridge, MA, 1985.

23. J. Yamato, J. Ohya and K. Ishii. Recognizing Human Action in Time-Sequential Images using Hidden Markov Models. In *Proceedings of the IEEE Intl. Conference on Computer Vision and Pattern Recognition*, pp. 379-385, 1992.