Interface and Language Issues in Intelligent Systems for People with Disabilities

Kathleen F. McCoy

Computer and Information Sciences Department and
Applied Science and Engineering Laboratories,
University of Delaware/duPont Hospital for Children,
Newark, DE 19716
mccoy@cis.udel.edu

1 Introduction

The papers in this section describe a diverse set of applications of various Artificial Intelligence (AI) techniques. The overriding theme of the papers is that of interfaces to language/communication for people who have disabilities which make it difficult for them to communicate using spoken language, or interfaces that use spoken language or some other means (e.g., eye-tracking) as one kind of input to controlling an environment for people whose physical disability precludes them from physically manipulating their environment.

Several of the papers can be seen as falling into the area of Augmentative and Alternative Communication (AAC), and many use some processing methodologies from the AI area of Natural Language Processing (NLP). Some represent "mature" technologies that have been tested on actual users, while others involve the development of technologies which hold future promise.

In this paper I will attempt to give an overview of the area to which many of these papers can be fit – pointing out places on which the papers in this volume can be seen as focusing and where application of AI technologies might continue. Next an overview of NLP will be provided (again pointing out which aspects the papers in this volume have emphasized). Finally, other AI areas emphasized in these papers will be discussed.

2 Augmentative and Alternative Communication (AAC)

AAC is the field of study concerned with providing devices or techniques to augment the communicative ability of a person whose disability makes it difficult to speak in an understandable fashion.

A variety of AAC devices and techniques exist today. Many of these are aimed at people who have severe speech impairments (such that their speech cannot

⋆ This work has been supported by a Rehabilitation Engineering Research Center grant from the National Institute on Disability and Rehabilitation Research (#H133E30010). Additional support was provided by the Nemours Foundation.

be reliably understood) and whose muscular control makes typing on a standard keyboard difficult (if not impossible). Some devices designed for such populations are non-electronic word boards containing words and phrases in standard orthography and/or iconic representations. A user of such a non-electronic system points to locations on the board and depends on the listener to appropriately interpret the selection. Electronic word boards may use the same sorts of selectable items, but may also include speech synthesis. These presumably provide more independence for the user who does not need to rely on a partner to interpret the selections. However, these systems may place more burden on the user who must be aware of the actual strings associated with each selection and must ensure that the synthesized string be an appropriate English sentence. Since the system will only "speak" what has been selected, generally more selections are required per sentence and speed of selection becomes more crucial.

3 Computer-Based Augmentative and Alternative Communication

A traditional computer-based AAC system can be viewed as providing the user with a "virtual keyboard" that enables the user to select items to be output to a speech synthesizer or other application. A virtual keyboard can be thought of as consisting of three components: (1) a physical interface providing the method for activating the keyboard (and thus selecting its elements), (2) a language set containing the elements that may be selected, and (3) a processing method that creates some output depending on the selected items. All three of these elements must be tailored to an individual depending on his/her physical and cognitive circumstances and the task they are intending to perform.

For example, for people with severe physical limitations, access to the device might be limited to a single switch. A physical interface that might be appropriate in this case involves row-column scanning of the language set that is arranged (perhaps in a hierarchical fashion) as a matrix on the display. The user would make selections by appropriately hitting the switch when a visual cursor crosses the desired items. In row-column scanning the cursor first highlights each row moving down the screen at a rate appropriate for the user. When the cursor comes to the row containing the desired item, the user hits the switch causing the cursor to advance across the selected row, highlighting each item in turn. The user hits the switch again when the highlighting reaches the desired item in order to select it. For users with less severe physical disabilities, a physical interface using a keyboard may be appropriate. The size of the keys on the board and their activation method may need to be tailored to the abilities of the particular user.

One of the papers in this volume, [Gip98], involves an intelligent eye-tracking system which can be viewed as a physical interface to an AAC system. The system allows a user to control a computer through five electrodes placed on the head. Users can be taught to control their muscles and to use head movement to control a cursor on the screen. The use of this EagleEyes system with appropriate applications (e.g., language sets) has enabled several users to communicate and exhibit intelligence that was previously locked out because their disabilities precluded their use of other traditional interfaces.

