A Machine Translation System for the Target Language Inexpert

Xiuming Huang
Department of Computer Science
Melbourne University
Parkville, Vic. 3052
Australia*
xiuming@trlamct.trl.oz.au

1 Introduction

Almost all commercial machine translation (MT) systems to date are designed for multilingual users, normally with post-editing facilities (Melby, 1987). To cater for the needs of users not skilled in either the source language or the target language an MT system must provide facilities for the user to check into the quality of the translation, without knowing the foreign language. In this note two presumptions are made: first, we discuss only translation from the user's native language into a foreign language (or "export translation"); second, we will restrict the quality checking to accuracy checking only, leaving style considerations aside.

In the following sections we will first discuss approaches to assuring correct translation, and then describe the selective confirmation approach adopted in our MT system.

2 Approaches to Assuring Correct Translation

2.1 Back Translation

The seemingly most natural way of finding out whether the translation into the target language (henceforth *forward translation*) is correct is to translate the forward translation back into the source language (*back translation*). This approach can indeed expose the wrong forward translations if the back translations differ significantly from the original, but it has its complexities and is not reliable. An MT system may either employ grammars designed to be reversible, i.e., to be used for both generation and analy-

sis, or employ separate grammars for generation and analysis. In the former case, theoretically whatever the forward translations are, the back translation will always produce sentences in the source language that are very close, if not identical, to the original. As a result we have no way to tell whether the forward translation is correct or not.

If the back translation is done in a system where separate grammars are used for generation and analysis, then the back translation itself may be incorrect inasmuch as the forward translation might be, so that a *correct* forward translation may be *wrongly* translated back, or vice versa. In either case there is no guarantee of accuracy.

2.2 Paraphrasing

In this approach, the system generates paraphrase(s) for the original sentence, based upon the syntactic and semantic analysis result, before passing it for further processing (by the transfer and/or generation component). The user checks the paraphrase and gives the system directives by either confirming or rejecting it. For example, for the input sentence

(1a) The man saw the woman in the park with the telescope.

the system might produce the following paraphrase for the user to check:

(1b) With the telescope, the man saw the woman who was in the park.

The problem with the paraphrasing approach is that the recovery may come in late, after a considerable amount of time is spent doing all the syntactic and semantic analysis required for generating the paraphrase. If the input sentence is long and complex the cost can be high. Furthermore, it may not be unusual that only after several trials the user finally accepts an appro-

*Current address: AI Systems Section, Telecom Research Laboratories, 770 Blackburn Rd, Clayton 3168, Australia

priate paraphrase, with all the previous efforts wasted.

2.3 Pre-editing

This is the approach adopted by many commercial MT systems (very often paired with post-editing). It requires the user to edit the input sentences before they are passed to the system for analysis and subsequent processing. The machine has a predefined translation capacity which must be known to the user so that anything in the text which may cause difficulties to the system will be removed or rewritten. For example, for the following sentence

(2a) The woman cannot bear children.

if the user knows the system would have difficulties resolving the ambiguous word "bear", s/he can rewrite the sentence as follows (if that is what s/he wants to say):

(2b) The woman cannot give birth to children.

To apply pre-editing, a set of rules must first be devised to set up lexical and structural constraints, then the user must keep the rules in mind and apply them consistently. This may involve expensive training of system users and impose strong restrictions on the system use in practice.

2.4 Interactive Disambiguation

Interactive disambiguation can take place at two levels. At the lexical level, whenever an ambiguous word is encountered, the user can be asked to help. (Carbonell and Tomita, 1987) gives an example of resolving word-sense ambiguity in this approach:

The word "pen" means:

1) a writing pen

2) a play pen

NUMBER?>

The problem with this approach is, given that any natural language is highly polysemous, the frequent occurrence of ambiguity at the lexical level will unnecessarily prolong the translation process and easily bore the user. Moreover, with lexical items with many different senses, it may become very difficult to pinpoint one in particular from a screenful of choices.

At the structural level, ambiguities can also be referred to the user for disambiguation, as is done in *Ntran*, an English-to-Japanese prototype MT system for monolingual users (Wood and Chandler, 1988). For the sentence

(3) The cursor corresponds to the puck position on the tablet.

Ntran asks the user to choose either 1 or 2 from the following:

1 on is location of position

2 on is location of correspond

(here "on" represents "on the tablet".) The user must be more or less familiar with linguistics to make correct choices. This may inconvenience some users; but the more severe problem is that the number of possible interpretations of an ambiguous structure can reach the hundreds (Church and Patil, 1982), making their handling very difficult.

3 Selective Confirmation

The basic idea behind our approach, described in this section, is to let the machine do most of the work without human interference, and only at certain decision making points ask for human assistance. We choose phrases as the level for user confirmation because at this level the system both avoids frequent and unintelligent questioning of the user (as is the case with the interactive disambiguation approach at the lexical level) and does not suffer from late recovery (as is the case with the paraphrase strategy). As an example, let us consider the following sentence.

(4) The talented conductor dated a young star.

The system does not ask for the user's help the first time it sees the word "conductor", but waits till the analysis of the NP with "conductor" as its head is complete just before the tree representation is built. At this point it asks:

Does "conductor" here mean "an official on a bus or train or tram who collects fares"? (y/n)

Here the system selected one of the "human being" senses of the word "conductor" as a result of carrying out semantic matches of the modifier "talented" and the head noun (a tram conductor can certainly be talented). The order of selecting the "bus conductor" sense vs. the "orchestra conductor" one is arbitrary, though with more domain information and statistical consideration preference can be given to one over the other. If the user answers "no" the system backtracks to find another "human being" interpretation of the word:

Does "conductor" here mean "a person who conducts an orchestra" ? (y/n)

The answer from the user at this point is likely to be "yes". If, however, the user insists on

"no" the system will relax the semantic constraints (Huang, 1988) and accepts the "substance" sense of the word, treating "talented" as used metaphorically.

