

Languages and machines

John Hutchins

It was under this title that the Automated Language Processing Advisory Committee (ALPAC) of the US National Academy of Sciences reported in 1966 on the prospects for machine translation (MT). The first suggestion that automatic translation was an appropriate task for computers had been made by Warren Weaver of the Rockefeller Foundation some twenty years earlier, but research began in earnest on a large scale only after a demonstration of a small Russian-English translation program in 1954 had impressed the general public with the possibilities and encouraged substantial funding from US governmental agencies. After ten years, however, there were no machine translation (MT) systems in operation producing translations of a quality comparable with those of human translators. The ALPAC report concluded that progress on MT had been unsatisfactory, and it recommended the cessation of funding for MT in the United States; instead, it recommended the support of fundamental research in computational linguistics and the development of computer-based aids for translators.

This report was interpreted as a definitive judgment of the ‘failure’ of MT. To this day, it is widely believed (not just by the general public but by linguists, translators and computer scientists) that MT failed. But what ALPAC announced was not the failure of MT as such but the failure of the excessively ambitious and unrealistic ideal of fully automatic translation of texts on any subject. Even at the time of ALPAC there were working MT systems producing useful (by no means perfect) translations which in unedited rough form were serving the needs of information gatherers and evaluators, or which after appropriate human revision ('post-editing') could be used as the basis for adequate publishable translations.

After the ALPAC report, research continued on MT with markedly less ambitious aims. There have been four main strands. Firstly the development of computer-based tools for translators (as recommended by ALPAC), e.g. multilingual word processing, concordances, glossaries, terminology management facilities, and now the successfully marketed translation workstations. In this sphere Alan Melby has himself been an influential advocate and active developer for many years.

The second strand has been the continued development of MT systems with post-editing facilities for large organisations: e.g. the European Communities, multinational companies such as General Motors and Aerospatiale. The major systems in this sphere have included Systran and Logos, building upon MT expertise going back to the pre-ALPAC period. From these systems are now produced millions of translated pages every year, and the numbers are increasing¹.

The third major strand has grown primarily since the mid 1960s. This is the design or adaptation of MT systems to operate within specific sublanguages or domains, or with controlled language input. The earliest success was the Meteo system for translating weather reports in Canada from English into French. Another early significant example was the use of Systran by the Xerox company for the translation of technical manuals written in a controlled subset of English. As Melby illustrates in this book, controlled domain-specific MT is capable of good quality output with little or no human

revision or interaction. In recent years, many MT systems have been specifically designed to take advantage of this feature. And we are now seeing research specifically on controlled-language MT, e.g. systems for business people ignorant of target languages who want to communicate 'standard' messages, and (more ambitiously) systems for translating spoken telephone messages to make conference registrations and hotel bookings.

The final strand is the appearance of systems for dealing with unrestricted input of colloquial language. With the increasing use of global telecommunications (e.g. the Internet) there is clearly a huge need for the translation of short electronic messages on almost any topic. (On CompuServe, such a service is already available.) In addition, the appearance on the market of cheap PC-based translation software has stimulated a latent demand for 'occasional' personal and business translation by monolinguals and non-translators. Furthermore, there is the continuing need to translate documents for information purposes. For reasons which Melby explains, the MT of unrestricted text is much more difficult and less successful than the MT of controlled texts. In this area, then, MT is providing a crude tool with many imperfections which is nevertheless satisfying identifiable needs.

Alongside these clear demonstrations of the practical everyday usefulness of MT, there continues to be research on MT with the clear and obvious aim of improving the quality of MT in general. It is here that the idealism of the early period still has echoes, and it is this field of activity which is the main focus of Melby's book.

The first efforts in MT research were necessarily constrained by limitations in hardware, inadequate storage facilities and the lack of suitable programming languages. More crucially there was no help from the language experts: syntax was a relatively neglected area of linguistic research and semantics was virtually ignored in the United States thanks to the behaviourist inclinations of the leading scholars. However, within a few years there were dramatic improvements in hardware and software and there were exciting developments in linguistic theory which gave rise to the belief that good quality output comparable to that of human translators was achievable within a matter of a few years. It became widely assumed that the goal of MT must be the development of fully automatic systems producing high quality translations. The use of human assistance was regarded as an interim arrangement: post-editing ought to wither away as systems improved. The emphasis of research was therefore on the search for theories and methods for the achievement of 'perfect' translations.