Great challenges in this work include (1) use of appropriate sensors, (2) developing methods for determining when eye gaze is being used purposefully (i.e., dealing with the "midas touch"), (3) accuracy and control, (4) developing augmentations such as mouse clicks via eye-blinks. Various training methods for using the interface are discussed and various applications developed and tailored to individuals.

The physical interface is also somewhat of an issue in [SWP98] also in this volume. This paper focuses on a modern AAC device which uses vision techniques to recognize sign language. The eventual application is that of translating the signed material into spoken text allowing the person who is signing to be understood by people who do not know sign language.

The interface issue in this paper is that the authors envision the recognition system to be a "wearable computer" which is worn by the signer. The camera for the system described in the paper is worn on a cap and has a view of the signer's hands (which are tracked by the vision system). While the authors note that the eventual system will need to capture facial expression, the cap mounted system has shown greater accuracy in picking out and following the hands then their previous attempts. This has led to greater overall system accuracy.

While not an AAC system, [KCB$^+$98] is an excellent example of a system whose interface combines several modalities (e.g., spoken and gestural) in allowing the user to control a robot to manipulate the environment.

Independent of the physical interface in an AAC system is the language set that must also be tuned to the individual. For instance, the language set might contain letters, words, phrases, icons, pictures, etc. If, for example, pictures are selected, the processing method might translate a sequence of picture selections into a word or phrase that will be output as the result of the series of activations. Alternatively, consider a language set consisting of letters. A processing method called abbreviation expansion could take a sequence of key presses (e.g., `chpt`) and expand that set into a word (e.g., `chapter`).

The use of a computer-based AAC device generally has many trade-offs. Assuming a physical interface of row-column scanning, a language set consisting of letters would give the user the most flexibility, but would cause standard message construction to be very time consuming. On the other hand, a language set consisting of words or phrases might be more desirable from the standpoint of speed, but then the size of the language set would be much larger causing the user to take longer (on average) to access an individual member. In addition, if words or phrases are used, typically the words would have to be arranged in some hierarchical fashion, and thus there would be a cognitive/physical/visual load involved in remembering and accessing the individual words and phrases.

One kind of language set that has been found to be very effective is an iconic language set. An iconic language set must be coupled with a processing method to translate the icon sequence selected into its corresponding word/phrase. A challenge in developing an iconic language set is to develop a language that can be easily used. In particular, the user must be able to recall the sequences of icons that produce the desired output. [ACP98] is a paper in this volume concerned with a design methodology for developing iconic languages. In the methodology icons in the language are associated with a set of semantic features which capture the various semantic concepts inherent in an icon. A set of relations is described which allow the meanings of individual icons to be combined in various fashions. The kinds of combinations available and the resulting semantic inferences can be used to establish meaningful sequences of icons and to predict the resulting intuitive meaning.

[VP98] (this volume) is concerned with several aspects of an AAC system that must be tuned if the language set consists of phrases rather than individual lexical items. In particular, when phrases/sentences are used the number of items to be accessed is quite large and the time spent navigating to the phrase must be minimal. This is because if the phrase takes longer to access than it would have taken to compose it from scratch, there is no savings!

The key idea in [VP98] is to store the text needed for a typical event (e.g., going to a restaurant) with the typical sub-events for which the text might be needed. For example, if a typical restaurant script has an entering, ordering, eating, and leaving scene, the text needed for each of those scenes would be stored with the scene. Thus the user could access the appropriate text by following along the script. Such a system puts certain requirements on the system interface, and some of these are explored in the paper (as well as a preliminary evaluation of the use of schemata to store prestored text in a communication aid).

The paper [PM98] is focused on the processing aspect of an AAC system. One issue that must be faced concerns literacy skills for people who use AAC systems. Because of the enormous time required to communicate with an AAC device, many users develop strategies for getting across their functional communication using telegraphic utterances. While this is a very beneficial strategy, it may cause non-standard English to be reinforced. The idea in this paper is to use processing on the telegraphic selections given by the user in order to give correct English sentence feedback. Such a system may have the benefit of raising the literacy skills of the user. The expansion of telegraphic input into full English sentences has been discussed in previous papers by this group. The focus of [PM98] is on additions (such as a user model which captures levels of literacy acquisition) which would be necessary for this new application.