Suppose the user has chosen the "orchestra conductor" sense. The system continues the parsing to process the finite verb "dated" and upon the completion of subject-verb match for deciding a proper sense for the verb which suits the selected sense for "conductor", it asks the user to confirm its choice:

Does "date" here mean "to go out on dates with"? (y/n)

If the user answers "yes" the system carries on to find an interpretation for the object NP, using the chosen sense of "date" to help disambiguate "star":

Does "star" here mean "a celebrity"? (y/n)

Similar to the confirmations described above, when prepositional phrases are processed confirmation is carried out after the semantic matches have resolved the attachment ambiguity of the PPs. Confirmations would also be needed at points where semantic matches are carried out to resolve the coordinate conjunction ambiguity, such as contained in the phrase "the man and the woman with an umbrella" ("[the man] and [the woman with an umbrella]" vs. "[the man and the woman] with an umbrella")¹ (Huang, 1983).

An important consideration in designing the interactive system is the number of questions asked of the user. The less questions asked, the more productive the system will be, and the less bored the user will become.

Ideally the system should be intelligent enough not to ask for confirmation about the word "dated" when processing

(5) ANU dated the world's oldest rock.

or

(6) Ann dated the department's oldest professor.

Whereas if the input sentence is

(7) Ann dated the town's oldest coach.

the user may not feel it unreasonable if the system asks him/her to confirm the disambiguation of "dated", "oldest", and "coach". What makes difference here is the so called "genuine" ambiguity (sentences for which more than

one interpretation is meaningful): sentences (5) and (6) should have only one valid interpretation each according to our common knowledge even though when standing in isolation the word "dated" is ambiguous, whereas sentence (7) might have two meaningful readings.

If user confirmation is required only for sentences containing genuine ambiguity, the system will become much more efficient, without endangering the quality of the translation. But to decide when a word or a structure is "genuinely" ambiguous may involve more computing resources than is worthwhile employing. First of all it has to exhaust the whole search space, both syntactic and semantic, to find out whether more than one meaningful interpretation exists, despite the fact that often the first such interpretation might well be the only one. Secondly it can be very tricky to draw a clear line between "meaningful" and "meaningless" interpretations.

For example, the sentence

(8) John cannot bear children.

might be judged as not "genuinely" ambiguous because common sense dictates that since John is a male (based on world knowledge of names), it is impossible for him to give birth to children, and therefore "bear" can have only the interpretation "to tolerate" in this context. But then if you tell your granddaughter "John cannot bear children because he is a male", you are using "bear" in the other sense, although "tolerate" is still an acceptable interpretation for "bear".

For this reason at the current stage we have not attempted at singling out genuinely ambiguous sentences but instead we employ certain domain information and statistical consideration in arranging the system's lexicon, so that for sentence (4) ("The talented conductor dated a young star.") it may well be the case that the user need only answer "yes" once to each of the questions the system asks. The worst case is when the intended meaning of the sentence is "abnormal" or metaphorical, for example when it is meant for sentence (4) to be interpreted as "the talented tram conductor ascertained the age of a new celestial object". In such cases the user will have to answer the system's questions several times over the same word.

With richer knowledge bases and more powerful inference engines and computing facilities it may become practical to first recognize genuinely ambiguous sentences and then to assign scores to different interpretations of such sentences, and fi-

¹ At this stage only confirmations involving NPs, NP-Verb pairs and Verb-Np pairs are implemented, although in the automatic version of the system all semantic matches have been executed ((Huang, 1987) and (Wilks et al., 1985)).

nally present them to the user for confirmation in the order of their scores.

References

- Carbonell, J. and Tomita, M. (1987). Knowledge-based machine translation, the CMU approach. *Machine Translation: Theoretical and Methodological Issues*, pages 68–89.
- Church, K. and Patil, R. (1982). Coping with syntactic ambiguity or how to put the block in the box on the table. *Technical Report MIT/LCS/TM-216*.
- Huang, X.-M. (1983). Dealing with conjunctions in a machine translation environment. *Proceedings of the 1st Meeting of the European Chapter of the Association for Computational Linguistics*, pages 81–85.
- Huang, X.-M. (1987). XTRA: The design and implementation of a fully automatic machine translation system. *Unpublished Ph.D. thesis*. Also as Memorandum in Computer and Cognitive Science, MCCS-88-121, Computing Research Lab, New Mexico State University, Las Cruces.
- Huang, X.-M. (1988). Semantic analysis in XTRA, an English - chinese machine translation system. *Computers and Translation*, 3(2):101–120.
- Melby, A. (1987). On human-machine interaction in translation. In Nirenburg, S., editor, *Machine Translation: Theoretical and Methodological Issues*, pages 145–154. Cambridge: Cambridge University Press.
- Wilks, Y., Huang, X.-M., and Fass, D. (1985). Syntax, preference and right attachment. *Proceedings of IJCAI85 (1985 International Joint Conf. on Artificial Intelligence)*, pages 779–784.
- Wood, M. M. and Chandler, B. (1988). Machine translation for monolinguals. *Proceedings of COLING88*, pages 760–763.