Machine-assisted translation with a human face

by KENNETH R BEESLEY

Machine-assisted translation and machine translation, with their inevitable acronyms MAT and MT, are terms commonly used to describe a wide variety of computer systems which help translate natural languages, such as English to French.

Abstract: Machine-assisted translation (MAT) has come under much criticism, both from government investors and human translators. However, modern MAT systems are of a high quality, although not perfect. Most work has been directed at automatic machine translation, which does not require human help. The model used is naive for human and machine translators. Some systems divide the work between the machine and the human. In evaluating any MAT or translator aid, the user needs to consider the overall process from speed to type of text to audience. The most neglected aspect of MAT is the needs, attitudes and sensitivities of the human translators.

Keywords: data processing, automatic translation, machine-assisted translation.

Kenneth Beesley is in the research and development department of ALP Systems.

Other terms found in the literature include machine-aided translation, human-assisted machine translation, machine-assisted human translation, computer translation, computer-assisted (or computer-aided) translation and automatic translation. Just as there is little agreement about which terms to use, and exactly what they mean, there are sharp disagreements about the value and future of the whole field. Some researchers, and even more translators, remain convinced that computers cannot possibly be of help in a craft, or art, as subtle as language translation. they cite famous, but probably apocryphal, examples of computer systems translating 'The spirit is willing, but the flesh is weak' into Russian and back as 'The liquor is strong, but the meat is insipid' and a dozen other suspicious variations. On the other side, proponents maintain that computers are not only valuable in translation, but that they are becoming absolutely necessary to help handle the world information explosion.

In response to the sceptics, there is one unassailable fact: computer systems are being used right now, every day, around the world, for automating translation in government agencies, military installations, translation houses and dozens of commercial companies. In most cases, these users are reasonably happy with the results, and they save money and/or time in the total translation process. ALP Systems is one of several software companies which specialize in creating and selling translation systems.

Machine-assisted translation is today a reality. In the face of this optimism, it should be noted that MT and MAT are not magical or perfect. In the majority of commercial applications, where translations must be of very high quality, human translators must work together with computers to get quality output. There are many ways for people and machines to divide the work, and there are correspondingly many approaches to translation automation, some helpful and some only frustrating.

Short history of MAT

The history of MT and MAT in the Western world goes back at least as far as World War II, when both the UK and the USA built primitive computers to help break codes and ciphers. People began to realise that machines could manipulate letters and words, as well as numbers. As early as 1946, researchers on both sides of the Atlantic were talking about using these new machines and code-breaking techniques to translate documents from one natural language into another.

In the 1950s and 1960s, there were numerous MT research projects at US universities, often funded by government intelligence agencies. Not surprisingly, the emphasis during the Cold War was on translating Russian into English.

Because computers, software and computational linguistics were all in

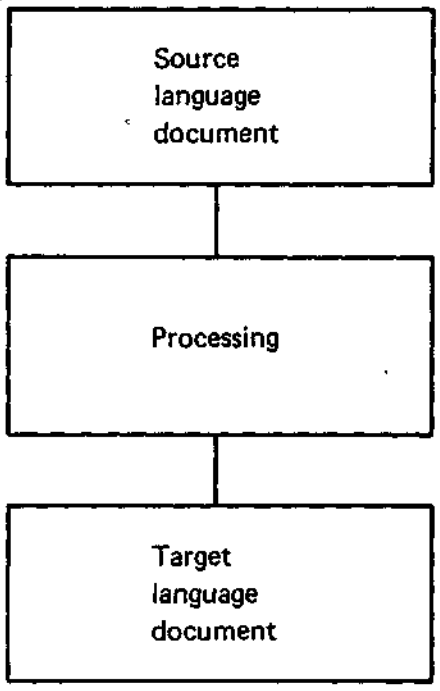

Figure 1. A simple model of translation

their theoretical infancy, the early systems were naturally rather primitive. Translation was little more than word-for-word substitution with *ad hoc* adjustments for inflection, conjugation and word order.

From about 1954 to 1964, the US government was enthusiastic about MT, pumping $20M into various projects. However, by 1964 the government was starting to get suspicious. It wanted practical results and seemed to be getting none. The Automated Language Processing Advisory Committee (ALPAC) was formed to look at MT and make recommendations. Its report, issued in 1966, was very negative, concluding that MT produced very poor output, cost more than normal human translation, and was slower than human translation. It also claimed that MT was not necessary, because the translation load was being handled by human translators without difficulty. Finally, it recommended that no more money be spent on MT.

Although the ALPAC report was, and still is, widely criticized for being one-sided and for relying on outdated information, the fatal recommendations were accepted. Many research projects were cut off and MT research became not only scarce but rather disreputable. Work continued in Canada, France, Germany and in a few commercial companies in the USA.

To this day there is a dearth of university research in MT and MAT in the USA. The University of Saarbrucken in Germany has a research system, as does the University of Grenoble in France. There are reportedly 18 systems under development in Japan and a few in Eastern Europe. The German firm Siemens has funded research at the University of Texas. The European Community is currently funding a vast research project in the hope that MAT will help handle the translation load that comes with having seven official languages, which means 42 language-translation pairs.

