[Proceedings of the *Conference on Theoretical and Methodological Issues in Machine Translation of Natural Languages*, Colgate University, Hamilton, New York, August 14-16, 1985]

Machine Translation as an Expert Task

Rod Johnson & Pete Whitelock
Centre for Computational Linguistics
University of Manchester Institute of Science and Technology
P.O Box 88
Manchester M60 1QD
United Kingdom

Acknowledgement

The work reported in this paper is supported by International Computers Limited and the UK Science and Engineering Research Council under the Alvey program for research in Intelligent Knowledge Based Systems. We are grateful to ICL and to SERC, as well as to our colleagues in UMIST and the University of Sheffield, for their help and support in putting our ideas into practice.

0. Introduction

The case against fully automatic high quality machine translation (FAHQMT) has been well-canvassed in the literature ever since ALPAC. Although considerable progress in computational linguistics has been made since then, many of the major arguments against FAHQMT still hold (a good resume is given by Martin Kay (1980)).

It is not our intention to reopen the case for FAHQMT here. Rather, we contend that, accepting that FAHQMT is not possible in the current state of the art, it is both feasible and desirable to set up R & D programmes in MT which can both produce results which will satisfy sponsors and provide an environment to support research directed towards bringing MT closer to the ultimate goal of FAHQMT.

This paper describes the rationale and organisation behind one such programme, the UMIST English-Japanese MT project.

1. MT as simulation of translator behaviour

Since an ideal MT system will probably be expected by consumers of translations to exhibit the functional input-output behaviour of an ideal human translator, it is not unreasonable to look to translators as a primary source of information about the problems of MT. Note that we are not saying here that an ideal MT system should necessarily be designed to model every aspect of the behaviour of a human translator. We do believe, though, that important insights into the organisation of MT systems can be gleaned from studying how translators operate — and, more importantly, what kinds of knowledge translators use — when they do translation.

What this claim comes down to is the assertion that translation as currently practised is a task entrusted to experts — the translators. What we try to do when we build an MT system is to incorporate all or part of the translator's expertise into a computer program. If we were able to characterise all of the expertise of the ideal translator in such a way that the characterisation could be expressed as an executable computer program then, presumably, we would have attained FAHQMT.

Since we do not yet know how to achieve such a characterisation, we look for a model which partitions translation knowledge in such a way as to maximise the efficiency of the human/machine collaboration, while at the same time facilitating transfer of responsibility from man to machine as our understanding of the act of translation improves.

2. Knowledge in translation

We postulate that the professional (technical) translator has access to five distinct kinds of knowledge: target language (TL) knowledge; text type knowledge; source language (SL) knowledge; subject area ('real world') knowledge; and contrastive knowledge.

We assume that the first four of these are not contentious: a translator must know both the language in which the translation is to be produced and the language in which the source text is written; (s)he should have sufficient command of the subject area and its associated stylistic conventions to make sense of the source text and to produce a target text which is acceptable to a subject expert TL speaker. It is worth noting here, in passing, that a good translator is normally expected to be able to compensate for lack of expertise in all of these except (typically) the first two, by appropriate use of external sources like native (SL) informants, monolingual subject specialists and reliable reference works. We shall return to this question in section 5.

The question of contrastive knowledge is a little more delicate. Many workers in MT advocate a two-stage translation model in which source and target texts are mediated by a linguistically neutral 'interlingua'. In such a model there is clearly no place for contrastive knowledge, or rather the relevant contrasts are between SL objects and interlingual objects, on the one hand, and TL objects and interlingual objects on the other.

What we intend by contrastive knowledge is present typically in the so-called 'transfer' models of translation, where both SL and TL components map between texts and 'deep' representations or 'interface structures' (IS). An SL (resp TL) IS, although it abstracts away from superficial idiosyncratic properties of texts, is still recognisably an SL (resp TL) representation. The role of contrastive knowledge — which in the limit case may be restricted to simple lexical equivalence — lies in determining how a given SL IS 'translates' to the corresponding (set of) TL IS.

We do not want to enter here into the debate on the relative merits of interlingual versus transfer organisation in models of MT (the issue will doubtless arise many times during the course of this symposium). As will transpire from the rest of the paper, it makes little difference to our organisational proposals whether contrastive knowledge mediates between abstract SL and TL representations or between some SL linguistic representation and some interlingua. The main difference lies in the ease and consistency of formulation of the necessary knowledge by experts in the domain (linguists, lexicographers and translators).