While not a traditional AAC system, [TO98] focuses on the processing required to translate Japanese into Japanese Sign Language (JSL) so as to make Japanese communication accessible to a person who is deaf and cannot understand Japanese. The basic methodology in translating between the two languages involves a translation of lexical items (word-for-word translation) and a translation of syntactic structures between the two languages. One problem that is of concern is that there may not be a lexical item in JSL corresponding to a particular lexical item in Japanese. The paper describes a method for finding a similar word based on some meaning information contained in the dictionaries for Japanese and for Japanese Signs contained in the system. They have evaluated their system on some news stories and are reaching translation accuracy rates of 70%.

Another paper whose aim is similar to [TO98] is [Cha98] which is concerned with knowledge bases necessary for a vision system to translate American Sign Language into English. In particular, the focus of the paper is at the word level and it is concerned with capturing information which would allow the signs to be translated into their corresponding English word equivalents. This is done using a feature-based lexicon which captures linguistically motivated features of the signs (which may be recognized by the vision system). This allows the vision system to search for the word corresponding to a sign in an efficient manner.

4 The Application of NLP

The fields of Natural Language Processing (NLP) and Computational Linguistics attempt to capture regularities in natural (i.e., human) languages in an effort to enable a machine to communicate effectively with a human conversational partner [All87], [All95], [GM89], [Gri86]. Areas of research within NLP have concentrated on all levels of processing – from the sub-word level (e.g., phonology, morphology) all the way up to the discourse level.

[Cha98], in this volume, takes advantage of linguistic work concerning individual signs and their components in American Sign Language (ASL). The goal is to develop a sign language lexicon which can be used to recognize ASL signs (and to translate them into their English equivalents). The lexicon would be used by a vision system, and it indexes the signs by both manual and non-manual information. The manual information includes information about the movement, location (with respect to the signer's body), handshape, and hand orientation used in making the sign. Non-manual information includes facial characteristics (such as raised eyebrows) or body orientation during the sign. These choices were motivated by sign formation constraints.

Above the word level, three major areas of research in NLP and Computational Linguistics (syntax, semantics, and pragmatics) deal with regularities of language at different levels. Various techniques have been developed within each which will be useful for application to various AAC technologies.

4.1 Syntax

The syntax of a language captures how the words can be put together in order to form sentences that "look correct in the language" [All87]. Syntax is intended to capture structural constraints imposed by language which are independent of meaning. For example, it is the syntax of the language that makes:

"I just spurred a couple of gurpy fliffs."

seem like a reasonable sentence even if some words in the sentence are unknown, but makes

"Spurred fliff I couple a gurpy."

seem ill-formed.

Processing the syntax of a language generally involves two components: 1) a grammar which is a set of rules that refer to word categories (e.g., noun, verb) and various morphological endings (e.g., +S for plural, +ING) that capture the allowable syntactic strings in a language and; 2) a parser which is a program that, given a grammar and a string of words, determines whether the string of words adheres to the grammar. (See [All87], [All95], [GM89], and [Win83] for examples of various parsing formalisms and grammars.)

Using a grammar and parser an AAC system would be able to: 1) determine whether or not the utterance selected by the user was well-formed syntactically, 2) determine valid sequences of word categories that could form a well-formed sentence, 3) given a partial sentence typed by the user, determine what categories of words could follow as valid sentence completions, 4) determine appropriate morphological endings on words (e.g., that a verb following the helping-verb "have" must be in its past participle form), and 4) determine appropriate placement of function words which must be added for syntactic reasons (e.g., that certain nouns must be preceded by an article, that the actor in a passive sentence is preceded by the word "by").

Syntactic knowledge is currently being successfully applied in a number of AAC projects. For example, several word prediction systems use syntactic information to limit the words predicted to those which could follow the words given so far in a syntactically valid sentence [SAN87], [VMD92], [Van91]. To some extent, many grammar checkers available today and systems aimed toward language tutoring (e.g., [SM93a], [SM93b], [MNBR92], [WBB+92]) also use syntactic information, though there is still great room for improvement.