Although the effect of ALPAC was to encourage more realistic objectives (with the successes outlined above), MT research continued on much the same lines as before. There was a firm belief that the systems developed during the "first generation" (before ALPAC) and which were now being brought into operation (Systran, Logos, etc.) were inherently incapable of further improvement. They were considered to be deficient from the perspective of theoretical linguistics. Thus arose after ALPAC the "second generation" architectures in which source languages were analyzed into abstract representations of the meaning of sentences, from which target texts were generated. The systems were explicit applications of models, rules and formalisms derived directly or indirectly from linguistics. Some argued that the representations should be interlingual language-neutral or language-independent representations, in order that new language pairs could be readily added to the system architecture. Others held that representations should be disambiguated abstractions of source sentences, and that they should then be transformed first into similar unambiguous representations of

the target language before generation into target sentences. In this case, lexical items (root forms) and basic structures were converted by a bilingual 'transfer' module of relatively simple form. Both approaches assumed that lexical and structural equivalences can be formulated in computationally appropriate rules.

Many models of linguistic formalism have been explored in MT (Transformational Generative Grammar, Categorical Grammar, Lexical-Functional Grammar, Head-Driven Phrase Structure Grammar, Principles-and-Parameters Grammar, etc.); indeed from one perspective it can be argued that MT has often served as a testbed for theories developed elsewhere - for the simple reason that whether a translation succeeds or not is relatively easy to evaluate. It is of course quite another and more complex matter to judge the quality and appropriateness of translations for specific needs; but, as far as the theorists are concerned, the production of any reasonable translation is acceptable - improvements, they would contend, can come over time with greater sophistication and elaboration of the rules.

However, it is precisely this last assumption that has proved to be mistaken time and again; beyond a certain level of quality, it has been found that rule-based systems become increasingly more complex and more difficult to improve. The solution of one outstanding problem tends to degrade some other aspect of the system previously working satisfactorily. In this respect the linguistics-based systems of the second generation appear to be no more capable of improvement than those of the first generation.

In recent years, MT researchers have looked to other methodologies. First, knowledge-based methods familiar in AI research were adopted; but ultimately, this variant of the rule-based approach showed similar problematic aspects; in particular, generalisability from relatively small-scale systems has proved difficult. It is notable that the best known group using in this approach (at Carnegie Mellon University) is now developing a controlled-language domain-specific system in collaboration with the Caterpillar Corporation.

More significantly, however, there have emerged in the last five years corpus-based methods which explicitly reject the rule-based methods of the linguistic systems. There are two main directions: the statistical methods developed by an IBM group on the basis of the vast Canadian Hansard corpus of English and French texts; and the example-based methods adopted by many Japanese research groups which use aligned bilingual corpora of (usually human) translated text to match segments of source sentences against possible target outputs. In neither case is there any linguistic analysis in the traditional sense (i.e. morphological, syntactic, semantic analysis) of the source texts².

Nevertheless, few MT researchers believe that these corpus-based methods will supersede completely the rule-based methods. Since systems would have to be 'trained' on specific text corpora, it would seem on the face of it that corpus-based approaches will be most appropriate in restricted domains; and that they will be no more (and perhaps less) successful with unrestricted general language corpora than rule-based systems. Most believe that future MT research systems will be hybrids, selecting the best and most effective aspects from both rule-based and corpus-based methods.

Whatever the future development of MT research, there remains the fundamental question which

Alan Melby addresses in this book. Is it possible in principle to design and construct a fully automatic system to translate unrestricted general language at a performance level comparable to that of human translators? We know that it cannot be done at present, but is it perhaps just a matter of time (perhaps a very long time) before methods are developed to deal with everyday language? And if not, then why not?³

Melby's answer begins by marshalling the major objections (some well known, some less known) to the 'universalist' theoretical linguists who believe that they are investigating the very foundations of all human language. By focusing on the universals of language they have neglected the variability and differences of languages and the communicative contexts of language itself. By seeking models of linguistic competence and language acquisition they have neglected models of linguistic performance. In brief, they have developed theories of language which are divorced from languages in use. Consequently linguistic theory is of little value to the processing of texts and to the treatment of bilingual and multilingual communication essential for translation⁴. Indeed, Melby contends that linguistic theory is irrelevant to human translation because it cannot account for ordinary normal language and communication. And it is relevant in MT only for those systems which deal with regularised and controlled 'artificial' special languages.

