[Published in *Machine Translation*, vol. 7 no.4 (1992/1993), p.259-264]

And after the second generation...

In discussion of the future development of research in machine translation (MT) it is probably useful to make a preliminary, if somewhat simplistic, distinction between two basic types of activity. On the one hand, there are research programmes dedicated to the exploration of methods, techniques and programming languages which are untried in the field. These embrace the small-scale projects which explore the application of a linguistic theory, such as LFG or GB, to some problematic area of MT, the projects which investigate example-based or statistics-based methods, the projects which explore the feasibility of a particular programming environment, such as Prolog, connectionism, etc. The essence of this type of research ('exploratory research') is to push back the frontiers, to open up the field to innovative approaches.

This mode of research may be contrasted with efforts to build prototype MT systems, which may or may not form the basis for operational MT systems. The objective in these projects -- typically on a larger scale and funded for longer periods -- is the development of systems which undertake the whole translation process (with or without human intervention). This type of research (‘prototype research’) is inevitably less innovative. The basic aim is to cover a wide range of linguistic phenomena and to establish the framework for large dictionaries. It has to be based on a relatively solid well-establish foundation of reliable and tested methods and techniques. It cannot change radically in response to the latest advances. Instead, this type of MT research must attempt to provide a flexible but stable framework; its conservatism is a strength rather than a weakness.

It should be remembered that any prototype system with pretensions to eventual commercial viability must work reliably and efficiently, and this means it must have been well tested. Bringing a MT system out of the laboratory into the market place can take many years; Siemens estimated that they would need 15 years to convert the LRC METAL prototype into a commercial system -- in fact they took less time, but nevertheless this, the most sophisticated commercial MT system, is based on linguistic and computational techniques now ten or more years old. It ought to be no criticism of commercial systems that from the perspective of current research they are outmoded. The criteria for their evaluation should not be innovativeness but the quality of output and whether actual performance matches the claims of makers.

In the late 1970s there was a general consensus that the model characterised by Harold Somers as the ‘second generation’ type was the most suitable basis for the development of systems which promised higher quality output than most pre-ALPAC designs. It was linguistics-based (essentially syntax-based), non-interactive (batch processing with post-editing), stratificational. During the 1980s the range of options broadened and it is now obvious that this is not the only possible design for MT ‘research prototypes’.

The evident passing of a long-standing orthodoxy as far as MT prototyping is concerned raises inevitably questions about the foundations of MT research itself. What is often raised is the lack in MT research of an agreed theoretical basis. Researchers borrow much of the theoretical baggage and apparatus of related disciplines, notably artificial intelligence and linguistics. But MT research should not be regarded as simply the application or testing of theories from outside -- even if, legitimately, it can be regarded as an excellent testbed for formal linguistic theories, since the output can be evaluated independently of knowledge about the theoretical basis of the system. Whereas AI research may be seen as the development of computational models which simulate human behaviour in some way, MT research cannot be defined as the development of computer programs which simulate human translation. Rather, MT research aims to develop programs which actually translate; and which are certainly not simulations of human translation. More crudely, while AI may be a theory of intelligent behaviour, MT is not a theory of translation; it *is* translation.

What kind of theory would be relevant for MT research? If MT is seen as essentially an 'applied' science (the application of computational linguistics, etc. to translation) then a theory could be defined in practical terms as the embodiment of those techniques, methods and approaches which are known to give desired results. On the basis of such a theory it should be possible for any researcher to predict the effect of using a particular technique in a given circumstance, e.g. using a particular parsing method to produce a particular desired representation. There may be a considerable body of knowledge which is relevant, both implicit and explicit, and which may be derived from the experience of MT researchers, but what is lacking is the appropriate abstraction. A theoretical basis of this kind could serve, for example, as a corpus of theses or hypotheses for refutation or for confirmation in experimental research projects.

