Providing Intelligent Language Feedback for Augmentative Communication Users

Christopher A. Pennington and Kathleen F. McCoy

Applied Science and Engineering Laboratories
Department of Computer and Information Sciences
University of Delaware / duPont Hospital for Children
Wilmington, DE 19899
penningt@asel.udel.edu

Abstract. People with severe speech and motor impairments (SSMI) can often use augmentative communication devices to help them communicate. While these devices can provide speech synthesis or text output, the rate of communication is typically very slow. Consequently, augmentative communication users often develop telegraphic patterns of language usage. A natural language processing technique termed *compansion* (compression-expansion) has been developed that expands uninflected content words (i.e., compressed or telegraphic utterances) into syntactically and semantically well-formed sentences.

While originally designed as a rate enhancement technique, compansion may also be viewed as a potential tool to support English literacy for augmentative communication users. Accurate grammatical feedback from ill-formed inputs might be very beneficial in the learning process. However, the problems of dealing with inherently ambiguous errors and multiple corrections are not trivial. This paper proposes the addition of an adaptive user language model as a way to address some of these difficulties. It also discusses a possible implementation strategy using grammatical mal-rules for a prototype application that uses the compansion technique.

1 Introduction

People with severe speech and motor impairments (SSMI) can often use augmentative communication devices to help them communicate. While these devices can provide speech synthesis or text output, the rate of communication is typically very slow (most users average less than 10 words per minute). Consequently, augmentative communication users can often develop telegraphic patterns of language usage, especially if the disability occurs at an early age.

* This work has been supported by a Rehabilitation Engineering Research Center Grant from the National Institute on Disability and Rehabilitation Research of the U.S. Department of Education (#H133E30010) and by NSF Grant # IRI-9416916. Additional support has been provided by the Nemours Research Program.

Although this functional style of communication is perfectly adequate for many situations, there are circumstances in which complete, grammatical English sentences are necessary to ensure proper communication and understanding. In addition, there are several obvious educational and psychological reasons for providing the ability to communicate in a literate manner. One in particular is to help dispel the general tendency of our society to automatically associate an inability to speak (or speak understandably) with a cognitive impairment or lack of intelligence.

To help address these concerns, a natural language processing technique termed *compansion* (compression-expansion) has been developed that expands uninflected content words (i.e., compressed or telegraphic utterances) into syntactically and semantically well-formed sentences [DM92], [MDJ+94]. For example, given the input `John go store yesterday`, an intelligent augmentative communication system using compansion might produce "John went to the store yesterday."

Originally, compansion was designed as a rate enhancement technique for word- or symbol-based augmentative communication systems; that is, its primary purpose was to enable users to express themselves more quickly either using a speech synthesizer or for performing writing tasks. However, compansion can also be viewed as a potential tool to support English literacy efforts for augmentative communication users. This paper discusses the mechanisms needed to provide compansion with an enhanced ability to identify and correct language errors. A parallel effort for improving the written English of deaf people who are American Sign Language natives is also in progress [SM93], [MPS96].

2 Issues in Providing Intelligent Feedback

By providing accurate, grammatical feedback from ill-formed input, the compansion technique can be used to help facilitate the language development process, especially for users of symbol-based communication devices. At the very least, compansion can provide language reinforcement to the augmentative communication user through speech output and/or written text. This is analogous to the situation where a teacher or tutor would provide corrective instruction either verbally or visually (e.g., writing on a chalkboard).

Of course, there are several difficulties that must be dealt with to successfully provide *accurate* feedback. A basic issue is the ability to detect multiple errors in an ill-formed input. In addition, there may be potentially ambiguous interpretations of what those errors are, so properly identifying the errors is a major step. For example, `John gone to the store` could be incorrect because of a wrong past tense form ("John went to the store") or a missing auxiliary verb ("John had gone to the store").

Often, the combination of these factors will generate a whole set of possible corrections. Deciding which correction is the most appropriate can be very difficult. For example, `The girl like John` appears to have a subject-verb agreement error and could be corrected as "The girls like John" or "The girl likes

John." However, for certain augmentative communication users, it could also be interpreted as "The girl was liked by John" or "The girls were liked by John." In some instances, the best suggestion for correction may be partially dependent on the specific user's language patterns.[2] The compansion technique already addresses these issues to some degree; nevertheless, there are several limitations that must be overcome in order to give truly intelligent feedback.

