Intelligent AAC Systems: What Can Be Done Now

Kathleen F. McCoy, ASEL, University of Delaware/A.I. DuPont Institute
Eric H. Nyberg 3rd, Carnegie Mellon University/Semantic Compaction Systems
Bruce R. Baker, Semantic Compaction Systems

Introduction

With the advent of smaller, more powerful computers, it is becoming possible to develop AAC systems which take advantage of Artificial Intelligence and Computational Linguistics techniques. In particular, we envision the development of an "intelligent" AAC system that can understand the operator's input well enough to infer what words or syntactic structure may be omitted or missing, enabling the device to output a complete and correct sentence from partial input. At some point in the future it may be possible to create an AAC system with all of the linguistic knowledge of a human speaker, but the limitations of existing theory and technology allow only a small subset of that knowledge to be incorporated in current systems. In this paper we explore some of the possibilities for AAC systems that could be developed given the present state of affairs. Of course, each of the possibil ities must contain trade-offs of one kind or another. We discuss various populations of potential AAC sys tem operators and show how different populations would benefit from different aspects of "intelligence" in their system.

Potential Operators

Hundreds of thousands of people could benefit from Artificial Intelligence in electronic communication aids. This population is heterogeneous and dispersed. Nevertheless, certain groups do cleave together. A discussion of two of these groups can illuminate the considerations involved.

One group includes ambulatory, cognitively impaired adolescents with severe language disability. Consider a 14-year old boy with Down's Syndrome, who has a language age of 4 years and 8 months. His actions are socially appropriate, he has frequent communicative intent, but does not speak and has failed to master signing. His vocabulary needs for school and daily living activities might include 50-75 function words and at least 300 content words. Such an operator would be willing to allow a communication device to determine the syntactic organization of his utterances. The finer points of style and pragmatics are less im portant to such a person. He may even acquire some syntax and independent speech through observation of his device's output.

Another group includes cognitively intact, educated adults who have recently acquired a severe speech im pairment. Consider the situation of a 44-year old fe male lawyer. She has suffered a brain stem stroke and is cognitively normal, yet physical disabilities preclude all but row-column scanning access to a communication aid. We might think that a template system giving access to very rapid yet general language would meet this individual's needs. Her perceptions would probably be different. Cognitively intact individuals who have few opportunities to communicate want complete control of their language in each communication opportunity. Style and pragmatics are of paramount importance to users in this group, who would be unwilling to have the communication device provide the overall organization of their utterances.

System Knowledge Sources

We envision systems in which operators will choose partial representations of their final communicative intentions using some input paradigm. It will be the job of the system to "understand" the intended meaning and eventually generate a well-formed utterance based on this meaning (Baker, 1982), (Demasco et al., 1989). Since the system's knowledge of the discourse context will be incomplete, however, there is a possibility that there may be more than one interpretation of the input.

Consider the following input: "ate duck". Notice that this is certainly not a well-formed utterance. However, one may be able to infer the intended meaning by taking into account various sources of knowledge.

Semantic Knowledge -- information about the semantic categories of the words involved would be helpful in determining what role was being played by each word. For instance, knowing that "ate" is an action performed by animate objects might lead one to conjecture that the intended utterance was "The duck ate".

Pragmatic Knowledge -- knowledge about the discourse situation might aid in inferring the intended meaning. For instance, if the preceding discourse has been discussing John and what he did last night, a more reasonable rendition of the input might be: "John ate duck". The assumption on the user's part was that the agent of the eating could be inferred from context.

Notice that as more sources of knowledge are applied, the system is able to do more inferencing toward the intended meaning. This, in turn, allows the user to provide less information about his/her intended utterance. At the same time, as the system does more inferencing the user is given less control over the actual utterance.

The above example shows a rather simple sentence in terms of its syntactic structure. It is a simple declarative sentence with simple noun phrases. Such a sentence may require little inferencing on the part of the system because it contains only one main verb and its simple noun phrases are given in canonical order -- usually AGENT-ACTION-THEME. Ambiguity resolution can become more complicated if the input is intended to be realized by more complicated syntactic structures.

Consider the problems with allowing syntactic structures which do not maintain canonical word order (such as a passive sentence as in "The duck was eaten by John"). Notice to generate these sentences the system may not assume that the input given by the user is AGENT-ACTION-THEME but must infer the role being played by each noun phrase. In addition, the system must have some reason for choosing to use a passive sentence over a simple declarative sentence.

The user may be able to provide the system with some information to help avoid extensive inferencing. For instance, in the above case, the system and user might follow the convention that all input will be given in canonical order, but a special marker will be used to indicate a passive sentence is desired. Of course, following such a convention may be difficult or impossible for some user groups.

The addition of relative clauses (as in "the man *that knows my father* helped me") adds a great deal of complication because the input from the user may now contain two verbs -- one to be taken as the main verb of the sentence and the second to be part of the relative clause modifying one of the noun phrases of the main sentence.

As the list of possible syntactic structures grows, the job of the system becomes more complicated. With the added complication comes the potential for ambiguity which, without aid from the user, may not be resolvable by the system given the current state of affairs in computational linguistics.

Fitting the System to the User

f we are to develop an intelligent system given today's technology, decisions must be made concerning both permitted syntactic structures (the more complicated the more powerful the inferencing must be) and the permitted semantic sophistication (the more complicated/extensive the domain of discourse the more detailed the required knowledge must be) and the locus of control given to the user (some of the inferencing required by the system can be eliminated by requiring the user to be more explicit about intended meaning).

If we consider the first operator group discussed, it is clear that an AAC system must assume a great degree of control and perform a large amount of inference to produce a correct utterance from a few content words entered by the operator. Such a system may not be able to provide much syntactic sophistication, and its word list may not be that extensive, but it can certainly meet the needs of operators in this group by cutting down on the number of actuations required for communication and producing a correct sentence from partial input.

If we consider the second operator group, we see that an AAC system must address a different need, cutting down on physical actuations without taking too much control from the user, who is still free to determine the overall structure of each utterance. This user will demand syntactic sophistication from the system as well as a large word list. This system will require that the user follow conventions for marking various syntactic desires so that the system can generate the intended utterance and the user can be saved the physical hardship that would be required to fully specify the utterance. On the other hand, this user is quite capable of following the necessary conventions. For her, it is a small price to pay in order to maintain control and still get the desired sophistication within a relatively short amount of time.

Conclusion

While current technology in Artificial Intelligence, Linguistics, and Computer Science will not allow us to build a vastly intelligent AAC system containing sophisticated syntactic, semantic, and pragmatic capabilities, it is possible to provide systems with limited intelligence. In this paper we have shown that this intelligence must be tailored to the wants and needs of each particular group of potential users.

Acknowledgments

This work is partially supported by Grant #H133E80015 from the National Institute on Disability and Rehabilitation Research. Additional support has been provided by the Nemours Foundation.

References

B. Baker. Minspeak. Byte, 186ff, September, 1982.

Intelligent AAC Systems

P. Demasco, K. McCoy, Y. Gong, C. Pennington, and C. Rowe. Towards more intelligent AAC interfaces: the use of natural language processing. In *Proceedings of the 12th Annual Conference*, pages 141-142, RESNA, New Orleans, Louisiana, June 1989.

Kathleen McCoy
Dept. of Computer and Information Sciences
University of Delaware
Newark, DE. 19716
Email: mccoy@udel.edu