3. A model of translation

The basic model we propose, in over-simplified form, is the familiar transfer scheme diagrammed in Figure 1.

The idea is that some device A(nalysis) applies SL knowledge to a source text to produce a source IS; a second device T(ransfer) applies contrastive knowledge to the source IS to produce a target IS; and finally S(ynthesis) applies TL knowledge to the target IS to produce a target text.

Symbolically:

(1) S(TLK,T(CK,A(SLK,ST) = TT

Figure 1

```
                                              Source Text (ST)
                                                     |
SL Knowledge (SLK) ------------------------- >       A
                                                     |
                                                     V
                                                 Source IS
                                                     |
Contrastive Knowledge (CK) ----------------->        T
                                                     |
                                                     V
                                                 Target IS
                                                     |
TL Knowledge (TLK) --------------------------->      S
                                                     |
                                                     V
                                              Target Text (TT)
```

In addition (not shown in the figure) all three of SL knowledge, Contrastive knowledge and TL knowledge may be modulated by text-type knowledge.

In practice, as we all know, the model (1), even when enriched by text-type knowledge, is pathetically inadequate. For (1) even to have a chance of being useful, we should have to require that all of S, T and A be total and functional. In practice, we know that this is unlikely ever to be the case with natural text.

Thus, we expect that the mapping computed by A (and its 'inverse' S) will be many-to-many (one text may have many corresponding IS, many texts may have the same IS). Similarly, T is likely to be many-to-many, even if T only involves lexical substitution (consider wall vs muro/parete, or veal/calf vs vitello}. Moreover both A and T will almost certainly in practice turn out to be partial (some texts will be ill-formed wrt available SL knowledge, some source items will not be assigned target equivalents by the available contrastive knowledge). The only thing we can reasonably enforce is that S should be total, by placing the requirement on T that it produce only well-formed IS representations.

There is, however, an important difference between the non-determinism inherent in A and T, on the one hand, and S on the other. If A, for a given text, produces multiple IS representations, then we assume that choice between them is not arbitrary and may have significant consequences for the correctness of the translation, although the available SL knowledge is inadequate to distinguish. Similarly, the available contrastive knowledge may be inadequate to disambiguate multiple IS representations, although again the disambiguation may be important for the adequacy of the translation. On the other hand, once we have a target IS, the assumption is that all texts generable from that IS within the constraints of a given text type will be equivalent with respect to translation. If it is possible to derive more than one text from the IS then in principle we do not need extra information to choose between the possibilities and the choice can be purely arbitrary. **

Thus we come up against two kinds of situations where linguistic knowledge in the system is potentially inadequate to meet the requirements of acceptable translation: when SL knowledge cannot disambiguate SL texts or contrastive knowledge cannot disambiguate IS representations; and when with the SL and contrastive knowledge available the system fails to produce any result at all.

In some cases we can remedy such failures simply by adding to the available stock of linguistic knowledge, as when the system fails to translate some text portion just because a word is missing from the dictionary. In many others there is no plausible linguistic solution; these are the cases where it is recognised that what is needed is an injection of subject-area or real-world knowledge.

Unfortunately, there does not exist, to our knowledge, any semantic-pragmatic theory which is sufficiently general and well-defined to allow incorporation into an MT model. Existing MT programs do what they can with what linguistics they have and leave the rest to human intervention. Our own view is also that practical MT should for the foreseeable future be a collaborative enterprise between human and machine. We want to claim, however, that current MT systems are generally not organised so as to make most efficient use of the human contribution. Moreover, we suggest that a well thought out design for an MT system should not only allow more efficient use of human resources but should also provide a useful research environment aimed at enhancing our understanding of the knowledge needed for translation.

In the next section, we look at conventional ways of organising the man-machine partnership, before going on to our own design.

4. The division of labour in MT

Suppose we have a machine which can perform some part of the translation task, assisted by a human expert. There are essentially three points in the translation process when the human can intervene: after the machine has finished, while the machine is operating, or before the machine starts. It is worth remarking that once human intervention has ceased and the machine is left on its own, the machine's knowledge of what remains of the translation task must be complete.

** Footnote

Actually, the situation with respect to T is not so clear-cut. Louis des Tombe has pointed out to me (personal communication) that, under certain very reasonable conditions, for any lexical item which is apparently unambiguous in the SL but ambiguous in the TL (eg the wall vs muro/parete case above) if the ambiguity is resolvable with respect to the source IS then the information for resolving it should also be present in the target IS. Under these circumstances there is no reason why 'disambiguation' should not be done by S, provided the available TL knowledge is sufficiently precise to rule out the inappropriate case. On this view we would also have to accept that S may be partial.