3 Overview of Compansion

The core of the compansion approach is a semantic parser that interprets input based on the use of case frames [Fil68], [Fil77]. Case frames are conceptual structures that represent the meaning of an utterance by describing the semantic cases or roles that each of the content words has in relationship with the others. In practice, the semantic parser designates the primary verb as the main component of the expression: all other words in the input are used to fill semantic roles with respect to the main verb that is chosen. While not an exhaustive or necessarily ideal list, we have adopted a set of semantic roles that have proven adequate for our purposes.

AGEXP (AGent/EXPeriencer) is the object doing the action, although for us the AGEXP does not necessarily imply intentionality, such as in predicate adjective sentences (e.g., *John* is the AGEXP in "John is happy"). THEME is the object being acted upon, while INSTR is the object or tool etc. used in performing the action of the verb. GOAL can be thought of as a receiver, and is not to be confused with BENEF, the beneficiary of the action. For example, in "John gave a book to Mary for Jane", *Mary* is the GOAL while *Jane* is the BENEF. We also have a LOC case which describes the event location (this case may be further decomposed into TO-LOC, FROM-LOC, and AT-LOC), and TIME which captures time and tense information (this case may also be subdivided).

As an example, given the input `John go store`, *go* would be selected as the main VERB, *John* would fill the role of AGEXP, and *store* would be assigned a LOC (LOCation) role.[3]

The semantic parser attempts to construct the most likely interpretation by filling in the various semantic roles in a way that makes the most sense semantically. It does this by employing a set of scoring heuristics which are based on the semantic types of the input words and a set of preferences which indicate which roles are the most important to be filled for a particular verb and what

[2] Of course, the context in which the expression occurs is extremely important; however, in many cases it is not possible for an augmentative communication system to have access to both sides of the entire conversation, although advances in continuous speech recognition appear promising. At this point we only draw simple inferences from the partial context (e.g., tense changes); therefore, unless otherwise noted, utterances are considered in isolation.

[3] Note that it is ambiguous at this point whether *store* should be a TO-LOC or a FROM-LOC.

semantic types the fillers of the roles should have. The parser relies on a set of scoring heuristics (based on the preferences) to rate the possible interpretations (i.e., different ways of filling in the case frame) it comes up with [JDMP91]. "Idiosyncratic" case constraints specify which roles are mandatory or forbidden given a specific verb (or class of verbs). This captures, for example, the difference between transitive and intransitive verbs, where (in general) transitive verbs are required to have a theme, but a theme is forbidden with intransitive verbs. Other heuristics reflect general case preferences, including case importance (e.g., most verbs prefer THEMEs to be filled before BENEFiciaries), case filler (e.g., action verbs prefer animate AGEXPs), and case interactions (e.g., a human AGEXP might use an INSTRument, but an animal like a dog probably would not).

After all of the ratings for the various case preferences are assigned, they are combined together to produce a final score for each possible interpretation that the semantic parser produces. Any interpretation with a value less than a specified cut-off value is discarded, and the rest are ordered according to score and passed on for further processing. So, two possible parser interpretations of the input `apple eat John` might look like the following (DECL denotes a declarative expression):

```
(70 DECL
    (VERB (LEX EAT))
    (AGEXP (LEX JOHN))
    (THEME (LEX APPLE))
    (TENSE PRES))}

(20 DECL
    (VERB (LEX EAT))
    (AGEXP (LEX APPLE))
    (THEME (LEX JOHN))
    (TENSE PRES))
```

The first interpretation corresponds to a sentence such as "John eats the apple" while the second, lower rated interpretation corresponds to the interpretation "The apple eats John". Obviously, "John eats the apple" should be preferred over "The apple eats John" (in fact, the latter interpretation would almost always be discarded because of the previously mentioned cut-off value). The more likely interpretation is chosen because the preferences associated with the VERB *eat* strongly indicate a preference for an animate AGEXP (i.e., a human or animal) and a THEME that is edible.