Since 1966, when ALPAC concluded that MT was not cost-effective, a number of advances have changed the picture. For one thing, hardware is vastly more powerful and less expensive that it was in 1964. Optical character recognition and word processing cut down on expensive manual typing and retyping. Designers now know how to make computer systems that are easy to learn and to use. Finally, the last 20 years have seen the growth and fruition of theoretical disciplines such as formal linguistics, computational linguistics and artificial intelligence. Machine analysis of language has become increasingly sophisticated.

One sign of the changing climate is the growth of custom-designed, special-purpose MAT systems in organizations such as The Canadian Weather Bureau, the Pan-American Health Organization, The Chinese University of Hong Kong and the French Textile Institute. It is important to note that these are not research projects but actual production systems which are used routinely to help translate specialized internal documents.

Perhaps the clearest sign that MAT has arrived is the rise and persistence of commercial MAT companies. It is

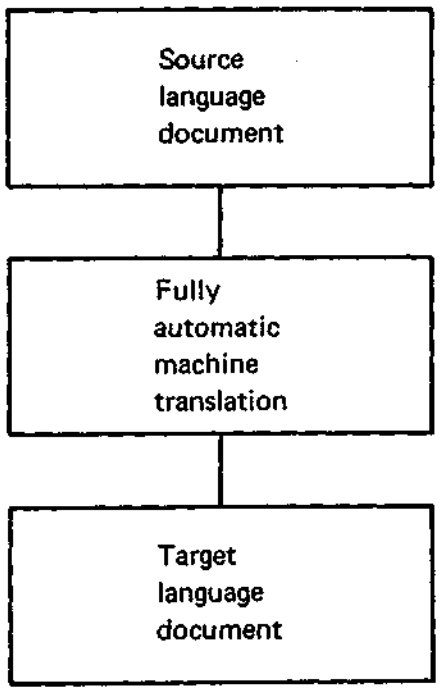

Figure 2. Original idealistic approach to translation automation

now possible for a company or government agency to buy a MAT system 'off the shelf', tailor it for their internal style and vocabulary, and then proceed with some part of their translation process automated. ALP Systems, for example, sells general-purpose translation systems which meet the requirements of the international business community. Other companies selling MAT systems include Logos, Systran and Weidner. The companies compete not so much against one another as against the general inertia behind doing translation the old way. That is, their biggest competitor is always the absence of MAT.

Models of commercial MAT

Looked at in its most simple and simplistic form, any kind of translation, human or automated, starts with a source-language text, puts it through some kind of processing, and ends up with a target-language text (see Figure 1). In a human translation, this processing would seem to be largely mental. When considering how a computer, a mechanical brain, could help in translation, the first and

