An Intelligent Tutoring System for Deaf Learners of Written English

Lisa N. Michaud
Kathleen F. McCoy
Dept. of Computer and Information Sciences
University of Delaware, Newark, DE 19716
+1 302 831 1956
michaud@cis.udel.edu, mccoy@cis.udel.edu

Christopher A. Pennington
AgoraNet, Inc.
10 Polly Drummond Hill Rd., Newark, DE 19711
+1 302 224 2475
penningt@agora-net.com

ABSTRACT

This paper describes progress toward a prototype implementation of a tool which aims to improve literacy in deaf high school and college students who are native (or near native) signers of American Sign Language (ASL). We envision a system that will take a piece of text written by a deaf student, analyze that text for grammatical errors, and engage that student in a tutorial dialogue, enabling the student to generate appropriate corrections to the text. A strong focus of this work is to develop a system which adapts this process to the knowledge level and learning strengths of the user and which has the flexibility to engage in multi-modal, multi-lingual tutorial instruction utilizing both English and the native language of the user.

Keywords

Intelligent tutoring systems, user modeling, English literacy, second language acquisition, American Sign Language

INTRODUCTION

This project seeks to implement an innovative strategy for supporting English literacy skills, specifically writing, for deaf students — especially those who are more comfortable with American Sign Language (ASL) than with English. The problem of deaf literacy in English has been well-documented and can affect every aspect of deaf students' education (including the learning of Science and Math). Although data on writing skills is difficult to obtain, we note that the reading comprehension level of deaf students is considerably lower than that of their hearing counterparts, "... with about half of the population of deaf 18-year-olds reading at or below a fourth grade level and only about 10% reading above the eighth grade level..." [57] In our own observations, the writing of this learner population also exhibits marked differences from that of their hearing peers, containing myriad errors which are different from those committed by hearing writers. These errors include dropped "be" (*She really pretty*), and missing possessives (*She age is 13*), subject/verb agreement, plural markers, and determiners (*She really like go with friend to mall*).[1]

ASSETS'00, November 13-15, 2000, Arlington, Virginia.

One explanation for this disparity is the distinction between ASL and English. ASL is a full, natural, and independent language distinct from English [56], [34], [39], [40], [5], [63], [13], [12]. Although there is a common misperception that ASL and English are identical except for the medium through which they are delivered, the word and sentence structures of ASL are actually so different from those of English that Jacobs [30] has argued it should be considered a Truly Foreign Language, as different from English as Chinese or Swahili.

Because of this difference between the two languages, a major factor in our system's design has been the application of second language acquisition theory to aspects of our system concept and architecture, as well as an investigation into how best to approach computer-generated instruction (which in traditional systems is entirely text-based) in the context of this learner population's special needs. While ASL is the first and strongest language of many deaf children, and a strong background in any native language is generally helpful in the acquisition of a second one, vast differences between two languages can pose serious challenges for a learner. For this and other reasons which will be discussed below, the deaf population of English learners places some rather unique demands on a language-instruction system.

Signing and Education

Because they receive little if any useful auditory input, prelingually deaf children are unable to acquire English in the way that hearing children do. Speechreading, also known as lipreading, provides highly ambiguous input; e.g., "pat" and "man" look the same on the lips. Moreover, in casual speech, only about 40% of English phonemes are visible [31], [64].

Such input is insufficient for most children to learn English naturally. Therefore, these children have to learn English in the classroom at the same time that they learn to read. This is far more difficult than transferring literacy from a native language to one being acquired [38]. The results are often less than ideal: the average deaf high school graduate has only a fourth grade reading ability [69], [18], [11], [57].

Traditionally, there have been three primary methods of communication in classrooms with deaf children: spoken English, signed English systems, and, to a much lesser

1. These examples were compiled from one of our actual writing samples from the learner population.

extent, ASL. The oralist philosophy, which propounds the use of spoken English without the use of any signing, has achieved very mixed success. It requires the students to rely on speechreading for their main means of receiving language communication. As mentioned earlier, this is an exceedingly difficult task. Some individuals have a talent for it and excel academically, but the rest are cut off from usable linguistic input.

Manually Coded English (MCE) or signed English systems are artificial codes developed to make the word and sentence structure of English visible. Each unit of meaning (morpheme) in English is given its own sign, and these signs are supposed to be articulated in English morpheme order. Many of the signs used are borrowed directly from ASL, but some were invented specifically for the MCEs.

Although such systems might sound like an excellent pedagogical tool in theory, in practice they have not lived up to their potential. It takes two to three times as long on average to articulate a sign as it does to say a word [6], [7], [26], and yet ideas can be expressed in approximately the same amount of time in natural spoken and signed languages. This is because signed languages employ a simultaneous presentation of morphemes, whereas spoken languages primarily present morphemes sequentially. MCEs, however, require the sequential presentation of morphemes in a signed mode; therefore, people using MCEs would have to take an uncomfortably long period of time to convey propositions if they were to sign every English morpheme [24].

Lieberman [41] states that if propositions are conveyed too slowly, messages become hard to process. In practice, users of simultaneous communication (the act of speaking and signing at the same time) typically leave out many morphemes in their signing, and sometimes incorporate aspects of ASL syntax and morphology instead [14], [2], [9], [43], [35], [44], [58]. It appears that a more "natural" communication paradigm is needed.