Notice that the semantic reasoning that is the heart of compansion does not address syntactic realization issues. The semantic parser indicates that it prefers the interpretation where `John` is doing the eating and the `apple` is the thing being eaten. This preferred interpretation has many different syntactic realizations such as: "John eats the apple", "The apple was eaten by John", "It was the apple that John ate", etc.

4 Improving the Scoring Methodology

This rating system has proven useful for developing an initial research prototype of the compansion technique, allowing distinctions to be made about some important conceptual relationships. However, it must be improved upon if it is to be used to provide appropriate corrective feedback for augmentative communication users in the process of developing literacy skills. In addition, as pointed out above, these scores only capture preferences among semantic interpretations (e.g., "John likes pizza" vs. "John is liked by pizza") and provide little help with differentiating among several potential syntactic expressions of the same (or very similar) concept (e.g., "John likes pizza", "John likes the pizza", "The pizza is liked by John", etc.). This latter issue will be addressed in the next section.

Currently, most of the preference ratings for cases are based on intuition and the rules for combining scores are somewhat arbitrary. This is not sufficient to ensure a consistently reasonable set of possible corrections. Statistical data from tagged corpora should be used to provide better supported values for the ratings. Methods outlined in [All95] suggest taking context into account as well as frequency when computing probabilities. A specific treatment of this approach for verb subcategorization is detailed in [UEGW93] and appears to be quite in line with our purposes. Information from lexical databases such as WordNet [Mil95] is also being integrated to help improve part-of-speech and word sense preferences, as well as semantic classification information.

Furthermore, the functions used in combining scores should reflect an appropriate and well-established probabilistic method (see [Cha93] for an overview of several possible algorithms). Related to this, the final scores should be normalized to provide a general measure of the appropriateness of an interpretation as well as to allow more objective comparisons between sentences.

Since the primary goal in this case is to promote literacy and not necessarily rate enhancement, a comprehensive list of choices should always be generated. This will increase the chances of augmentative communication users always finding a correct representation of what they want to express.[4] This does not detract, however, from the goal of presenting the best correction first whenever possible.

5 Accounting for Syntactic Variations

Besides the semantic parser, compansion also contains some rudimentary inferencing principles based on general observations of "telegraphic" forms of expression found in some sign languages and pidgins. For example, if no main verb is found, it will attempt to infer the verbs *be* and/or *have*, taking into account the possible roles of the other words in the input. In a similar manner, if there is no valid agent, it will infer the pronouns *I* or *you*, depending on whether the

[4] Of course there will always be instances in which compansion may be unable to correctly interpret the user's intended meaning. Even humans have a difficult time with that task from time to time.

input is a statement or a question. These techniques allow us to interpret input like `happy ?` to mean "Are you happy?" (as one reasonable possibility). At this point, we are beginning to reason about mainly syntactic distinctions and in fact, compansion uses a "word order" parser that attempts to account for various expected telegraphic and word order variations that may be common in this population (e.g., determining whether the output should maintain the input word order or not which would dictate whether the system should generate "John likes Mary" or "Mary is liked by John").

Additional research has begun that investigates more fully the often telegraphic language patterns of augmentative communication users [MMP+94]. Knowing more about general language expressions used in augmentative communication should enable compansion to make better choices among syntactic variants of the user's intended communication. The proposed methodology for accomplishing this is to group the common language variations into a taxonomy that can assist error identification [SM93].

Although there may be general language variations that occur, it is also likely that each individual will have idiosyncratic patterns of expression (e.g., some users may never construct passive sentences), including commonly made errors. This information could be very useful for error identification and for determining the most appropriate correction(s). Thus, there is a need for both an individual and a general user language model [Chi89]. In addition, there is the possibility that an augmentative communication user's language abilities and preferences will change, especially if they are in the process of learning English literacy skills. This argues for a language model that can adapt to the user over time. This model will be essential for generating better interpretations, handling language errors intelligently, and providing additional feedback that may be helpful to the user.

