The Problem of Metaphor in Intelligent AAC Systems

Mark Jones, Kathleen McCoy & Patrick Demasco

Applied Science and Engineering Laboratories
University of Delaware/ A.I. duPont Institute
Wilmington, Delaware USA

Abstract

Intelligent AAC systems, which encode an understanding of language, cannot afford to ignore the problem of metaphor. Here we discuss the nature and promise of future AAC systems that handle some metaphorical statements. We describe the problems that metaphor creates for such systems, as well as a special class of metaphors called transparently-motivated metaphors, which are particularly important. Finally, we point to a possible solution to the problem.

Background

The majority of present AAC systems use numerous techniques at the surface level of an expression to facilitate the communication of the user. More intelligent AAC systems of the future will be able to take advantage of an understanding of what meaning the user intends to convey in order to facilitate the communication of the user in ways that present systems cannot.

These more intelligent AAC systems are distinguished by having a deeper understanding of the expression the user wishes to convey. They must have some understanding of what is actually meant by the input from the user and the language generated.

These systems will rely upon advances from the field of Natural Language Processing (NLP). Techniques from NLP can empower a computer system to understand abbreviated input, track previous expressions, recognize communicative goals, and disambiguate expressions with the help of specified world knowledge. The goal is to have a computer system that can help a user communicate with insight similar to that of a human aid.

<u>Example</u>

Consider how one could understand, and perhaps rewrite in sentence-form, the ideas behind the abbreviated utterance "Buy 3000 Exxon." Several elements come into play here. The context could tell us whether this is a command or a question. Our knowledge that (unlike buying a pizza) Exxon is a company that can only be purchased in part, via stock, is important. Knowing that stocks are bought with money is also important. Finally, knowledge about the user's wealth would indicate whether 3000 shares or dollars is intended (if indeed it is a command to execute a purchase). The equivalent sentence may be "Buy 3,000 dollars worth of Exxon stock."

<u>Complementing existing technology</u>

The advances represented by present surface-level systems need not be replaced, but can be combined with the benefits of more intelligent AAC systems; they can work in tandem. For example, the Sentence Compansion System (Demasco et al., 89, McCoy et al., 90) is designed to work in tandem with word selection and word prediction systems.

Statement of the Problem

The knowledge that more intelligent systems are based upon must be specific (e.g., stocks are purchased with money, hammers are tools, windows are fragile) to be useful to a computer system.Therefore, it may be necessary for a system to require that metaphorical expressions be excluded. Otherwise they will be unable to behave adequately, and would fail their task.

This may at first seem reasonable, if the user wants to write poetry, the user should know that the more intelligent system should be switched off. Unfortunately, there is a class of metaphors that are not easily recognized as metaphorical (Lakoff and Johnson 80). We cannot expect the user to recognize the metaphorical nature of such expressions. These types of expressions are called *transparently-motivated metaphors* (Jones and McCoy 1992).

Expressions such as "Put $3000 into Exxon stock." occur quite often and, at first glance, may not appear to be metaphorical. However, the speaker does not literally put money into the stock, but rather buys the stock with the money.

<u>Example</u>

Consider how the system of the first example would handle "Put 3000 Exxon." Here we see where the traditional approach to encoding the knowledge breaks down. The system would quickly become confused because there is only one Exxon. Furthermore, Exxon is hardly the type of object that can be put somewhere. The highly structured knowledge may be quite powerful, but it is also very inflexible to metaphors, even those which we use without a second thought.

Consider just a few more metaphorical expressions that often pass without notice. "Savings and Loan bail out" and "Chrysler bail out," although salient at different points in recent history, appear to come from the same conceptual roots. In literal terms, an entity is in trouble, and is getting outside (government) help. The term bail-out that is used to describe this situation is hardly literal, and evokes the image of something more concrete. Finally, many of us have heard the progress of someone's career described in more concrete terms by saying something similar to, "it is moving {slowly / forward / backward}."

It might be reasonable to expect the user to avoid giving the more intelligent AAC system a blatant metaphor, however the cognitive load associated with recognizing transparently-motivated metaphors is unreasonable. We will need more robust solutions.

Approach

In order to approach a solution to this problem, it is useful to consider why these metaphors are hard to recognize and are such a part of our language environment.

Our understanding of the world influences the way we express ourselves. Metaphorical expressions often reflect conceptual models which are the basis for how we understand the world. Mark Johnson (Johnson 87) has made some interesting observations about the building blocks of thought, most notably that the building blocks of thought are based closely to our bodily experience. Among the building blocks he has described are attraction, blockage, and containment.

This metaphor of "putting money into stock" is based in the simplifying concepts that represent investments as containers which can hold money. The metaphor of the S&L bail-out is based on the much more concrete situation of that of helping to bail out another's boat before it sinks into the water. The example of "a career moving slowly" suggests a background model in which people can view progress in terms of a moving vehicle.

When we write and talk we automatically use non-literal expressions that reflect our common conceptual groundings. These lead to very natural and easily understood expressions because we (speaker and hearer) share these common conceptual groundings. Because they are grounded in the building blocks of our thought, these metaphorical mechanisms and concepts can operate in our minds without our conscious knowledge of them doing so.

Our approach is to supply a more intelligent AAC system with the necessary conceptual groundings to understand/generate these types of metaphors. Doing this will allow a user to express ideas in a way that is most natural. This addition will allow metaphoric expression in places where the system's knowledge would previously allow only strictly literal statements.