A relatively recent concept in deaf education is the "Bi-Bi" philosophy; deaf children are seen as being bi-lingual and bi-cultural. The primary mode of communication in a Bi-Bi classroom is ASL, and English is taught as a second language [23], [55], [4], [49], [57], [32], [64], [19], [27].

The use of both languages is motivated by the argument that even in classrooms where hearing students are learning a second oral language, pure immersion in the second language might not be as effective as strategic use of each of the students' two languages. For example, DiPietro [17] wrote, "The belief remains strong that the target language should be used as much as possible right from the first day. Such a practice supposedly forces students to 'think' in the target language. However, the breakdown in communication caused by this narrow view of language use may outweigh any benefit." Furthermore, children who have a strong foundation in ASL outperform their less fluent peers with regard to English literacy [59]. Also, studies involving hearing children learning English as a second language show that children in bilingual programs have better comprehension of spoken English than do those in immersion ones [21], [66], [37], [15].

Preliminary reports indicate that one such English as a Second Language (ESL) program for deaf students has shown some success [47]. Similar programs have been established in other parts of the world [16], [52].

The system we are developing endeavors to be consistent with the Bi-Bi philosophy. Specifically, we are taking into account the similarities and differences between English and ASL and using this information throughout our intelligent tutoring system. Later, we will also discuss some future goals which include providing at least some of the instructional feedback in the student's native language (i.e., ASL) in addition to the tailored English text explanations.

SYSTEM OVERVIEW

The envisioned system is called ICICLE (Interactive Computer Identification and Correction of Language Errors) [45], [51], [46]. It can be viewed as two interacting modules: the error identification module, which is responsible for analyzing the text and identifying errors to be corrected, and the response generation module, which is responsible for generating and presenting the tutorial dialogue. Both components will make use of knowledge bases which cover both general information about how to analyze and explain grammatical structures and specific information about the current system user, his level of knowledge, and how he has reacted to system efforts in the past.

The system-user interaction in ICICLE centers around a cycle of user input and system response. The entry into the cycle occurs when the user inputs a piece of writing (either directly typed into the system, or through loading a prepared text file) for the system to analyze. The system identifies the grammatical errors in the text and then presents those errors which are relevant for tutoring to the user. (Relevance for tutoring is defined in terms of the user's preparedness for learning a concept; this will be discussed more in-depth later.)

Once relevant errors have been identified, they are passed to a response generation module which brings the student through a review of those aspects of his or her writing which need to be improved, after which the student is free (and encouraged) to make changes and request a new analysis.

This type of tutorial instruction, where the system reviews the performance of a user after a task has been completed, is motivated by the theory that the cognitive demands of some tasks are so intense that learning is hampered during their execution [48], [62]. It is our belief that the composition of original text in a non-native language is a task of this level of cognitive difficulty. Sources such as [10] emphasize the usefulness of post-performance review or reflection as an approach to learning, and stress the suitability of computer systems in this kind of task, since they can perfectly capture the user performance and then review any aspect of it. Our system therefore endeavors to accomplish this goal and to optimize the knowledge derived from the composition experience by reviewing it with the user, an approach also found in some other language tutors such as MILT [33].

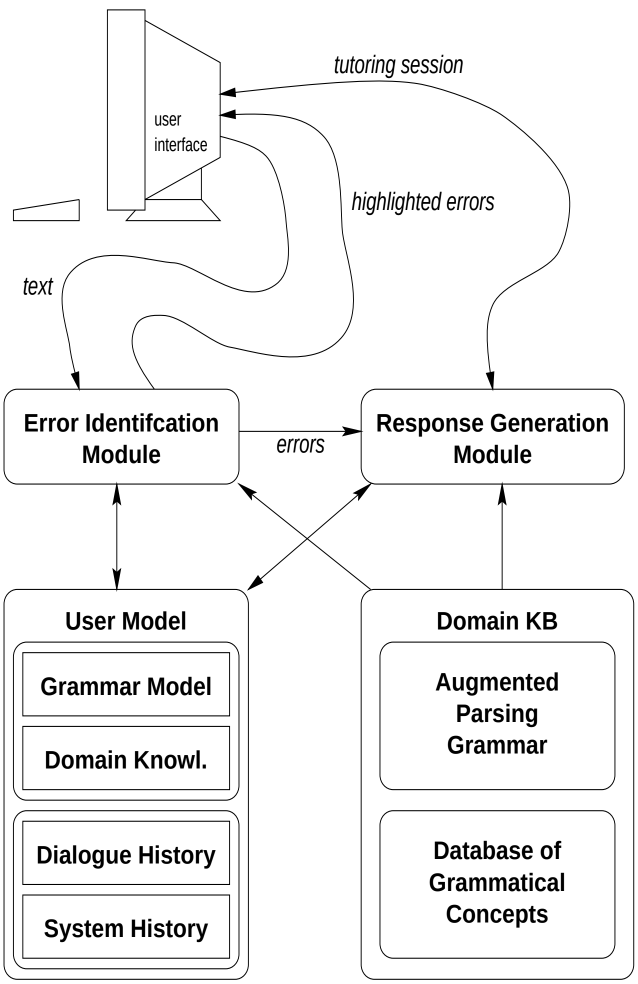

Figure 1: ICICLE System Diagram

While providing this tutorial feedback, another goal of the ICICLE system is to satisfy the deaf learner's need for understandable English input. Deaf learners receive nearly all of their English input through written material, often academic texts aimed at the comprehension level of their hearing peers [1], and yet the consensus among most researchers in Second Language Acquisition [36], [65], [67], [28] holds that second language input at or near the learner's level of existing proficiency is most beneficial for learning. We therefore intend for our system's generated text to use grammatical constructions that involve those aspects of English the student is currently attempting to master in his or her writing, while also using constructs from his or her area of well-established knowledge. Because most researchers hold that the reading proficiency level of second language learners is in advance of their current production proficiency, we expect the student to benefit from this much-needed level of language exposure while having little difficulty with reading comprehension.