6 An Adaptive User Language Model

In this section, we focus primarily on modeling syntactic expectations, given a specific user. We propose an adaptive user language model that requires several steps and relies on several different components to capture expectations for a particular user. First, a general language assessment model must be developed. This model will capture typical levels of literacy acquisition and indicate syntactic constructions that a person is expected to be able to use at a given level of acquisition. Second, this stereotypical model must be (possibly) modified for particular classes of users who may share common factors that could influence the process of language acquisition for that group (e.g., language transfer from a first language, prior literacy training, or speech therapy). The intermediate result of this design will be a model that captures expectations about syntactic structures being used by individuals that fall into various pre-defined levels of literacy expertise. The last component will be able to determine at what level to place a particular user, and to update the placement based on a carefully tracked history of interaction with that user. The final language model can then be used to help determine which suggested corrections are the most appropriate given the user's linguistic abilities and past language use.

6.1 SLALOM - A Language Assessment Model

Intuitively, people express themselves relatively consistently with a certain degree of language sophistication. Some grammar checkers rely on this concept when they assign a general "grade level" to a user's composition. Often this evaluation is based primarily on average sentence length (which is a very rough measure of the text's syntactic complexity). Knowing a person's writing "grade level" could help us immensely in choosing among various possible syntactic expressions. For instance, we would not expect someone who generally writes at a second grade level to use a complicated tense (e.g., past perfect) because that complexity is apparently beyond their current writing ability.

What we need is a mechanism that organizes syntactic constructions that are likely to be used together and can serve as the means for evaluating and predicting each user's English language proficiency. This "profile" can then be used to help determine a preferred interpretation when either the error or its underlying cause is ambiguous (e.g., when results from error identification suggest more than one possible correction for a single error).

To accomplish this, we propose the development of a language assessment model called SLALOM ("Steps of Language Acquisition in a Layered Organization Model") that is based on linguistic theories describing the processes by which people learn language. There is considerable linguistic evidence from research on both first language acquisition and second language acquisition that the acquisition order of language features is relatively consistent and fixed [Ing89], [DB74], [BMK74]. In fact, a stronger version of this statement is one of the central tenets of universal grammar theory (see for example, [Haw91] and [KH87]).

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) that are ordered from least to most complex. Features of similar complexity but in different hierarchies are arranged in layers that are then linked together to represent stereotypical "levels" of language ability.

Figure 1 contains an conceptual illustration of a piece of SLALOM.[5] We have depicted parts of four hierarchies in the figure: morphological syntactic features, noun phrases, verb complements, and various relative clauses. Within each hierarchy, the intention is to capture an ordering on the feature acquisition. So, for example, the model reflects the fact that the +ing progressive form of verbs is generally acquired before (and thus considered "easier to acquire" than) the +s plural form of nouns, which is generally acquired before the +s form of possessives, etc.

[5] Specific details of the feature hierarchies have been simplified and are given for example purposes only.