In this volume syntactic processing of spoken language is used in [KCB+98] in order to understand the user's intentions. In that system, a side-effect of parsing is the computation of meaning. Following a grammar of sign language sentences (in this case, there is only one sentence pattern) is used in [SWP98] in order to aid the Hidden-Markov-Model to recognize the signs. Finally, [TO98] use syntactic translation rules as one step in translating Japanese sentences into Japanese Sign Language sentences.

4.2 Semantics

The area of semantics deals with the regularity of language which comes from the meanings of individual words and how the individual words in a sentence form a meaningful whole. A problem in semantics is the fact that many words in English have several meanings (e.g., "bank" may refer to the edge of a river or to a financial institution). In Computational Linguistics the use of selectional restrictions [KF63], case frames [Fil68], [Fil77], and preference semantics [Wil75] is based on the idea that the meanings of the words in a sentence are mutually constraining and predictive [SR82]. When the words of a sentence are taken as a whole, the meanings of the individual words can become clear.

Consider the sentence "John put money in the bank." Here the financial institution meaning of "bank" can be inferred from the meaning of the verb "put" (which expects a thing to be put and a location to put it in) and the fact that "money" is the appropriate kind of object to be put in a financial institution.

Note that in order to take advantage of semantics, a natural language processing system must (1) have rules (selectional restrictions, case frames) which capture the expectations from individual words (e.g., "eat" is a verb that generally requires an animate agent and an object which can be classified as a food-item), and (2) have a knowledge base that contains concepts that are classified according to their meanings (e.g., "apples" are food-items, "John" is a person, and "people" are animate).

The Compansion system [DM92], [MDJ$^+$94], which is the underlying system referred to in [PM98], has made extensive use of semantic information to transform telegraphic input into full sentences. In this volume it is suggested that the full sentence constructed might be used as a literacy aid. Semantic information is also a main component of the PROSE [WBN92] system developed at the University of Dundee. PROSE is intended to give the user access to prestored phrases/sentences/stories which can be accessed according to their semantic content. The basic idea is that sets of phrases, stories, sentences etc. will be input by the user (in advance) along with some semantic information about their content. PROSE will then store this information in an intelligent way according to the semantic information given. The system will then retrieve the pre-stored material, based on minimal prompting by the user, in semantically appropriate contexts.

Both syntax and semantic information are used in the project described in [Cop96], [CFS97] involving "co-generation" (where the generation of a natural language sentence is shared between the user of the system and the system itself). This project attempts to speed communication rate by allowing sentences to be generated with fewer selections on the part of the user. Here the user fills in a "semantic template" with desired content words. The system then generates a full grammatical sentence based on the semantic information specified by the user.

In this volume two papers make extensive use of semantic information for diverse purposes. [ACP98] uses semantic information associated with icons and combination methods to determine "natural meanings" in sequences of icons. It is suggested that the combination rules can be useful in developing intuitive iconic languages.

In describing a machine translation system between Japanese and Japanese Sign Language (JSL) [TO98] uses semantic information in order to find an appropriate translation for a Japanese word when there is no corresponding word in JSL. In order to do this, they look for similar Japanese words (i.e., those with the same concept identifier) in an extensive Japanese dictionary and attempt to find a word that does have a translation in JSL. Failing this, they attempt using words in the dictionary definition of the word with no JSL equivalent. Finally, if this fails as well, they attempt to use the superconcept. If all of these methods fail, the system resorts to using finger-spelling. The finding of an appropriate substitution word is possible because of the semantic information encoded in their Japanese dictionary.

4.3 Pragmatics

Pragmatic information refers to the broad context in which language and communication takes place [All87], [JWS81], [Lev83]. Situational context and previous exchanges produce conversational expectations about what is to come next. Natural language processing has concerned itself with developing computational mechanisms for capturing these same expectations in a computer.

A great deal of AAC work that takes advantage of pragmatic information has come from the University of Dundee. Their CHAT system [ANA87] is a communication system that models typical conversational patterns. For example, a conversation generally has an opening consisting of some standardized greeting, a middle, and a standardized closing. The system gives users access to standard openings and closings (at appropriate times). In addition (for the middle portion of a conversation) it provides a number of "checking" or "fill" phrases (e.g., "OK", "yes") which are relatively content free but allow the user to participate more fully in the conversation.