Another way in which we anticipate that our system will address the unique needs of the deaf population is by providing the user with feedback on his or her writing without involving a human teacher. Some students might prefer this mode of feedback since they would not risk feeling a "loss of face" as they might with a human tutor. The hope is that this will get the students to write more.

CURRENT IMPLEMENTATION

Our work to date has primarily focused on the error recognition component. In the current implementation, the user has the ability to load in a pretyped text or to type the sentences directly into the system. When the user requests an analysis, the system determines what grammatical errors are in the text and highlights those sentences containing errors. When the user clicks on a sentence that has been highlighted, a "canned" explanation of the error or errors involved appears. The user can then make corrections to the sentence and resubmit it for further analysis. Thus the rudiments of our "user input, system response" cycle are in place, although the form of instruction is very crude as all errors are explained in one-sentence statements that have been prepared by the system designers ahead of time.

Grammar Coverage

The implemented analysis component of ICICLE uses an English grammar augmented with error-production rules, or mal-rules [53], [68]. These mal-rules allow sentences containing errors to be parsed with the grammar, and enable the system to flag errors when they occur. The mal-rules contain annotations that indicate what error they "realize." They were originally derived with the assistance of an error taxonomy which was the result of an analysis of writing samples from our learner population. The error taxonomy was also influenced by an analysis of how ASL knowledge might influence written English [60], [61] and other information about ASL.

The ICICLE grammar itself is a broad-coverage grammar designed to parse a wide variety of both grammatical sentences and sentences containing errors. The first implementation was built around the COMLEX Syntax 2.2 lexicon [25], which contains approximately 38,000 different syntactic head words. There is a simple set of rules that allows for inflection, thereby doubling the number of noun forms, while providing three to four times as many verb forms as there are heads. Thus ICICLE can handle approximately 40,000 noun forms, 8,000 adjectives, and well over 15,000 verb forms. In addition, unknown words coming into the system are assumed to be proper nouns, thus expanding the number of words handled even further.

This grammar contains approximately 25 different adjectival subcategorizations, including subcategorizations requiring an extraposed structure (the "it" in "it is true that he is here"). It also includes half a dozen noun complementation types. It has approximately 110 different verb complementation frames. The grammar is also able to account for verb-particle constructions. Additionally, the grammar allows for various different types of subjects, including infinitivals with and without subjects ("to fail a class is unfortunate", "for him to fail the class is irresponsible"). It handles yes/no questions, wh-questions, and both subject and object relative clauses.

The grammar has only limited abilities concerning coordination — it only allows limited constituent coordination, and does not allow non-constituent coordination (e.g. "I

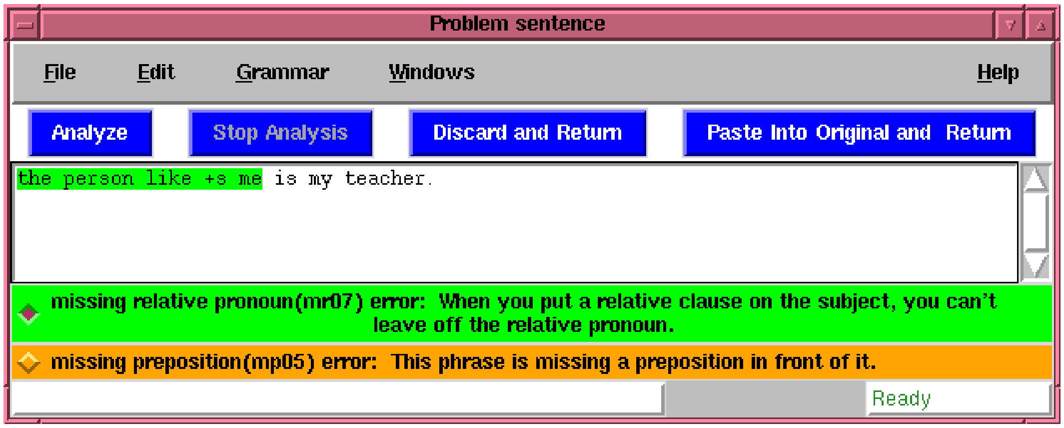

Figure 2: User Interface Example

saw and he hit the ball") at all. It is also fairly weak in its handling of adjunct subordinate clauses. The population we are concerned with has significant trouble with this; in particular there is a strong propensity towards overusing the word "because". Adverbs are also problematic, in that the system is not yet able to differentiate what position a given adverb should be able to take in a sentence, so no errors in adverb placement can be flagged at this time. We are presently in the process of integrating a new version of the lexicon that includes features specifying what each adverb can attach to. Once this is done, we expect to be able to process adverbs quite effectively.

Another ongoing modification to the existing ICICLE implementation is a change over to a new grammar developed by Rose [50] and specifically modified for use with the ICICLE system. Since this grammar introduces the ability to relax certain grammar requirements rather than explicitly modeling ungrammatical structures through mal-rules, it may eliminate many of our mal-rules and make repetition of different permutations of the same basic error unnecessary.