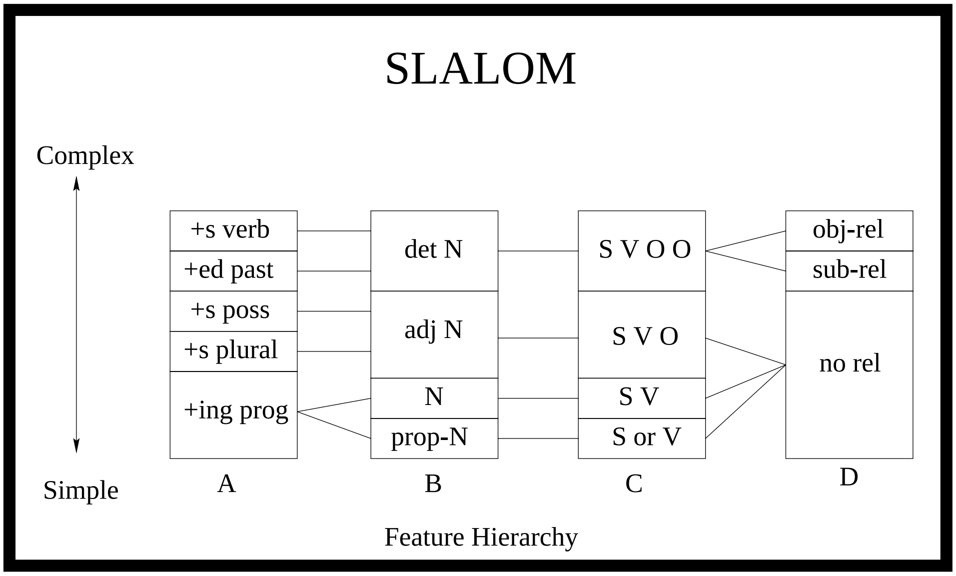

Fig. 1. Example of feature hierarchies in SLALOM

As mentioned before, the connections among the hierarchies are intended to capture sets of features which are acquired at approximately the same time. So, for example, the figure indicates that while the +s plural ending is being acquired, so too are both proper and regular nouns, and one and two word sentences. During this time, we do not expect to see any relative clauses.

These connections are derived from work in language assessment and grade expectations such as found in [Lee74], and [Cry82]. A combination of existing assessment tools will be needed to ensure adequate coverage of English language features and in some cases additional linguistics research may be required to develop an accurate and useful default model.

6.2 Customizing the Language Model

Once the default has been established, there must be a method for customizing the model on the basis of user characteristics which might cause either the simple/complex order within a hierarchy to change, or cause the predefined "levels" across the hierarchies to be altered. For instance, the model might be altered as a result of a study of texts from augmentative communication users. This study might result in an error taxonomy (mentioned earlier) which gives evidence (contrary to the views of universal grammar proponents) that the steps of language learning for augmentative communication users (as a group) were differently ordered from the standard expectations of English language acquisition. SLALOM can also be tailored to the needs of individual users via a series of "filters", one for each user characteristic that might alter the initial generic model. One possible filter could reflect how much and what kind of formal instruction the user has had in English literacy. For example, if the user's educational program stressed subject-verb agreement, this relatively complex feature might have already been learned, even though other features designated as "simpler" in the original model may remain problematic.

Once SLALOM has been completed for the population under consideration, we will presumably have a model of the order in which we expect our users to learn the English language. Essentially, we will need to "place" a particular user in the model. With this placement we will have a model of (1) what features we expect the student to have mastered and to be using consistently – these are features below the user's level in the model, (2) what features we expect the user to be using or attempting to use, but with limited success – these are features at the user's level, and (3) what features we do not expect to see used (correctly) – these are features above the user's level.

The initial placement of the student user in SLALOM will most likely be based on an analysis of a first input sample. Once this initial determination is made, further input from the user, as well as feedback given during any correction and tutorial phases, could cause the system to update the user's level in the model. It is important to note that although the default levels (i.e., cross-hierarchical connections) for the process of language acquisition will be somewhat predefined, the model is flexible enough to allow and account for individual variations beyond those represented by the initial model and its filters. In other words, additional information about each user's language usage gathered over time should provide a better and more accurate reflection of the current set of language features they are acquiring.

6.3 Adaptation Mechanisms

To realize a flexible model, a good history mechanism must be provided that can assist the language model in adapting to each individual's abilities and preferences. The history mechanism's responsibility is to update information in the user model based on experience with the augmentative communication user. Most of this information will be derived implicitly (e.g., analyzing expressive output to discover an especially problematic language feature), although a particular interface may allow explicit changes to the model.[6]

Potentially, there is a need for both a short-term and a long-term history mechanism. Short-term frequency data for errors and successes could be used to reassess the user's language abilities, especially when determining whether or not a specific language feature is known or in the process of being learned. This could be very helpful for deciding among several possible corrections as well as moving the user along the "steps" of the language model. Also, the

[6] This becomes more relevant if a tutoring component is being used to provide corrective responses.

prototypical language levels allow a system using this model to make reasonable default inferences when little knowledge is available. For example, if the user has not expressed a language feature before, the system can assume its acquisition level based on other features that are known.[7]

A long-term history mechanism would provide additional evidence for language change, as well as providing a way of adapting to the user's idiosyncratic language patterns. In addition, for tutorial purposes it might be useful to look for the user's avoidance of certain linguistic structures[8] since not all language difficulties are evident through error identification.

7 Developing a Compansion-Based Prototype

Up to this point, we have discussed two different aspects of a possible literacy aid for augmentative communication users: a semantic aspect based on conceptual information associated with words used in an input sentence, and a syntactic aspect based on expectations derived from the use of an adaptive language model. In this section, we discuss how these two aspects are being combined into an augmentative communication prototype.