We look briefly at each of the three possibilities in turn.

4.1. After -- post-editing

The safest way to organise man-machine cooperation in translation is to use a human post-editor to verify the output of an MT program, as is done in many large organisations using MT, especially where post-editing of work done by human translators is anyway the norm.

Post-editing is a highly skilled task: a post-editor needs to be an expert in:

- the subject area
- the target language
- the text-type
- contrastive knowledge.

In effect, the post-editor should be at least as skilled in all of these domains as the original translator. When the task of the translator is being done by a machine, it is not at all evident that we can claim that the machine is usefully extending expert capabilities to non-experts. At best, the computer is being used as a tool for the expert to increase productivity.

4.2. During -- interactive MT

A number of systems currently in use display the source text in the screen and provide facilities to allow the operator to build up a translation interactively, usually in a second window. Typically, the facilities provided include a window-oriented word-processor and on-line bilingual glossaries. In addition, such systems tend to offer an interactive 'translation' mode, in which the machine attempts a sentence-by-sentence translation, pausing to prompt the operator to choose from among possible translation options; for example, the system might prompt:

> Shall I translate 'party' as
>
> 1. partido
> 2. fiesta

This way of working does not really differ from the post-editing scenario above. The possibility of interaction is only used to reduce the size of text fragments to be post-edited from full texts to sentence-sized units. Thus, although it appears to increase productivity (Hundt, 1982), it does not relieve the operator of any responsibility for any part of the translation task. The human end of the collaboration still needs to be carried out by an expert operator, who needs to possess all the expert skills of a translator.

4.3. Before -- pre-editing

In the pre-editing case there is at least some part of the translation task for which the machine is totally responsible (that part which happens after the last human intervention). Typically, in pre-editing environments, documents have to be specially drafted in a limited language using a restricted syntax and restricted vocabulary. The bargain is that the user guarantees only to submit input in the restricted language; the system guarantees that it will translate any valid text in that language.

The division of expertise here is quite different. Now the human needs only active, expert knowledge of the restricted language; all other aspects of translation expertise are supplied by the machine.

The neatness of this partition is somewhat illusory, however. The success of such an arrangement depends on being able to design a restricted language which ensures that all of the machine's inherent knowledge sources can operate infallibly: (passive) SL knowledge, subject-area knowledge, contrastive linguistic knowledge, text-type knowledge and (active) TL knowledge. As a consequence these restricted source languages tend to become so specialised and unnatural as to place unreasonable demands on the expertise of the pre-editor.

5. Distribution of knowledge in human and machine translation

None of these characterisations seems to us to offer a completely satisfactory framework for designing MT systems in such a way that they can be made to approximate more and more closely to the performance of an ideal translator.

To get closer to this goal, we look again at the question of the use a human translator makes of available knowledge, with a view to finding a more productive basis for the sharing of expertise between man and machine.

A human translator is, first and foremost, a target language expert, as is evidenced by the practice of large organisations which require translators to translate only into their native language. It is rare for translators also to have expert knowledge of the subject area of the documents which they translate: they are normally expected to compensate for any deficiencies in their expertise by having extremely good contrastive knowledge and by consulting informed sources (reference works and/or subject experts). They are, however, expected to have good knowledge of the text types which they have to translate, since they largely bear the responsibility for the stylistic appropriateness of the translations they produce. Source language knowledge is also required, of course, but that knowledge need only be passive, and can be limited to experience of the written form in the relevant class of text types.

It is instructive to see how this use of knowledge compares to the presuppositions which seem to be built into the majority of commercial MT systems. In both the post-edited and the conventional interactive schemes, it appears that users expect to have to massage the machine's output to make it more acceptable stylistically . "Style clearly seems to be the main problem in post-editing" (Lavorel, 1982). This view is certainly not consistent with the idea of an MT system as a target language expert.

MT systems with only pre-editing come much closer to treating the machine as an expert translator. Where they differ from human translators is in placing strong, even perhaps unreasonable, requirements on the originators of documents as a means of circumventing their own deficiencies.

6. Towards more productive interaction strategies

The model we propose is intermediate between the pre-edited and interactive styles of MT. If the machine is to behave functionally as far as possible like a human translator, then we would like to free the user from any need to know about the target language, so that the machine has to be a TL and a contrastive expert, as well as having text-type knowledge built in. On the other hand, while we anticipate that the system will be more or less deficient in knowledge of the user's SL and in subject-area knowledge, we assume that these deficiencies can be remedied in consultation with a (SL) monolingual operator. In terms of the model of section 3., we now have the picture in Figure 2.