<u>Metaphorical Domains</u>

A metaphor maps between two domains. In the "putting money into stock" example, the finance domain is the *tenor* domain and the container domain is the *metaphorical* domain. Computationally, it is very useful to limit the number of possible metaphorical domains. Unfortunately a metaphor can potentially use anything as a metaphorical domain. However, in the case of transparently-motivated metaphors, we can be more specific about what can qualify as a potential metaphorical domain. We have recognized that they must be universal, concrete, and have specialized lexical expressions to clearly identify the domain (e.g., the expression *put in* identifies the container domain).

<u>Metaphorical Domain Selection Rules</u>

Recall that earlier we discussed how conceptual groundings motivate transparently-motivated metaphors. The purpose of the metaphorical domain selection rules is to represent those concepts computationally. We envision rules that will embody the notion that "Stocks can be described as containers for money." Other selection rules would include: "arguments are often described in terms of war", "argument structure can be described in terms of buildings", and "progress can be described in terms of a vehicle moving toward a goal".

Each of these selection rules must have associated with it more specific mapping information. In the stock example, we would need to specify that it is the purchase price of the stock that is put into the stock. We would like such information to be general enough so it could also account for the meaning of "I moved $3,000 from Exxon to a money market account." Here, the concepts associated with *put* and *take* are both found in the verb *move*.

Consider the rule stating that progress can be described in terms of a moving vehicle. Included in (or derived from) the information closely attached to this rule should be the following:

Tenor Domain	Metaphorical Domain
progress	forward
negative progress	backward
no progress	still
unsatisfactory progress	slow

The Problem of Metaphor in AAC Systems

Generating from More General Principles

We are excited about the potential for abstracting beyond the level of information presented above. Recall the "career as moving object" example. Note that the moving object has some starting point, some goal and some points on its path. With time involved, it also has some speed. It appears that with a sophisticated model of this behavior in the metaphorical domain and rules linking it to appropriate types of tenor domains, that the above four mappings could be derived. Interestingly, a more general structure matched with reasoning could yield other derived expressions. With the knowledge that energy is required to move objects, and given that a prototypical moving object is a car that runs on gas, we could hope to handle "My career is running out of gas" *without* an explicit rule for this occurrence. As the level of the information that motivates the metaphors reaches a higher level of abstraction, a greater degree of flexibility is anticipated.

Not all selection rules lie at the same level of abstraction. The selection rules, and their associated mapping knowledge may need to be encoded hierarchically. For example, the concept of possessions as containers of their purchase price may be a more specific version of the general concept that things are containers of what goes into their existence (e.g, "I put a lot of (time, energy, love, money) into this relationship!").

Benefits

It would be undesirable for a more intelligent AAC system to attempt to handle metaphorical expressions without a model of metaphor. One might propose that in order to handle metaphor all we need is to relax some of the knowledge in the system. If nothing else fits, doubt the knowledge in the system about the words *put Exxon* and *3000*. Then, even though the system knows better, it can forget that restrictive knowledge and allow 3000 to be put into Exxon, without a sophisticated model of metaphor.

Unfortunately, this naive approach cannot work satisfactorily. First, if we merely relax, or doubt the knowledge in the system, we have little or no idea what the words mean. Such ambiguity would then cause such a system to merely try a great number of combinations. It would generate all kinds of expressions, perhaps: "Exxon has been put on top of 3000", "Put 3000 Exxons" and "Exxon has put 3000." Note that these examples occur while still assuming that the verb put is still the same, that only the properties of the other words are in flux. If the verb were also relaxed, even stranger constructions would be expected.

Second, the power of such AAC systems comes not only from their static knowledge about the world, but also from the dynamic knowledge that comes from tracking the discourse by understanding previous expressions. This discourse tracking would be diminished, along with the performance of the system, even if it were possible to generate the right metaphor without understanding it.

It would be different and more reasonable to consider ways to use lesser approaches to the problem if it were believed that the phenomena of metaphor was arbitrary with no identifiable structures behind it. Fortunately, the most important metaphorical problem to AAC systems, transparently-motivated metaphors, are based upon a firm structure, as has been described earlier. We do, however, need to do much more work on the problem of specifying the mechanisms behind transparently-motivated metaphors computationally.

Implications

Since transparently-motivated metaphors are often produced without being recognized as non-literal, it is unreasonable to require users of future AAC systems to recognize them. Fortunately, the behavior of transparently-motivated metaphors is not arbitrary, and therefore can be incorporated into a more intelligent AAC system. This will allow for much better behavior of the system, lessen the cognitive load, as well as facilitate completely natural expression.

Acknowledgments

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

References

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.

M. Johnson. The Body in the Mind: The Bodily Basis for Reason and Imagination. University of Chicago Press, Chicago, 1987.

G. Lakoff and M. Johnson. Metaphors we live by. University of Chicago Press, Chicago, 1980.

K. McCoy, P. Demasco, M. Jones, C. Pennington, and C. Rowe. A Domain Independent Semantic Parser For Compansion. In *Proceedings of the 13th Annual RESNA conference*, pages 187-188, RESNA, Washington, DC., June 1990.

M. Jones and K. McCoy. Transparently-Motivated Metaphor Generation. In *Proceedings of the 6th International Workshop on Natural Language Generation*, 1992

Contact

Mark Jones
Applied Science and Engineering Laboratories
A.I. duPont Institute
P.O. Box 269
Wilmington, DE 19899
Email: jones@udel.edu