The driving force of the process is a (primarily) syntactic grammar that is implemented in the Augmented Transition Network (ATN) formalism. An ATN parses sentences by encoding a grammar as a network of nodes and arcs (connecting the nodes) that is traversed as decisions are made about the input words (e.g., is it a noun, verb, etc.?). Registers containing specific information about the words and the parse may be passed through the network (and subnetworks), providing a powerful mechanism for reasoning about natural language. This formalism also allows arbitrary tests and actions to be performed during the parsing process; thus, many semantically-oriented tests can be incorporated directly into the grammar. Using this approach, we have encoded many aspects of the compansion technique into the grammar for this system; thus, the semantic "score" may be calculated as the grammar network is traversed. Below is a discussion of the changes needed to integrate an adaptive user language model into this application.

7.1 Using Mal-Rules to Encode Language Variations

The first step is to develop a syntactic grammar that is enhanced to capture the regular variants in the language use of augmentative communication users. A conceptual mechanism that could be used to simulate the language patterns would be mal-rules [Sle82], [WS83]. Mal-rules are grammar rules specifically coded to accept ill-formed structures that reflect expected language errors; however, additional information can be associated with these rules that indicates an error has been made and what the possible cause(s) might be. The rules would

[7] At this time it is not clear if the best strategy would be to assign the default as the minimum level, the highest level, or an average level.

[8] That we expect to see based on the perceived language level of the user.

handle observed telegraphic conventions (e.g., omitting forms of *be*) as well as any commonly detected irregularities (e.g., inverted word order). A similar method has been used for second language learning [Sch85].

A possible implementation of this approach is to construct a core grammar representing standard grammatical English and a separate set of mal-rules that captures common language variations of augmentative communication users. These mal-rules can be realized as an overlay of alternate arcs at the appropriate nodes within the ATN grammar. The resulting "modularity" will allow association of additional information with the mal-rules in a group-specific manner; for instance, we could construct appropriate error messages in this manner. If designed carefully, it should also be possible to (easily) use a different set of mal-rules (e.g., language patterns of a deaf person learning English as a second language) with the core grammar.

7.2 Implementing the User Language Model

In essence, this combination of mal-rules with the standard grammar comprises a "grammar" for all prototypical augmentative communication users. However, in order to individualize the grammar to specific users, a weighted grammar that assigns relative scores to each of its rules is proposed.[9] Usage frequency information from corpora and research data will be used as the initial weights for both the arcs of the standard grammar and the set of mal-rules. However, one complicating factor is that no large corpora exist for the language use of augmentative communication users. Thus, we must be careful in how the probabilities for the mal-rules are determined and rely mostly on data from standard text corpora.[10]

One possibility is that the initial values for the mal-rules will be predominantly stereotypical (i.e., reflecting the general relationships of the error taxonomy instead of being strictly frequency-based) and more sensitive to changes based on the user's interactions with the system. Some of the methods for dealing with sparse data [Cha93] may also be helpful. In addition, features representing the relative complexity of acquisition will be attached to the nodes of the grammar. In the absence of other information, this value may be helpful in discriminating among multiple interpretations.

Once this default structure has been defined and initialized, the scores and features of the grammatical arcs (including those representing the mal-rules) may be modified by interactions with a separate user model that contains the individual's specific language characteristics (with respect to SLALOM). This model will consist of long-term information including the following: what language features are known, unknown, or in the process of acquisition; an overall

[9] The most likely implementation is a probabilistic context-free grammar similar to those described by Charniak [Cha93] and Allen [All95].

[10] In a related effort, we are engaged in an ongoing process of collecting conversational data (keystrokes, "spoken" text and some video) from a few augmentative communication users and hope to take advantage of this information at some point. It is unclear if this will be generalizable, though.

measure of the user's language level (derived from the known language features); and historical data reflecting the user's language usage and error patterns. The latter information will be used to make changes to the grammar for each particular user.