User Interface

The user interface is designed with a windowing environment implemented in Tcl/Tk. Its primary window has a space for text entry and controls for initiating analysis, loading prepared files, and other such functions. When parsing is complete, sentences in the main window are highlighted with different colors corresponding to the different types of errors that have been found. From a list summarizing the different errors that were found, the user can select one, and the system will highlight all occurrences of that error in the text. When the user double-clicks on a sentence, a separate "fix-it" window is displayed with the sentence in question, along with descriptions of the error or errors in the text. The user can click on an error and the system will highlight the part of the sentence where the error occurred, focusing the user's attention on where the correction must take place.

For example, in the sentence "The person likes me is my teacher," when the user selects the error describing the missing relative pronoun, only the ungrammatical relative clause will be highlighted. (See the figure above.) The "fix-it" window also allows the user to change the sentence and then re-parse it. If the changes are acceptable to the system and desired by the user, the new sentence can be substituted back into the main text. In the eventual system, the tutorial interaction, which will evolve into a far more involved explanation generated specifically for an individual learner, will take place through a similar interface.

TOWARD MODELING THE USER

One problem that we face with the target population is that some of the texts have many errors and there is a great deal of difference among the texts produced by different learners (i.e., some are quite good while others contain many fundamental errors). This leads to the realization that in order for the system to make intelligent analyses of the writing of different students and to deliver instruction to such a broad range of students, it is going to have to model its users and use that model at multiple levels of its interaction with that individual.

There are several attributes of the user which ICICLE should model in order to improve its suitability for use by a range of students. For instance, modeling the user's history of system use and how different explanations have succeeded or failed to influence subsequent performance would be greatly useful if the system wishes to concentrate on tutorial techniques which work well with that particular individual.

The focus of current work has been on another aspect of the student which is very relevant to tutorial instruction, namely the level of existing knowledge, and specifically the level of mastery the user has achieved in the different grammatical structures involved in putting together English sentences. Information on the user's proficiency in English can be used to solve two problems faced by the current implementation of the system. The first problem is how to handle the multiple parses of a sentence often obtained by the error analysis component; with a model of the user's grammatical proficiency, the system can select the parse that best reflects the performance that can be expected of the user at this stage in his or her acquisition process. The second problem is how to narrow the tutorial feedback to the subset of errors which are relevant to the user; "relevance" is defined as excluding those errors which occur on well-mastered structures (and which are really just careless mistakes, so instruction would be unnecessary as the student already has that knowledge) and those which are far beyond the user's ability to grasp at this level of understanding (and thus should not be explained, lest the user become confused and frustrated). We are presently developing a user model component for the ICICLE system to address those two issues.

Three Levels of Knowledge

The user model is essentially a representation of the user's ability to correctly use each of the grammatical "features" of English. These features include aspects of grammar such as pluralizing a noun with "+S," or making appropriate use of the past tense. The information stored about each of these features represents the observations made by the system based on the performance it has observed over the submission of multiple pieces of writing by a given user. If the user typically uses a given feature correctly, its corresponding element in the model will be marked "acquired." Conversely, consistent violation of a grammar rule will cause it to be marked "unacquired."[1]

We also wish to represent a third realm of proficiency, based on Vygotsky's observations about the acquisition of cognitive skills. He used the term *Zone of Proximal Development* (ZPD) to capture that subset of the skill which the learner is about to master [67]. Krashen's observation that at each step of language learning there is some set of grammar rules which the learner is "due to acquire" [36] effectively applies this theory to our domain. While acquired and unacquired constructions can be identified through consistent positive or negative performance, Ellis [22] identifies those structures on the verge of acquisition as exhibiting free variation in use, much of which will be grammatically inappropriate, before the user's level of understanding reaches the point where he can use them consistently correctly. Observation of inconsistent behavior should therefore clearly flag for ICICLE those features which should be marked "ZPD" for this learner.

An established set of these tags should enable ICICLE's error identification process to proceed on the premise that future user performance can be predicted based on the patterns of the past. If the tags have been assigned in the model based on the performance of the user to date, then if a feature has been marked "acquired," the user tends to execute it correctly, whereas a feature marked "unacquired" indicates that it is usually broken; the system can therefore generally prefer parses which use rules representing well-formed constituents associated with "acquired" features, mal-rules from the "unacquired" area, and either correct rules or mal-rules for those features marked "ZPD." If the system is then attempting to analyze the sentence, "My brother like baseball," and the model indicates that the user's mastery of subject/verb agreement is well-established but his mastery of plural nouns is not, it may prefer a parse containing a mal-rule which marks "brother" as a plural noun which is missing an "+S" ending over an interpretation in which agreement is missing from the verb.

The SLALOM model

We capture these different levels of proficiency on grammatical structures in a model called SLALOM ("Steps of Language Acquisition in a Layered Organization Model").The basic idea behind SLALOM is to divide the English language into a set of feature hierarchies (e.g., morphology, types of noun phrases, types of relative clauses). Within each hierarchy, the specific entries are ordered according to the order in which they are acquired. Evidence suggests that the acquisition sequence is relatively fixed regardless of the first language [29], [20], [3] and it may be very similar to the order in which the language is acquired as a first language. For example, evidence suggests that the +ING on verbs is acquired early, following it is the +S plural ending on NPs, followed by the +ED past tense ending, etc. This ordering will be reflected in the morphology hierarchy. We then coordinate the acquisition progression between the hierarchies by grouping items across multiple hierarchies into "layers" indicating sets of features which are generally acquired together. So, for example, if adjective noun clauses (from the NP hierarchy) are acquired at about the same time as irregular past tense forms (in the morphology hierarchy) then these forms would be positioned in the "same layer" across the hierarchies.