The TALKSBACK system [WAN90], [WBNA91], [WBN92] incorporates user modeling issues. It takes as input some parameters of the situation (e.g., the conversational partners, topics, social situation) and predicts (pre-stored) utterances the user is likely to want based on the input parameters. For example, if the user indicates a desire to ask a question about school to a particular classmate, the system might suggest a question such as "What did you think of the geography lesson yesterday?". In other words, the system attempts to use the parameters input by the user to select utterances that are pragmatically appropriate.

Pragmatic information is the key in [VP98] in this volume. In their system prestored text is stored in schema structures [Sch82], [SA77] which capture typical sequences of events. This should allow access to text appropriate for an event by allowing the system user to "follow along" the typical sequence.

Pragmatic information in the form of a user model is also a focus of [PM98]. Here the user model attempts to capture the level of literacy acquisition in an attempt to provide beneficial feedback to the user.

5 Other Artificial Intelligence Technology

Another AI technology prominent in this section is vision processing. [SWP98] uses vision technology in order to identify and track the hands on a video, and interprets sign language using Hidden Markov Models (HMM's). Evaluation of the results includes an experiment where gloves are worn (and the hand is tracked by color) and an experiment where the hands are tracked on skin tone. In both experiments the HMM is trained on 400 sentences and 100 sentences are used for testing. The training and testing sets use a limited number of signs, and the recognition rate is above 97% when the sentence structures are constrained to follow a predefined sentence pattern.

The paper [KCB+98] focuses on the integration of several different AI technologies in order to provide an interface that enables a person who has physical disabilities to manipulate an unstructured environment. The project combines a vision subsystem (which is able to identify object location, shape, and pose), an interface subsystem which interprets limited spoken commands combined with pointing gestures from a head-mounted pointing device, and a planning subsystem that interprets commands by the user and plans a method for carrying out the request. The user of the system provides some information to the system (such as indicating the class of particular objects) and then can ask the system to move objects around to various locations.

6 Conclusion

The papers in this volume provide us with a snapshot of the variety of issues that must be considered when applying AI technologies to projects involving people with disabilities. Here a focus is on controlling interfaces and language issues. The AI technologies used are quite varied; the resulting ideas have a great deal of promise and point us to future possibilities.

References

[ACP98] P. L. Albacete, S. K. Chang, and G. Polese. Iconic language design for people with significant speech and multiple impairments. In *This Volume*. 1998.

[All87] James Allen. *Natural Language Understanding*. Benjamin/Cummings, CA, 1987.

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

[ANA87] Norman Alm, Alan Newell, and John Arnott. A communication aid which models conversational patterns. In Richard Steele and William Gerrey, editors, *Proceedings of the Tenth Annual Conference on Rehabilitation Technology*, pages 127–129, Washington, DC, June 1987. RESNA.

[CFS97] Ann Copestake, Dan Flickinger, and Ivan Sag. Augmented and alternative NLP techniques for augmentative and alternative communication. In *Proceedings of the ACL Workshop on Natural Language Processing for Communication Aids*, Madrid, Spain, 1997.

[Cha98] Robbin Nicole Chapman. Lexicon for ASL computer translation. In *This Volume*. 1998.

[Cop96] Ann Copestake. Applying natural language processing techniques to speech prostheses. In *Working Notes of the 1996 AAAI Fall Symposium on Developing Assistive Technology for People with Disabilities*, Cambridge, MA, 1996.

[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.

[Gip98] James Gips. On building intelligence into EagleEyes. In *This Volume*. 1998.

[GM89] G. Gazdar and C. Mellish. *Natural language processing in lisp, an introduction to computational linguistics*. Addison-Wesley Publishing Company, New York, 1989.

[Gri86] R. Grishman. *Computational linguistics: An introduction*. Cambridge University Press, Cambridge, 1986.

[JWS81] A. Joshi, B. Webber, and I. Sag, editors. *Elements of discourse understanding*. Cambridge University Press, Cambridge, 1981.

[KCB+98] Zunaid Kazi, Shoupu Chen, Matthew Beitler, Daniel Chester, and Richard Foulds. Speech and gesture mediated intelligent teloperation. In *This Volume*. 1998.

[KF63] J. J. Katz and J. A. Fodor. The structure of semantic theory. *Language*, 39, 1963.

