Providing Intelligent Language Feedback for Augmentative Communication Users

Christopher A. Pennington (penningt@asel.udel.edu)
Applied Science and Engineering Laboratories
Department of Computer and Information Sciences
University of Delaware / A. I. duPont Institute
Wilmington, DE 19899

People with severe speech and motor impairments (SSMI) sometimes 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 would 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 IPG (Intelligent Parser/Generator), a prototype system that uses the compansion technique.

keywords: augmentative communication, natural language processing, compansion, literacy

Providing Intelligent Language Feedback for Augmentative Communication Users

People with severe speech and motor impairments (SSMI) sometimes 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 would 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 IPG (Intelligent Parser/Generator), a prototype system that uses the compansion technique.

1 Introduction

People with severe speech and motor impairments (SSMI) sometimes 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 (on the order of 2-10 words per minute) [Fou80]. Consequently, augmentative communication users can often develop telegraphic patterns of language usage, especially if the disability occurs at an early age.

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 non-speaking 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 [DemaMcCoy92]. 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; however, it 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 [Sur93].

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 additional 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.[1] 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 Limitations of the Current Compansion Approach

The core of the compansion approach is a semantic parser that interprets input based on the concept of case frames [Fil68, Fil77]. In short, the semantic parser designates the primary verb as the main case or role of the expression: all other words in the input are used to fill semantic roles with respect

1. Of course, the context in which the expression occurs is extremely important; however, in many cases it is not possible for a computational system to have access to both sides of the entire conversation. Thus, unless noted otherwise, utterances are considered in isolation.

to the main verb that is chosen. For example, given the input *{John go store}*, *go* would be selected as the main VERB, *John* would fill the role of AGEXP (AGent-EXPeriencer), and *store* would be assigned a LOC (LOCation) role.[2]

3.1 Improving the Scoring Heuristics

The semantic parser relies on a set of scoring heuristics 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. 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.

This rating methodology has proved useful for developing a research prototype of the compansion technique, allowing distinctions to be made about some important relationships. However, it must be improved upon if it is to be used in provide appropriate corrective feedback for augmentative communication users in the process of developing literacy skills.

First, 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. At the very least, 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 [UEG+93] and appears to be quite in line with our purposes.

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.[3] This does not detract, however, from the goal to present the best correction first whenever possible.

---

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

3.2 Improving the Inferencing Strategies

The semantic parser contains some rudimentary inferencing principles based on general observations of telegraphic forms of expression found in some sign languages and pidgins. For example, if the semantic parser cannot find a main verb, 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 the parser cannot find a valid agent, it will infer the pronouns *I* or *you*, depending on whether the input is a statement or a question. While these are useful “shortcuts,” more research has been started that investigates the general language patterns of augmentative communication users [MMP+94].

It is expected that these results will contribute to improving the inferencing strategies of the semantic parser. Knowing more about the telegraphic expressions used in augmentative communication should enable the parser to make better interpretations of the user’s intended communication. One possible 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, 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 [Chi86]. 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.

4 An Adaptive User Language Model

The following architecture is proposed for constructing an adaptive user language model: the design of a general language assessment model, an overlay model reflecting the knowledge stored in the error taxonomy, and appropriate history mechanisms to update the model. The 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.

4.1 SLALOM - A Language Assessment Model

A primary goal of this work is to develop a consistent and theoretically sound model that can be used to evaluate each user’s English language proficiency. This profile will 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). An accurate profile of the user is also necessary to ensure that the system’s corrective responses are both relevant and understandable.

There is considerable linguistic evidence (especially in first and second language acquisition research) 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]). These findings have become the foundation for the development of a language assessment 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) according to their relationships and complexity. Then features of similar complexity are grouped into layers representing stereotypical "levels" of language ability. Below is a conceptual diagram of what the assessment model will look like:

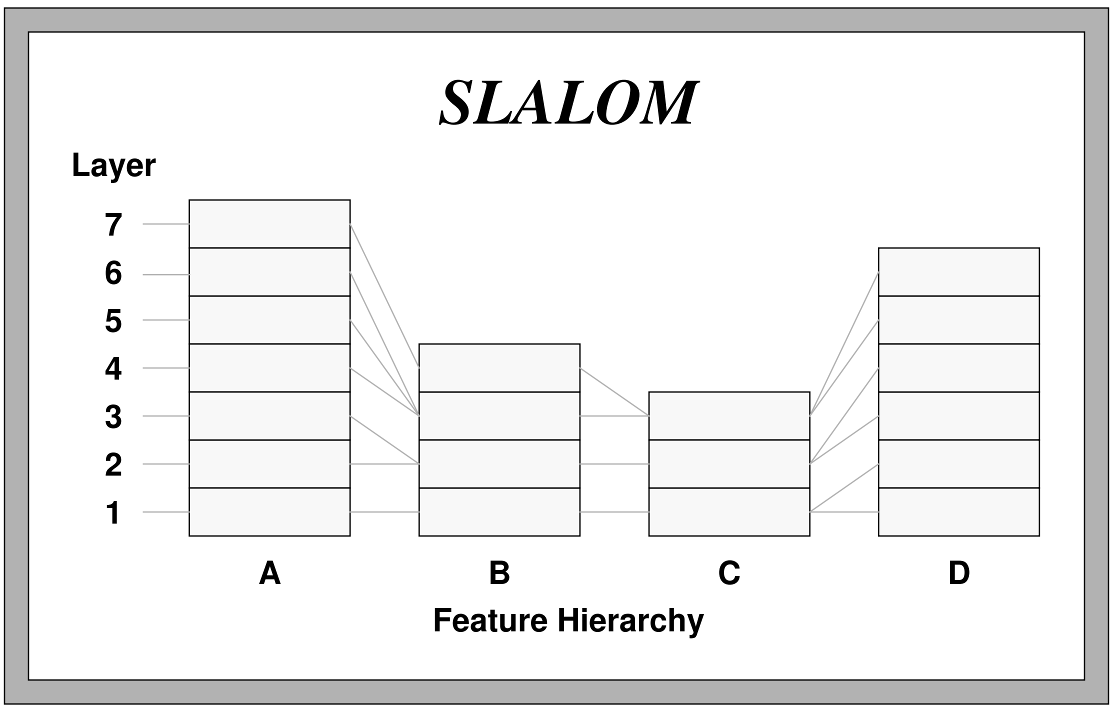

Each feature hierarchy is ordered according to complexity. So, assume A represents the morphology hierarchy: plurals are generally acquired before irregular past tense forms, so plurals might be at Layer 3 while irregular past tense would be at Layer 5. If adjective noun clauses appear at about the same time as irregular past tense forms then they might be positioned at Layer 5 in Feature Hierarchy D. Since separate hierarchies may contain different numbers of features, it is likely that some of the layers will be collapsed for certain hierarchies. Most of the evidence for the cross-hierarchical groupings are based on statistics and educational "grade" expectations. Possibilities for defining the default levels can be found in [Lee74] and [Cry82]. It is expected that a combination of existing assessment tools will be needed to ensure adequate coverage of English language features.

4.2 Using the Language Model

This adaptive language model provides consistent and reliable knowledge for user assessment, error identification, and correction selection. It also has great potential for use in tailoring explanatory responses to each user's language abilities and preferences. Although the defaults are based on linguistic and educational evidence, there is enough flexibility to allow for individual variation and change as a whole or within each feature hierarchy.

It is expected that these settings will be modified by an error taxonomy similar to that mentioned in Section 3.2, perhaps as an overlay on the default language model. By keeping the error taxonomy separate, it gives us the ability to associate additional information unique to the population of users we are assisting. This knowledge, in turn, could then be used by other mechanisms in the corrective process (e.g., a tutoring module).

Layers 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".[4] Also, information about typical telegraphic language patterns will be available for any reasoning processes.

4.3 Adaptation Mechanisms

A good history mechanism will 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.[5]

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. 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[6] since not all language difficulties are evident through error identification.

---

4. At this time it is not clear if the best strategy would be to assign the default as the minimum layer, the highest layer, or an average layer
5. This becomes more relevant if a tutoring component is being used to provide corrective responses.
6. That we expect to see based on the perceived language level of the user

5 Discussion of Enhancements to IPG

IPG (Intelligent Parser/Generator) is an augmentative communication prototype that is based on an Augmented Transition Network (ATN) formalism. Many aspects of the compansion technique have been encoded as a grammar for IPG. Below is a discussion of the changes needed to implement an adaptive user language model in IPG.

5.1 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 would be mal-rules [Sle82, WS83] to simulate the language patterns. They would handle expected telegraphic conventions (e. g., omitting forms of *be*) as well as any commonly observed 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. Advantages of this method include modularity that will allow association of additional information with the mal-rules in a group-specific manner; for instance, the assignment of semantic case frame information. This provides a way of interleaving semantic reasoning with the syntactic grammar. 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.

5.2 Implementing the User Language Model

In essence, this combination of mal-rules with the standard grammar comprises a default structure for a prototypical augmentative communication user. To instantiate this structure, a weighted grammar is proposed. Our expected approach is to implement a probabilistic context-free grammar similar to those described by Charniak [Cha93] and Allen [All95]. 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.

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. 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 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. Specific criteria for deciding when to change the feature acquisition levels (e.g., from "acquiring" to "known") have not yet been determined. 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.

5.3 Processing Considerations

After a sentence is parsed there will be an indication that errors exist and are 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 dealing with 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.

6 Future work

The most immediate future need is to further specify the relationships of features within the base language model 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 have been made to the IPG system. As discussed previously, it is thought that changes to IPG 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 on IPG that will adjust weights and other features based on experiences with the augmentative communication user's language choices and feedback selections.

Eventually, this work will be integrated back into a larger project called ICICLE (Intelligent 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 planned for ICICLE concerns adaptive tutoring and explanation. This module will rely strongly on 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., anaphor resolution). This is a major area of needed research.

7 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 system that uses the compansion technique.

8 Acknowledgments

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). Additional support has been provided by the Nemours Research Program.

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.

[Chi86] David N. Chin. KNOME: Modeling what the user knows in UC. In *Proceedings of User Modeling Workshop*, pages 1–36, 1986.

[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.

[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.

[Fou80] R. A. Foulds. Communication rates for non-speech expression as a function of manual tasks and linguistic constraints. In *Proceedings of the International Conference on Rehabilitation Engineering*, pages 83–87, Toronto, Canada, 1980. RESNA.

[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*, volume 49 of *Current Issues in Linguistic Theory*, 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.

[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 Compansion. In Mary Binion, editor, *Proceedings of the RESNA '94 Annual Conference*, pages 129–131, Arlington, VA, 1994. RESNA Press.

[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. A methodology for developing an error taxonomy for a computer assisted language learning tool for second language learners. Technical report 93-16, Department of Computer and Information Sciences, University of Delaware, Newark, DE, January 1993.

[Sur93] Linda Z. Suri. *Extending Focusing Frameworks to Process Complex Sentences and to Correct the Written English of Proficient Signers of American Sign Language*. Technical report 94-21, Department of Computer and Information Sciences, University of Delaware, Newark, DE, 1993.

[UEG+93] 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, 1983.