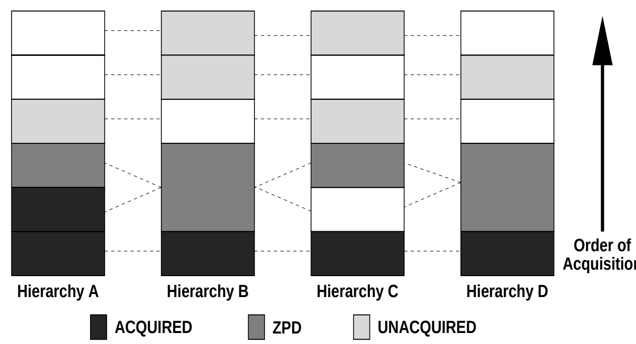

Figure 3: SLALOM Diagram

SLALOM's tags will be initialized following the first performance analysis of a new user's writing. Those features he or she has used consistently correctly will receive "acquired" tags, those used incorrectly "unacquired" tags, and those in variation "ZPD" tags. With each analysis of a new piece of writing from the student, these observations will be augmented with new and potentially different data, as features originally tagged as part of the ZPD exhibit correct usage and features originally tagged "unacquired" begin to show signs of variation and move into the ZPD. New data will result in the SLALOM tags being revised to reflect the user's developing knowledge. Because SLALOM represents an expected order of acquisition, the likely path of the ZPD would be to move "up" in the stacks.

1. Absence of a grammar feature in the writing of a user could either mean it is unknown and is being avoided or that there has been no opportunity yet to use it; how ICICLE decides between these possibilities is covered in the section "Reasoning on Partial Evidence" to follow.

We are aware of the difficulty of performing accurate parses in the initial evaluation of a user without the support of SLALOM's tags. We intend to investigate a two-pass approach, where a first pass through the piece evaluates such crude measures of writing competence as mean number of words per utterance or complexity of clause structure. Such observations woud provide a general idea of a global user competence level on which ICICLE could base its initial decisions.

Regardless of how SLALOM receives its initial markings, however, it is clear that its accuracy will improve greatly over time. Because the system is intended to be used by an individual over many pieces of writing, it will have access to a continually growing corpus of user-produced utterances.

Although it has been argued [54], [8] that creating and maintaining a detailed user model in a system involving natural language interaction is a very difficult task, and that evidence for user modeling in such a system is likely to be "poor in both quantity and quality," [54] we feel that SLALOM is not subject to the same limitations as the dialogue systems on which these observations were based for two reasons: first, the number of user utterances it has access to is much larger because they are not artifacts of natural interaction but fed to the system in large batches; and second, the user knowledge that is measured by SLALOM is not that which is *communicated* by these utterances (semantic content), but that which is *exhibited* by the utterances (syntactic content). This both increases the data extracted from each sentence and removes a lot of ambiguity, making it a far more accessible task for a machine to judge the extent of user knowledge.

Even under the best of circumstances, however, there will be times when the system will be inaccurate in its interpretations of the user's errors and their corrections (even as a human proofreader would). While we plan to provide ways for the user to provide feedback regarding the system's decisions (e.g., their choice of a possible correction), it is clear that any misinterpretation that remains would have serious implications for the updating of the language and user models within the overall system. Because of this concern, we are beginning to look at non-Bayesian statistical methods (e.g., Dempster-Shafer) that will allow us to incorporate confidence measures and partial belief into how we change our models over time.

Reasoning on Partial Evidence

Another problem with relying on SLALOM tags in order to make decisions based on user competency is the necessity of making judgments on user competence from incomplete information; ICICLE will not always have empirical data covering all features in SLALOM. An unacquired feature may have been absent due to avoidance, or an acquired feature absent due to lack of opportunity. We therefore must establish a method by which the system can infer a fuller description of user proficiency than is directly displayed in his or her past use of language forms.

Although the SLALOM feature relationships will be based on a general learner profile and not on the individual, they can serve to supplement the solid data we have on a specific learner. If an item in SLALOM has not yet received a tag, but it is below those items marked "ZPD" in SLALOM (or perhaps even below those that are "acquired"), it should be considered acquired. Likewise, one above the ZPD or above "unacquired" structures should be considered unacquired. Those at the same layer as the ZPD should also be part of it.

Once a user has begun to attempt a given construction, whether successfully or unsuccessfully, his or her performance will determine the marking on that construction in SLALOM and the model's organization will be irrelevant. The system therefore only has to rely on stereotypical data in novel situations; if a learner is acquiring features out of order due to instructional emphasis in the classroom, then his or her markings will reflect this and the system's decisions will be based on the individual, not the population of second language learners as a whole.

Developing the Model

We are applying both existing research on second language acquisition and analysis of our corpus of writing samples from our learner population toward establishing the relationships between features in the SLALOM model which will be essential for formalizing the design and implementing it within our system. Once a panel of human judges has identified general levels of ability for each writing sample, performance statistics from these samples will be obtained in order to determine which structures are being executed correctly and incorrectly at different stages of the learning process. Preliminary work on this corpus has indicated statistically significant groupings of errors at different levels of ability (indicating that writers of different skill levels commit different errors). Ongoing work will further specify these groupings in order to finalize the SLALOM design.