[Lev83] S. Levinson. *Pragmatics*. Cambridge University Press, Cambridge, 1983.

[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 (ASSETS94)*, pages 107–114, Marina del Ray, CA:, 1994.

[MNBR92] C. Morris, A. Newell, L. Booth, and I. Ricketts. Syntax PAL: A writing aid for language-impaired users. In *Presented at ISAAC-92*, volume 8, Philadelphia, PA, 1992. ISAAC, Abstract in AAC – Augmentative and Alternative Communication.

[PM98] Christopher A. Pennington and Kathleen F. McCoy. Providing intelligent language feedback for augmentative communication users. In *This Volume*. 1998.

[SA77] R. C. Schank and R. P. Abelson. *Scripts, plans, goals and understanding: An inquiry into human knowledge structures*. Erlbaum, Hillsdale, NJ, 1977.

[SAN87] A. L. Swiffin, J. L. Arnott, and A. F. Newell. The use of syntax in a predictive communication aid for the physically impaired. In Richard Steele and William Gerrey, editors, *Proceedings of the Tenth Annual Conference on Rehabilitation Technology*, pages 124–126, Washington, DC, June 1987. RESNA.

[Sch82] Roger C. Schank. *Dynamic Memory: A theory of reminding and learning in computers and people*. Cambridge University Press, NY, 1982.

[SM93a] Linda Z. Suri and Kathleen F. McCoy. Correcting discourse-level errors in CALL systems. In *Proceedings of The World Conference on Artificial Intelligence in Education (AIED-93) (refereed poster presentation)*, August 1993.

[SM93b] 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 TR-93-16, Dept. of CIS, University of Delaware, 1993.

[SR82] Steve Small and Chuck Rieger. Parsing and comprehending with word experts (a theory and its realization). In Wendy G. Lehnert and Martin H. Ringle, editors, *Strategies for Natural Language Processing*, pages 89–147. Lawrence Erlbaum Associates, 1982.

[SWP98] Thad Starner, Joshua Weaver, and Alex Pentland. A wearable computer based American Sign Language recognizer. In *This Volume*. 1998.

[TO98] Masaaki Tokuda and Manabu Okumura. Towards automatic translation from Japanese into Japanese Sign Language. In *This Volume*. 1998.

[Van91] Julie A. VanDyke. Word prediction for disabled users: Applying natural language processing to enhance communication. Thesis for honors bachelor of arts in cognitive studies, University of Delaware, Newark, DE, 1991.

[VMD92] Julie VanDyke, Kathleen McCoy, and Patrick Demasco. Using syntactic knowledge for word prediction. Presented at ISAAC-92. Abstract appears in Augmentative and Alternative Communication, 8., 1992.

[VP98] Peter B. Vanderheyden and Christopher A. Pennington. An augmentative communication interface based on conversational schemata. In *This Volume*. 1998.

[WAN90] A. Waller, N. Alm, and A. Newell. Aided communication using semantically linked text modules. In J. J. Presperin, editor, *Proceedings of the 13th Annual RESNA Conference*, pages 177–178, Washington, D.C., June 1990. RESNA.

[WBB+92] Andrew Wright, William Beattie, Lynda Booth, William Ricketts, and John Arnott. An integrated predictive wordprocessing and spelling correction system. In Jessica Presperin, editor, *Proceedings of the RESNA International '92 Conference*, pages 369–370, Washington, DC, June 1992. RESNA.

[WBN92] A. Waller, L. Broumley, and A. Newell. Incorporating conversational narratives in an AAC device. Presented at ISAAC-92. Abstract appears in Augmentative and Alternative Communication, 8, 1992.

[WBNA91] A. Waller, L. Broumley, A. F. Newell, and N. Alm. Predictive retrieval of conversational narratives in an augmentative communication system. In *Proceedings of 14th Annual Conference*, Kansas City, MI, June 1991. RESNA.

[Wil75] Y. Wilks. An intelligent analyzer and understander of English. *Communications of the ACM*, 18(5):264–274, 1975.

[Win83] Terry Winograd. *Language as a cognitive Process, Volume 1: Syntax.* Addison-Wesley Publishing Company, Reading, Ma, 1983.

This article was processed using the LaTeX macro package with LLNCS style