Of course, not all linguists share universalist aims; many are more concerned with the social, psychological and pragmatic dimensions of languages in action. It is in fact arguable that these other aspects of language, which are probably more relevant to translation, are in principle unformalisable – at least, not reducible to rules and procedures suitable for computational treatment. It is, however, the computationally attractive formalisations of the 'universalist' linguistic theories which have had and continue to have such influence on many MT researchers. In this book, Melby directs his argument mainly towards the Chomskyan brand of linguistic theory, but his criticisms are equally valid for many of the varieties of formal linguistics which have been applied and adopted in MT, since the issues concern the basic assumptions of these theories about the nature of language, meaning and communication.

Melby's argument is relevant specifically to the rule-based linguistics-inspired MT research. Its relevance to the newer corpus-based approaches is more indirect. It is a basic assumption of MT that it is possible for any given text to produce the 'correct' translation (or at least the 'best possible' translation) if the equivalences between source and target vocabularies and structures are correctly defined and recorded, and if the relevant contextual factors can be identified and brought into play appropriately. The assumption is as relevant to the statistically derived (probabilistic) equivalences and contexts of the corpus-based approaches as it is (more obviously) to the rule-specified equivalences and contexts of the linguistics-based systems. Melby's thesis is that underlying this assumption is a belief in the objective 'meanings' independent of human communication.

If the computer modelling or implementation of a theory can be regarded as proof of the validity of its assumptions (although this itself is disputable), then the automatic translation of ordinary unrestricted language as good as human translation would be proof of the 'objectivist' view of meaning which lies at the root of much of linguistics, science and philosophy. However, it does have to be translation from and into languages which are not restricted in any way to specific domains or sublanguages; automatic translation of 'unnatural' (standardised, homogenised, normalised, etc.) language is not sufficient.

On the other hand, if it can be proved that unrestricted MT is impossible, then (as Melby stresses) there are significant ramifications and implications, not just for translation both human and automatic, but also for linguistic theory, computer science and the philosophy of language. With such a thesis it will be a surprise if readers agree with everything Melby writes. Some will want to reject his main argument, others will dispute particular details or inferences. No readers are going to be neutral or indifferent. Melby's arguments deserve the most careful consideration by all those concerned with the fundamental aims and the future prospects of both human and machine translation.

Notes

¹. The rapid growth in the use of MT systems is reported in a survey by Muriel Vasconcellos: 'The present state of machine translation usage technology, or: How do I use thee? Let me count the ways!' in **MT Summit IV: International Cooperation for Global Communication**, July 20-22, 1993, Kobe, Japan, pp.35-46. Also published in **MT News International** 6 (September 1993), pp.12-17.

². For a critical review of recent developments in MT research see Harold Somers: 'Current research in machine translation' *Machine Translation* 7(4), 1993, pp.231-246. A general overview of current research can be found in John Hutchins: 'Latest developments in machine translation technology' in **MT Summit IV**, July 20-22 1993, Kobe, Japan, pp.11-34; and in John Hutchins: 'A new era in machine translation research' in **Translating and the Computer** 16...10-11 November 1994 (London: Aslib, 1994), pp.1-13.

³. The conundrum is akin to that faced by researchers in Artificial Intelligence where it has been found that systems can cope with formalised, rule-specifiable expert knowledge but that ordinary common sense has so far defied all attempts at computational treatment. Is this too impossible in principle? and if so, why? The answer may well be related to Melby's thesis in this book.

⁴. In this respect the title of the ALPAC report had an aptness which the authors were perhaps only dimly aware of: in the context of MT, computers must deal with the variations and differences of living **languages** in use and not with the generalities and universals of **language** in the abstract.