GENERATING A RESPONSE

The Response Generation component is another focus of current research. After the entire text has been analyzed by the Error Identification component, the text, along with the error results (for errors deemed at the user's level) and annotations from the mal-rules, will be passed to the Response Generator. The Response Generator is responsible for generating the tutorial feedback.

The Anatomy of a Response

We view the response itself as potentially differing along five dimensions: content, method, form, history, and manner. Content refers to the factual knowledge included in each response. It is the information that the system wants the student to gain from the correction. Content determination is influenced by the annotations on the errors that were passed from the Error Identification component, the probable source of the error, information about the languages involved, and information from the acquisition model.

Method describes the way in which the content is organized and presented to the listener. For example, the same basic content may be relayed to the student in various ways, such as by including information about similarities/differences between ASL and English, or by a discussion of a particular feature in English. The method chosen will be dependent on the kind of information available in the language model (e.g., about similarities/differences between ASL and English), the kind of information available about various English constructs involved in the error, and by the student's receptivity to various kinds of information in the correction. We have done preliminary analysis of text-only responses to actual sentences from our corpus of student writing in order to construct models of effective human correction techniques for this component of the response generator.

Form ensures that appropriate background knowledge has been included, and is responsible for structuring information to ensure that appropriate rhetorical resources are used in realizing the chosen method.

History is concerned with making the response contextually aware by making reference to earlier tutorial information and to knowledge known to the user.

Manner is concerned with the style of the actual language employed in the realization of the response as actual English text. As was the case for method, manner choices will partly depend on information in the user model. In particular, it is here that the user's language acquisition level will be taken into account to ensure that the constructions chosen are understandable to the student. Additionally, the system will favor using constructions that the student is currently acquiring when choices exist.

A Multimodal Response

There is some question concerning whether tutoring in the language being acquired will be effective [21], [66], [37], [15], [17]. One of the current directions of the project is to investigate presenting (some of the) tutorial information in a language that is closer to the "native" language of the students (i.e., ASL). This will hopefully make the information delivered easier for the student to understand and integrate, and may allow the student to feel more comfortable with the system and thus take more initiative without the intervention of human signers. As an example, in describing the evaluation of a grammar checker for deaf students, Loritz [42] indicates that some of the students would not correct their writing without the aid of a human who would interpret the system's messages to the student.

To address this concern we are looking at issues surrounding the incorporation of an animated (or possibly video-based) sign agent into our tutoring system. We must decide what information to sign and how to combine (both spatially and temporally) the signed and textual explanations. At first glance it may seem that the question of what information to sign is rather independent of the incorporation of a sign agent — i.e., it may appear that the sign agent could just be seen as an alternate method for displaying tutorial information. Upon further inspection one finds that it brings unique abilities involved in including information that may perhaps only be able to be rendered in sign. In addition it provides a framework for exploring appropriate media selection for different types of information and for exploring multimodal generation in visual modalities.

Acquisition of written English is a difficult task for people who are deaf because they do not have access to the auditory channel through which so much valuable language data is acquired. Here we attempt to provide more understandable tutoring to compensate for the "missing channel" of information (i.e., audio).

A series of studies investigating human-to-human tutoring are planned to determine when an animated sign agent should be used. The sign agent may have the advantage of being understandable when English text explanations might not be. In addition, some types of information may be inherently easier to describe in sign instead of text.

However, the visual nature of the signing medium presents some inherent challenges that must be addressed to gain the full benefit of this additional modality. We must first investigate how/whether it is possible for a user to attend to multiple information sources (e.g., an animated sign agent and a textual display) in the same medium and then determine how such coordination can be done.

CONCLUSION

We have focused in this paper on an approach to the problem of English literacy among deaf students: a tutorial tool which can supplement classroom instruction by providing a student with feedback on grammatical errors in a text while tracking that student's acquisition progress and tailoring the feedback to that individual. The eventual system will incorporate a unique user language model, flexible and robust tutorial planning, and multimodal capabilities via the inclusion of an animated sign agent.

ACKNOWLEDGMENTS

This work has been supported by NSF Grants #GER-9354869 and #IIS-9978021. It is part of the Center for Applied Science and Engineering in Rehabilitation of the University of Delaware and the duPont Hospital for Children, and has been supported by the Nemours Foundation.

REFERENCES

1. Anderson, J.J. *Deaf Student Mis-Writing, Teacher Mis-Reading: English Education and the Deaf College Student*. Linstok Press, Burtonsville, MD, 1993.
2. Baker, C. How does "sim-com" fit into a bilingual approach to education? In Caccamise, F. (ed), *American Sign Language in a Bilingual, Bicultural Context: Proceedings of the Second National Symposium on Sign Language Research and Teaching*, Silver Spring, MD, 1978. National Association of the Deaf.
3. Bailey, N., Madden, C., and Krashen, S.D. Is there a `natural sequence' in adult second language learning? *Language Learning*, 24(2):235-243, 1974.
4. Barnum, M. In support of bilingual/bicultural education for deaf children. *American Annals of the Deaf*, 129:404-408, 1984.
5. Bellman, K., Poizner, H., and Bellugi, U. Invariant characteristics of some morphological processes in American Sign Language. *Discourse Processes*, 6:199-233, 1983.
6. Bellugi, U., and Fischer, S. A comparison of sign language and spoken language: Rate and grammatical mechanisms. *Cognition*, 1:173-200, 1972.