Eventually, these changes will allow the grammar to adapt to the augmentative communication user's specific language style. Exact criteria for deciding when to change the feature acquisition values (e.g., from "acquiring" to "known") have not yet been determined, but essentially we can view the placement in SLALOM as highlighting language features (and corresponding mal-rules) that we expect the user to be learning at a given point in time. Thus, it gives us a glimpse into users' current language patterns by zeroing in on the mal-rules we expect them to be using at this point in their acquisition of English. Key to this process is the feedback provided by interactions where one of the suggested corrections is selected. This information will help to either confirm or modify the system's current "view" of the user. In any event, the mechanisms needed to implement these adjustments should be straightforward.

7.3 Processing Considerations

After a sentence is parsed, the identified errors will be returned and tagged appropriately with the mal-rule(s) thought to be responsible. In many cases, we cannot assume that there will be a one-to-one mapping between the identified mal-rules and the possible corrections. Confounding this issue is the strong possibility of multiple errors in each sentence, possibly interacting with each other; hence, it might be necessary to look at evaluating sets of mal-rules that are triggered instead of individual ones. At this time it is unclear what method will be best for determining the most likely set of mal-rules.

8 Future work

The most immediate need is to further specify the relationships of features within SLALOM and their likeliness to occur. In addition, while there is some evidence of what constitutes a "typical" telegraphic language pattern, more work must be done to classify these variations and to gain information on their frequency of use. Once this is accomplished, the data can be used in the modifications that will be made to the current compansion-based application as it integrates the adaptive language model. As discussed previously, it is thought that these changes will take the form of adding mal-rules and weighted features to the ATN, along with any necessary reasoning mechanisms. Adaptability will be addressed by superimposing a history mechanism that will adjust weights and other features based on experiences with the augmentative communication user's language choices and feedback selections.

Results from this work will be filtered back into a larger project called ICICLE (Interactive Computer Identification and Correction of Language Errors).

ICICLE currently encompasses the mechanisms for identifying errors in the written English of deaf people. As mentioned earlier, the design of corrective feedback mechanisms for that system is proceeding in parallel with the work described here. It is hoped that some of the semantic reasoning strategies in compansion will be of use to ICICLE as well.

Another essential component being designed for ICICLE concerns adaptive tutoring and explanation [MM97]. This module will be able to consult the adaptive language model for information to help customize its instruction for the individual user. Finally, at the present time, both ICICLE and the compansion technique are primarily concerned with clause- or sentence-level variations; however, it is important to note that many difficulties in English literacy occur at a discourse level (e.g., anaphora resolution). This is a major area of needed research.

9 Summary

The compansion technique has great potential for use as a tool to help promote literacy among users of augmentative communication systems. By providing linguistically correct interpretations of ill-formed input, it can reinforce proper language constructions for augmentative communication users who are in the process of learning English or who have developed telegraphic patterns of language usage. To accomplish this goal, several modifications to the existing compansion approach are proposed to improve the accuracy of the corrective feedback. The most significant change is the addition of an adaptive language model. This model initially provides principled defaults that can be used to help guide the identification and correction of language errors, adapting to each user's specific language abilities and patterns over time. Finally there is a discussion of using sets of grammatical mal-rules to integrate the language model into an existing application that uses the compansion technique.

References

[All95] James Allen. *Natural Language Understanding.* Benjamin/Cummings, Redwood City, CA, 1995.

[BMK74] Nathalie Bailey, Carolyn Madden, and Stephen D. Krashen. Is there a "natural sequence" in adult second language processing? *Language Learning*, 24(2):235–243, 1974.

[Cha93] Eugene Charniak. *Statistical Language Learning.* MIT Press, Cambridge, MA, 1993.

[Chi89] David N. Chin. KNOME: Modeling what the user knows in UC. In Alfred Kobsa and Wolfgang Wahlster, editors, *User Models in Dialog Systems*, Berlin Heidelberg New York Tokyo, 1989. Springer.

[Cry82] David Crystal. *Profiling Linguistic Disability.* Edward Arnold, London, 1982.

[DB74] Heidi C. Dulay and Marina K. Burt. Natural sequences in child second language acquisition. *Language Learning*, 24(1):37–53, 1974.

[DM92] Patrick W. Demasco and Kathleen F. McCoy. Generating text from compressed input: An intelligent interface for people with severe motor impairments. *Communications of the ACM*, 35(5):68–78, May 1992.