Figure 2

```
        System Contribution                                  User Contribution


                                 Source Text (ST)
                                        |
                                        |
SL Knowledge--------------------------> A < ------------------(non-expert) SL Knowledge
                                        |
                                        V
                                    Source IS
                                        |
                                        |
Contrastive Knowledge ----------------> T < -------------------- Subject-Area Knowledge
                                        |
                                        V
                                    Target IS
                                        |
                                        |
TL Knowledge (TLK) -------------------> S
                                        |
                                        V
                                 Target Text (TT)
```

It is, of course, one thing to say that the system makes up for its own shortcomings by consulting the operator. It is quite another to determine when and how such consultation should take place. Being able to determine when to trigger an interaction depends on an awareness on the part of the system that there is something which it does not know. We can distinguish two such situations:

(a) the input is ill-formed with respect to either A (the analysis) or T (transfer);
(b) the input is ambiguous with respect to either A or T

These two situations may occur, respectively, in cases where (a) A (resp T) is partial, or (b) A (resp T) is not functional.

Now we can (and should) arrange matters so that any construct produced by A can be transferred (i.e. T is total over the domain of outputs of A). This means that interactions triggered by ill-formed input (case (a)) can be localised within A only. We are not enthusiastic about attempts by the system to go it alone in 'repairing' ill-formed input (cf Arnold and Johnson (1984) for discussion), although this does not rule out use by the system of its own SL knowledge to propose plausible reconstructions of the input as prompts to the user.

Case (b) is interesting, in that some apparent ambiguities in analysis may carry over to the TL (and so the system should not waste the user's time trying to resolve them). This observation suggests that this type of interaction should be handled as a part of transfer, utilising contrastive information as a criterion.

In cases of type (a), the system has a text fragment which it 'knows' at analysis time it is unable to translate. The aim is thus to prompt the user to rephrase the input in a form which the analyser can recognise. Thus the system must indirectly use its contrastive knowledge (knowledge of what it can translate) to extract from the user, who has extensive, but non-expert, SL knowledge an acceptable formulation of the input text.

In type (b) cases, what has happened is that the purely linguistic knowledge available to the system is insufficient to distinguish between translationally distinct 'readings' of the text. Hence the appeal to the user's 'real-world' or subject-area knowledge to resolve the ambiguity.

7. Application

We believe the approach advocated in this paper to have two advantages over more orthodox MT system design: it encourages a more efficient and productive sharing of expertise between man and machine; and it provides a useful framework for MT research by allowing the role of the machine to be extended incrementally on the basis of systematic experimentation within an operational environment. Most of the ideas are not original -- indeed the basic principles go back at least as far as Kay (1973). The same principles also seem to have been applied to MT by Tomita (1983). In our case the application domain is an experimental English-Japanese translator for technical documentation. The linguistic knowledge representation and processing we have used is based largely on LFG (Kaplan and Bresnan, 1982). To do justice to the benefits of using a linguistic theory like LFG in an enterprise of this kind would take another paper, and so we do not discuss the question here, although we do want to stress the importance of a well-founded linguistic theory to support a system whose expertise depends crucially on linguistic knowledge.

The project is now one year old, and a small prototype is due for its first demonstration this Autumn.

References

Arnold D.J. & Johnson R.L. (1984)
Robust Processing in Machine Translation
In Proceedings of Coling84, pp 472-475

Hundt M.G. (1982)
Working with the Weidner Machine-Aided Translation System
In Lawson (ed), pp 45-51

Kaplan R.M. & Bresnan J. (1982)
Lexical Functional Grammar: A Formal System for Grammatical Representation
Chapter 4 of J. Bresnan (ed), The Mental Representation of Grammatical Relations, MIT Press, pp 173-281

Kay M. (1973)
The MIND System
In R. Rustin (ed), Natural Language Processing, Algorithmics Press, New York, pp 155-188

Kay M. (1980)
The Proper Place of Men and Machines in Language Translation
Working paper CSL-80-11, Xerox PARC, October 1980

Lavorel B. (1982)
Experience in English-French Post-Editing
In Lawson (ed), pp 105-109

Lawson V. (ed) (1982)
Practical Experience of Machine Translation
North Holland

Tomita M. (1984)
Disambiguating Grammatically Ambiguous Sentences by Asking
In Proceedings of Coling84, pp 476-480