7. Bornstein, H., and Pickett, J. Coordination of spoken and Signed English. Unpublished paper, Gallaudet College, Washington D.C., 1976.

8. Cawsey, A. *Explanation and Interaction: The Computer Generation of Explanatory Dialogues*. MIT Press, Cambridge, MA, 1993.

9. Cokely, D. *Pre-college Programs: Guidelines for Manual Communication*. Gallaudet College, Washington, D.C., 1979.

10. Collins, A., and Brown, J.S. The computer as a tool for learning through reflection. In Mandl, H., et al (eds), *Learning Issues for Intelligent Tutoring Systems*, 1-18, Springer-Verlag, NY, 1988.

11. Conrad, R. The reading ability of deaf school-leavers. *British Journal of Educational Psychology*, 147:138-148, 1977.

12. Corina, D., and Sandler, W. On the nature of phonological structure in sign language. *Phonology*, 10:165-207, 1993.

13. Coulter, G.R. On the relevance of traditional phonological analysis to the abstract pattern found in ASL and other sign languages. In Napoli, D.J., and Kegl, J.A. (eds), *Bridges Between Psychology and Linguistics: A Swarthmore Festscrhift for Lila Gleitman*. Lawrence Erlbaum Associates, Hillsdale, NJ, 1991.

14. Crandall, K. *A Study of the Production of Chers and Related Sign Language Aspects by Deaf Children Between the Ages of Three and Seven Years*. PhD thesis, Northwest University, 1974.

15. Cummins, J. The role of primary language development in promoting educational success for language minority students. In *Schooling and Language Minority Students: A Theoretical Framework*. California State University, Evaluation, Dissemination, and Assessment Center, Los Angeles, 1981.

16. Davies, S.N. The transition toward bilingual education of deaf children in Sweden and Denmark: Perspectives on language. *Sign Language Studies*, 71:169-195, 1991.

17. DiPietro, R.J. *Strategic Interaction: Learning Languages through Scenarios*. Cambridge University Press, 1987.

18. DiFrancesca, S. Academic achievement test results of a national testing program for hearing impaired students - United States: Spring, 1971. (series d, no. 9), Gallaudet College, Office of Demographic Studies, Washington, DC, 1972.

19. Drasgow, E. Bilingual / bicultural deaf education: An overview. *Sign Language Studies*, 80:243-266, 1993.

20. Dulay, H.C., and Burt, M.K. Natural sequences in child second language acquisition. *Language Learning*, 24:37-53, 1974.

21. Eddy, P. Does foreign language study aid native language development? *ERIC/CLL News Bulletin*, 1-2, 1978.

22. Ellis, R. *The Study of Second Language Acquisition*. Oxford University Press, Oxford, 1994.

23. Erting, C. Language policy and deaf ethnicity in the United States. *Sign Language Studies*, 19:139-152, 1978.

24. Gee, J., and Goodhart, W. Nativization, linguistic theory, and deaf language acquisition. *Sign Language Studies*, 49:291-342, 1985.

25. Grishman, R., Macleod, C., and Meyers, A. Comlex syntax: Building a computational lexicon. In *Proceedings of the 15th International Conference on Computational Linguistics*, Kyoto, Japan, 1994.

26. Grosjean, F. A study of timing in a manual and a spoken language: American Sign Language and English. *Journal of Psycholinguistic Research*, 8:379-405, 1979.

27. Gutiérrez, P. A preliminary study of deaf educational policy. *Bilingual Research Journal*, 18(3-4):85-113, 1994.

28. Hatch, E. Simplified input and second language acquisition. In Andersen, R. (ed), *Pidginization and Creolization as Language Acquisition*, 64-86, Newbury House, 1983.

29. Ingram, D. *First Language Acquisition: Method, Description, and Explanation*. Cambridge University Press, Cambridge; New York, 1989.

30. Jacobs, R. Just how hard is it to learn ASL? The case for ASL as a truly foreign language. In Lucas, C. (ed), *Multicultural Aspects of Sociolinguistics in Deaf Communities*. Gallaudet University Press, Washington, DC, 1996.

31. Jeffers, J., and Barley, M. *Speechreading (Lipreading)*. Charles C. Thomas Publisher, Springfield, IL, 1971.

32. Johnson, R.E., Liddell, S., and Erting, C. Unlocking the curriculum: Principles for achieving access in deaf education. *Gallaudet Research Institute Working Paper*, 89(3), 1989.

33. Kaplan, J.D., and Holland, V.M. Application of learning principles to the design of a second language tutor. In Holland, V.M., Kaplan, J.D., and Sams, M.R. (eds), *Intelligent Language Tutors: Theory Shaping Technology*, 273-287. Lawrence Erlbaum Associates, Mahwah, NJ, 1995.

34. Klima, E., and Bellugi, S. *The Signs of Language*. Harvard University Press, Cambridge, MA, 1979.

35. Kluwin, T. The grammaticality of manual representations of English in classroom settings. *American Annals of the Deaf*, 127:417-421, 1981.

36. Krashen, S.D. *Principles and Practice in Second Language Acquisition*. Pergamon Press, NY, 1982.

37. Legarreta, D. The effects of program models on language acquisition by Spanish speaking children. *TESOL Quarterly*, 13(4):521-534, 1979.