[Fil68] C. J. Fillmore. The case for case. In E. Bach and R. Harms, editors, *Universals in Linguistic Theory*, pages 1–90, New York, 1968. Holt, Rinehart, and Winston.

[Fil77] C. J. Fillmore. The case for case reopened. In P. Cole and J. M. Sadock, editors, *Syntax and Semantics VIII: Grammatical Relations*, pages 59–81, New York, 1977. Academic Press.

[Haw91] John A. Hawkins. Language universals in relation to acquisition and change: A tribute to Roman Jakobson. In Linda R. Waugh and Stephen Rudy, editors, *New Vistas in Grammar: Invariance and Variation*, pages 473–493. John Benjamins, Amsterdam / Philadelphia, 1991.

[Ing89] David Ingram. *First Language Acquisition: Method, Description, and Explanation*. Cambridge University Press, Cambridge; New York, 1989.

[JDMP91] Mark Jones, Patrick Demasco, Kathleen McCoy, and Christopher Pennington. Knowledge representation considerations for a domain independent semantic parser. In J. J. Presperin, editor, *Proceedings of the Fourteenth Annual RESNA Conference*, pages 109–111, Washington, D.C., 1991. RESNA Press.

[KH87] Edward L. Keenan and Sarah Hawkins. The psychological validity of the accessibility hierarchy. In Edward L. Keenan, editor, *Universal Grammar: 15 Essays*, pages 60–85. Croon Helm, London, 1987.

[Lee74] Laura L. Lee. *Developmental Sentence Analysis: A Grammatical Assessment Procedure for Speech and Language Clinicians*. Northwestern University Press, Evanston, IL, 1974.

[MDJ+94] Kathleen F. McCoy, Patrick W. Demasco, Mark A. Jones, Christopher A. Pennington, Peter B. Vanderheyden, and Wendy M. Zickus. A communication tool for people with disabilities: Lexical semantics for filling in the pieces. In *Proceedings of the First Annual ACM Conference on Assistive Technologies*, pages 107–114, New York, 1994. ACM.

[Mil95] George A. Miller. Wordnet: A lexical database for English. *Communications of the ACM*, pages 39–41, November 1995.

[MM97] Kathleen F. McCoy and Lisa N. Masterman. A tutor for deaf users of American sign language. In *Proceedings of Natural Language Processing for Communication Aids, an ACL/EACL '97 Workshop*, Madrid, Spain, July 1997.

[MMP+94] Kathleen F. McCoy, Wendy M. McKnitt, Christopher A. Pennington, Denise M. Peischl, Peter B. Vanderheyden, and Patrick W. Demasco. AAC-user therapist interactions: Preliminary linguistic observations and implications for Companion. In Mary Binion, editor, *Proceedings of the RESNA '94 Annual Conference*, pages 129–131, Arlington, VA, 1994. RESNA Press.

[MPS96] Kathleen F. McCoy, Christopher A. Pennington, and Linda Z. Suri. English error correction: A syntactic user model based on principled "malrule" scoring. In *Proceedings of UM-96, the Fifth International Conference on User-Modeling*, pages 59–66, Kailua-Kona, HI, January 1996.

[Sch85] Ethel Schuster. Grammars as user models. In *Proceedings of IJCAI 85*, 1985.

[Sle82] Derek H. Sleeman. Inferring (mal) rules from pupil's protocols. In *Proceedings of ECAL-82*, pages 160–164, Lisay, France, 1982.

[SM93] Linda Z. Suri and Kathleen F. McCoy. Correcting discourse-level errors in CALL systems for second language learners. *Computer-Assisted Language Learning*, 6(3):215–231, 1993.

[UEGW93] Akira Ushioda, David A. Evans, Ted Gibson, and Alex Waibel. Frequency estimation of verb subcategorization frames based on syntactic and multidimensional statistical analysis. In *Proceedings of the 3rd International Workshop on Parsing Technologies (IWPT3)*, Tilburg, The Netherlands, August 1993.

[WS83] Ralph M. Weischedel and Norman K. Sondheimer. Meta-rules as a basis for processing ill-formed input. *American Journal of Computational Linguistics*, 9(3-4):161–177, July-December 1983.