The Commercial Application of Natural Language Interfaces

Harry Tennant
Computer Science Center
Texas Instruments
Dallas, Texas

I don't think that natural language interfaces are a very good idea. By that I mean conventional natural language interfaces -- the kind where the user types in a question and the system tries to understand it. Oh sure, when (if?) computers have world knowledge that is comparable to what humans need to communicate with each other, natural language interfaces will be easy to build and, depending on what else is available, might be a good way to communicate with computers. But today we are soooooo far away from having that much knowledge in a system, conventional natural language interfaces don't make sense.

There is something different that makes more sense -- NLMenu. It is a combination of menu technology with natural language understanding technology, and it eliminates many of the deficiencies one finds with conventional natural language interfaces while retaining the important benefits.

Since a conventional natural language interface invites the user to type in whatever he wants, it is based on the assumption that it will be able to understand just about anything that the users are likely to say. All systems today deal with limited domains of discourse. I am convinced that users are likely to type in all kinds of things. The probability is high that users will not be understood. They could be taught the limitations of linguistic and conceptual coverage of the system, but a major motivation of building natural language interfaces is to allow effective use without training or retraining. So, it doesn't seem like a very good idea.

The assumption behind NLMenu is the opposite. It assumes that there are all kinds of things that the users would like to ask but that the coverage is so limited that it is best to reveal the coverage (limitations) to the user. He then can find quickly that what he wanted to ask cannot be asked of this system, so he'll give up quickly, minimizing his frustration. Or he might find that what he wanted to ask can be asked and the system helps him ask it in the way it will understand. There is another important advantage: there may be things that the user did not imagine that he could ask about. NLMenu reveals these to the user encouraging him to make full use of system capabilities. Conventional natural language systems do not.

NLMenu works by displaying a collection of menus of words and phrases on the screen. The user builds sentences by selecting words and phrases from the menus. The menus are driven by a parallel parser which maintains all parses of the sentence fragment constructed so far. After each word or phrase is selected, the parser

looks ahead in the grammar for all the phrases that could come next for each parse. A new set of menus is constructed with these phrases, which are then displayed to the user for his next selection. In this way, the NLMenu system is constraining the user to constructing only those sentences that are parsable with the grammar. By including semantic constraints, one can also constrain the user to the conceptual coverage of the system. In our current implementations, we use semantic grammars so syntactic and semantic constraints are conveyed simply. Many other techniques can be imagined.

Because the NLMenu technique rests on the same technologies as conventional natural language interfaces, they have the same expressive power -- one can say the same range of things with either approach. But since NLMenu is a system-initiated dialog, the system always knows what the user is trying to express. This makes it very easy to mix natural language (sentence building) with other interface techniques such as graphical input, form filling and others. For example, in an application involving airports, when the user was about to enter the location, a map popped up, the user pointed at the area of interest, then the map went away and the coordinates of that location were textually inserted into the sentence. The user then continued building the sentence. This allowed the use of text where appropriate (specifying runway lengths, location names, etc.) and graphics where appropriate (specifying locations).

It seems to me that there is much more that user interfaces can gain from natural language research. For example, cooperative response is a good idea independent of whether an original query was expressed in English or a formal query language or through some other means. Similarly, repeated reference is important in any extended dialog. Discourse objects should remain available for terse reference. There is nothing that limits repeated reference to natural language dialogs. Ideas based on focus and dialog structure can be applied to dialogs mediated through all sorts of interface languages. We seem to be concentrating on reproducing the form of human communication and ignoring the substance: large vocabularies, concept creation through reference, modification and analogy, mechanisms that use context to gain terseness and allowing dialogs to accommodate the nonlinear characteristics of human thought. Natural language research has much more to offer the world than simply a means for interpreting typewritten commands, yet we as a field have accomplished little toward influencing the other user interface technologies.