38. Lessow-Hurley, J. *The Foundations of Dual Language Instruction*. Longman Publishers USA, NY, 1996.
39. Liddell, S.K. *American Sign Language Syntax*. Mouton Publishers, 1980.
40. Liddell, S.K., and Johnson, R.E. American Sign Language: The phonological base. *Sign Language Studies*, 64:195-206, 1989.
41. Lieberman, P. *On the Origins of Language: An Introduction to the Evolution of Human Speech*. MacMillan Publishing Company, New York, 1975.
42. Loritz, D. GPARS: A suite of grammar assessment systems. In Holland, V.M., Kaplan, J.D., and Sams, M.R. (eds), *Intelligent Language Tutors: Theory Shaping Technology*, 121-133. Lawrence Erlbaum Associates, Mahwah, NJ, 1995.
43. Marmor, G., and Petitto, L. Simultaneous communication in the classroom: How well is English grammar represented? *Sign Language Studies*, 23:99-136, 1979.
44. Maxwell, M., and Bernstein, M. The synergy of sign and speech in simultaneous communication. *Applied Psycholinguistics*, 6:63-81, 1985.
45. Michaud, L.N., and McCoy, K.F. Planning text in a system for teaching English as a second language to deaf learners. In *Proceedings of Integrating Artificial Intelligence and Assistive Technology*, an AAAI-98 workshop, Madison, WI, 1998.
46. Michaud, L.N., and McCoy, K.F. Supporting Intelligent Tutoring in CALL by Modeling the User's Grammar. In *Proceedings of the Thirteenth International FLAIRS Conference*, Orlando, FL, 2000.
47. Neuroth-Gimbrone, C., and Logiodice, C. A cooperative bilingual language program for deaf adolescents. *Sign Language Studies*, 74:79-91, 1992.
48. Owen, E., and Sweller, J. What do students learn while solving mathematics problems? *Journal of Educational Psychology*, 77:272-284, 1985.
49. Quigley, S., and Paul, P. ASL and ESL? *Topics in Early Childhood Special Education*, 3(4):17-26, 1984.
50. Rose, C.P. A framework for robust semantic interpretation. In *Proceedings of the 1st Meeting of the North American Chapter of the Association for Computational Linguistics*, 311-318, Seattle, WA, 2000.
51. Schneider, D., and McCoy, K.F. Recognizing syntactic errors in the writing of second language learners. In *Proceedings of COLING-ACL*, 1198-1204, Montreal, Quebec, Canada, 1998.
52. Silvester, T. Managing a bilingual programme (BSL and English) in a school for deaf children. In *Proceedings of the International Congress on Education of the Deaf*, Tel Aviv, Israel, 1995.
53. Sleeman, D. Inferring (mal) rules from pupil's protocols. In *Proceedings of ECAI-82*, 160-164, Orsay, France, 1982.
54. Sparck Jones, K. Tailoring output to the user: What does user modelling in generation mean? In Paris, C.L., Swartout, W.R., and Mann, W.C. (eds), *Natural Language Generation in Artificial Intelligence and Computational Linguistics*, chapter 8, 201-225, Kluwer Academic Publishers, Boston, 1991.
55. Stevens, R. Education in schools for deaf children. In Baker, C., and Battison, R. (eds), *Sign Language and the Deaf Community*, 177-191. National Association of the Deaf, Silver Spring, MD, 1980.
56. Stokoe, W. Sign language structure. *Studies in Linguistics: Occasional Papers* 8, Department of Anthropology and Linguistics, University of Buffalo, 1960.
57. Strong, M. *Language Learning and Deafness*. Cambridge University Press, New York, 1988.
58. Strong, M., and Charlson, E. Simultaneous communication: Are teachers attempting an impossible task? *American Annals of the Deaf*, 132:376-382, 1987.
59. Strong, M., and Prinz, P.M. A study of the relationship between American Sign Language and English literacy. *Journal of Deaf Studies and Deaf Education*, 2(1):37-46, 1997.
60. Suri, L.Z. Language transfer: A foundation for correcting the written English of ASL signers. Technical Report TR-91-19, Dept. of CIS, University of Delaware, 1991.
61. Suri, L.Z., and McCoy, K.F. A methodology for developing an error taxonomy for a computer assisted language learning tool for second language learners. Technical Report TR-93-16, Dept. of CIS, University of Delaware, 1993.
62. Sweller, J. Cognitive load during problem solving: Effects on learning. *Cognitive Science*, 12:257-285, 1988.
63. Swisher, M.V. Similarities and differences between spoken languages and natural sign languages. *Applied Linguistics*, 9:343-356, 1988
64. Swisher, M.V. The language learning situation of deaf students. *TESOL Quarterly*, 23(2):239-257, 1989.
65. Tarone, E. Systematicity and attention in interlanguage. *Language Learning*, 32(1):69-84, 1982.
66. Troike, R. Research evidence for the effectiveness of bilingual education. *Nabe Journal*, 3:13-24, 1978.
67. Vygotsky, L.S. *Thought and Language*. MIT Press, Cambridge, MA, 1986.
68. Weischedel, R.M., Voge, W.M., and James, M. An artificial intelligence approach to language instruction. *Artificial Intelligence*, 10:225-240, 1978.
69. Wrightstone, J.W., Aronow, M.S., and Moskowitz, S. Developing reading test norms for deaf children. *American Annals of the Deaf*, 108:311-316, 1963.