However, such a conception of MT theory is unlikely to be either satisfactory or sufficient, since MT research is not generally seen as a field of 'engineering'. What is intuitively desired is a theory based on principles independent of particular practices. The Eurotra researchers developed a theory which defined translation as a series of transformations of linguistic objects (representations of texts) from a source to a target language. The Carnegie-Mellon University researchers define translation (broadly) as the extraction of a language-independent representation of the 'meaning' of a source text and the rendering of this meaning in a target language. These are both legitimate MT theories, but they are clearly specific to particular system types and cannot be regarded as generally applicable. What is required in MT is a set of principles of universal validity.

The task must involve an abstraction away from specific system types. But first we must be clear what the options are. Systems can be developed for use by translators, by non-translators knowing both source and target language, by people knowing only the target language or by monolinguals knowing only the source language. They can be designed to produce rough translations, which can then be revised or polished, or good quality translations. Input text can be undoctored (unrestricted) or edited or composed in a controlled vocabulary and syntax. Texts can be limited to a particular sublanguage or (in theory) can relate to any text-type within the subject domain of the system's lexicon. (Note that sublanguage systems are not necessarily 'controlled language' systems, and the latter are not necessarily limited to particular sublanguages.) Analysis and transfer can operate with or without human intervention or interaction. As for internal design features, systems can be of the direct 'first generation' type, can be based on an interlingua or can adopt the transfer-based approach. They may use essentially linguistic information or may refer also (or instead) to non-linguistic domain-specific data. They may employ statistical techniques (wholly or partially); they may refer to corpora of example translations (example-based translation), or to databases of set phrases (phrasebook or template translation). Each of these options may be used and combined in various ways, some as primary methods, others as supplementary techniques. It is evidently incumbent upon future researchers to state what type of system they are working towards, who the eventual users might be and what constraints and limitations on input and output they envisage. It is no more certain now than it was in the past that there will be one single solution, even for one specific translation need.

With such a variety of potential system types what can be said which is common to them all? Apparently, little more than the trite banality that MT is a computational process which relates linguistic objects (texts) in different natural languages which "have the same meaning". However, it does pinpoint two basic aspects: that MT must start and end with texts, and that the core processes involve lexical and structural transfer under conditions of 'meaning equivalence'.

Given the discussion by Somers the first point may appear disputable. In my view, the information conveyed in the source text must be the starting point for translation, whether by human or by machine. Many kinds of information may be brought to bear in the processing (interpretation) of the text: morphological, syntactic and semantic information, general knowledge of human behaviour, domain-specific subject knowledge, appropriate registers, statistical information about vocabulary frequencies, etc.; but the symbols of the text are the only possible

starting point. This applies also to messages composed at a terminal by monolinguals communicating with others in another language. The 'text' may be partially formed (ungrammatical, semi-coherent) but it is still a linguistic 'object', and at every stage what is expressed (the 'text') is the starting point for some process: an interaction with the computer for clarification and rewording, or the confirmation of a selected template text. Furthermore, what is actually translated (after interactive negotiation), i.e. what the computer program processes, is an agreed linguistic text. Likewise, the output of a translation program must be a linguistic 'text' (spoken or written, complete or partial); and the evaluation of the effectiveness or success of the computational processes can only be made on a comparison of the input and the output. If there is no source text or no target text then we are dealing not with translation but something else (nonverbal communication?).

The second point is that the core process of any MT system is the 'transfer' of lexical items and structural relations. The processes of analysis and generation, however essential, are supplementary. What distinguishes MT from other natural language processing activities is bilingual lexical and structural transfer, where transfer can be either direct (by substitution) or indirect (via an interlingual representation). In this sense, Somers is absolutely correct to emphasise the centrality of transfer in the design of an MT prototype: decisions about what in the source text is to be analysed and how target texts are to be generated should be conditional upon decisions about what has to be changed or retained in lexical and structural transfer -- whether this is accomplished via interlingua representations or via series of interfaces. In this perspective, monolingual analysis and generation are means of reducing the complexity of bilingual problems to a minimum. One of the most important activity of MT research should be the contrastive analysis of languages in order to isolate those features (syntactic, semantic, pragmatic, conceptual) which are constant across languages and those which must be present for translation to succeed. The aim should be the establishment of a set of feature parameters which are required for lexical and structural transfer and which are independent of any parameters required for interpretation or disambiguation in particular languages. As a simple example, the gender of French and German nouns (masculine, feminine) is not relevant for transfer, but the number (singular, plural) and the 'case' (genitive, or instrumental) is relevant. The differences in English between *fast*, *quick*, and *rapid* may be specific features of English and may be irrelevant for transfer, but the differences in German between *Wand* and *Mauer* may be relevant even if not made explicit in English. The real problems of lexical transfer are of course far more complex. The work of the Eurotra groups on many aspects of contrastive linguistics (e.g. in the area of tense and time) and the investigations of interlingual features in many projects are of course invaluable starting points for such research.

Whereas features of surface structure (e.g. whether a sentence is active or passive) are often expendable in transfer -- and indeed, as Somers makes clear, structure preservation of this kind should be a last resort -- there is at least one type of structure which should be retained. This is informational structure, i.e. the 'rhetorical' sequence in which the message is presented. Human translators preserve the informational structures of original texts fairly faithfully. This may entail destruction of surface syntactic structures. As a simple example, a German sentence such as *Den Studenten gab der Lehrer eine schwere Aufgabe* cannot be rendered in English in an active form without destroying the theme-rheme sequence; it must be a passive such as *The students were given a difficult exercise by the teacher*. For transfer then in this case, active/passive is an expendable feature, but theme/rheme is not. If informational structure is not preserved then it might be argued the result is not a translation but a paraphrase.

The crux of the matter is, of course, the principle of 'meaning preservation', and we are little further forward in defining this satisfactorily than thirty years ago. What the principle does support, however, is the argument that if for a lexical or structural ambiguity in the source text there is a fully equivalent lexical or structural element in the target language with the same range of ambiguity, then they is no absolute necessity to disambiguate the source. (For example, quantifier

scope ambiguity and the uncertainty of prepositional phrase dependency may be unchanged in transfer.) It supports also the argument that target texts should not be more explicit than their source (as might occur in conversion via a conceptual representation), since this implies non-equivalence of meaning.

In the first steps towards a theory for MT, we could do no worse than examine in detail why external theories are deficient and how they have to be adapted to work in the MT environment. What has to be added to Government-Binding Theory to make it suitable for MT? What is the peculiarity of MT that obliged the relaxation of the isomorphy principle in Rosetta? What features of the CMU interlingua distinguish it from semantic representations in AI monolingual projects? The aim should be to isolate what is apparently intrinsic to computer-based translation. On this basis, hypotheses can be stated and be open to test and refutation.

It is clearly desirable for the health of MT research that new 'research models' are developed, new approaches and techniques are explored (statistical analysis, sublanguage systems, reversible grammars, isomorphic grammars, systems for monolinguals, alternative interactive modes, etc.), that particular problem areas of MT are examined (particularly lexical and structural transfer, anaphora disambiguation, stylistic improvement in generation, etc.) However, these studies ought not to be conducted as independent exercises; at the very least they should be seen to be tests of clearly formulated hypotheses about some aspect of translation (not just about analysis or generation) and should provide the basis for comparison with other approaches. In this respect a standard corpus of source and target texts for benchmark tests is an urgent desideratum; without such measures there is no way of judging whether we are making progress or not. Equally desirable is a 'standard' dictionary base which could be used by researchers for testing new parsing or transfer techniques without having to build their own lexicon each time.

The development of 'prototype systems' has to be more conservative, based on well-tested techniques, and within clear frameworks. The 'second generation' design model has come to an end primarily because advances within and outside MT have presented so many alternatives to the basic premises that the whole framework is called into question. There is no obvious successor, since there is no agreement on the basic principles of MT, and there is no prospect of that in the immediate future. Just as many commercial systems are based on 'direct translation' models of the first generation -- because they use methods which are known to work -- so is there little doubt that basic features of the stratificational syntax-based model of the 'second generation' will also survive. Workers on MT prototypes are obliged to discover the good points of any system type, to find out what worked in older systems, what did not work and why; the historical perspective is important if mistakes are not to be repeated and wheels rediscovered ad infinitum. There is no cause for pessimism here, it is a normal characteristic of most research activity whether 'applied' or not.

John Hutchins
University of East Anglia
Norwich, England