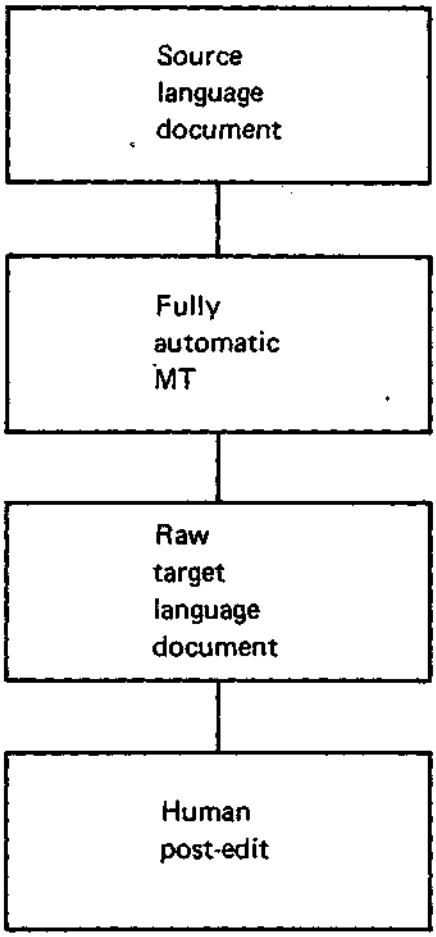

Figure 3. Automatic MT with post-editing the most common commercial approach

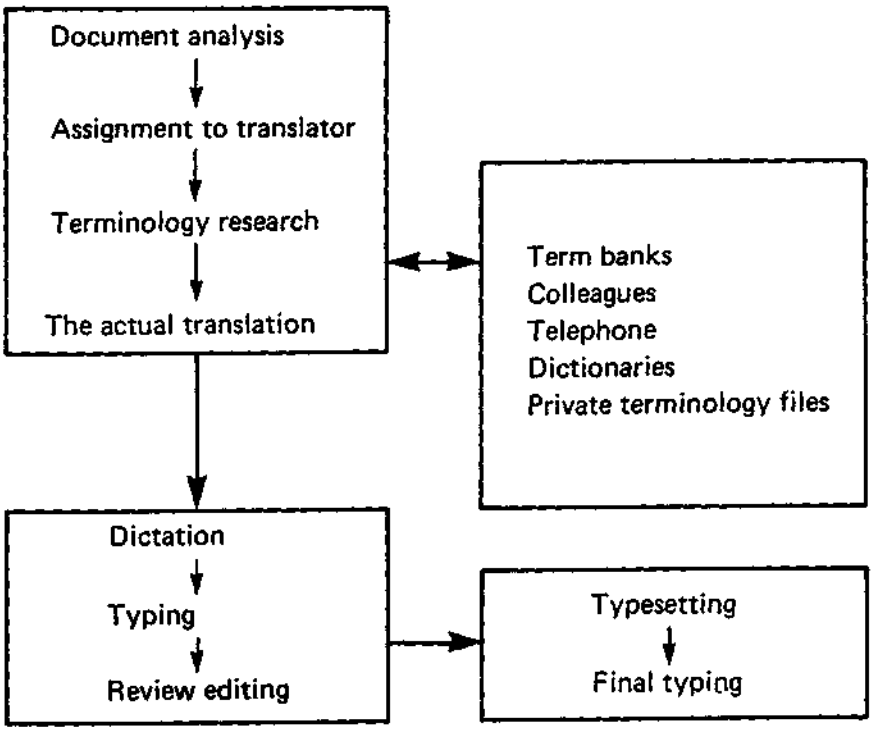

Figure 4. Translation is a complex process

most natural answer has been to make the computer handle all the processing automatically.

Almost all the work in machine translation has been directed at developing automatic machine translation, systems that will take an English text, for example, and automatically translate it into, say, French, without any outside human help at all (see Figure 2). This automatic translation model underlies most commercial MAT systems.

Obviously, the ideal for any automatic machine translation system is to produce perfect output, with accurate meaning, correct grammar and appropriate style. In reality, of course, the results are too often disappointing; translation is a difficult and often subtle task. This is not to say that the results of automatic machine translation are useless. Raw output may be sufficient for some purposes, such as information gathering. If more polished output is needed, the raw output can be placed in the hands of a human post-editor who fixes it up, perhaps online using a text editor. So a post-editing phase is usually added to the automatic MT model (see Figure 3).

Another common, but not necessary, feature of automatic translation systems is that they translate whole documents in batch mode. That is, the computer translates the whole document into the target language before the human post-editing starts. The post-editor, therefore, is presented with the whole target text as a *fait accompli*, and he or she must do with it what he can. He cannot influence the quality during the translation phase itself, and he is reduced to a machine output revisor.

Such a model of translation is naive for both human and machine translators. In reality, the traditional human process of converting a document from one language to another is highly complex, often involving several people and numerous dictionary or terminology sources (see Figure 4).

A source-language document to be translated by a translation house might first go through the hands of a translation manager, who will evaluate the document and decide which of his staff or freelance translators is best qualified to translate it. For instance, a translator adept at medical articles might be a failure with legal documents.

While translating, the translator may have access to general dictionaries, technical dictionaries and often a privately maintained file of new terminology (sometimes kept on cards in a shoebox). It may be necessary to undertake terminology research.

In addition, translators can consult their colleagues and even call up the author to ask him what he meant, some translators in Canada have listed the telephone book as their most important reference work.

Translators often dictate their translations, leaving them to be transcribed later by typists. First drafts may be revised by the translator and typed again. After the translator finishes his draft, it often goes through at least one post-edit to ensure accuracy and stylistic coherence. The final draft may be retyped or sent along for typesetting.

The automatic translation model concentrates the computer automation on the linguistic translation itself,

Taum-Metro	Weather Reports	Automatic Sorting	Easy text c. 80%	Automatic MT	Final text
			Hard text c. 20%	Human translation	
Titus, Colt	Controlled writing/ pre-editing		Automatic MT	Raw Output	Human post-editing
Most commercial systems	Normal technical text	Automatic MT	Raw output	Human post-editing	

Figure 5. Human translators and computers — ways to divide the work

probably the least understood step in the overall translation process. If one looks at what happens in the traditional human translation process, it becomes obvious that many steps are candidates for computer automation. If translators used powerful multilingual text editors instead of dictaphones, much time and expensive retyping could be saved. Text editors can also interface directly with computer typesetters, and electronic transmission of text controls the office paper shuffle. Writing aids such as spelling and style checkers can help both the translator and the review editors. Online term banks can replace the shoebox full of file cards.

While it is now respectable to talk about and produce a variety of translator aids, for years it has been labelled simplistic or uninteresting, the emphasis has been on forcing the machine to perform the actual translation without ever asking for help, something that even a human translator would find difficult.

Problems are swept under the carpet of post-editing, which is underplayed as much as possible. Granted, all the appeal lies in fully automatic translation: it is immensely gratifying to feed an English text into a machine and watch it print out totally automatically in French, especially if you cannot read the French which emerges.

It is crucial to evaluate any MAT system with the overall translation process in mind. The actual conversion from English words and grammar to French words and grammar is only one part of the obstacle course leading from source text to equivalent target text. Gains in this one area may be offset by the cost of putting the source text in machine-readable form. The cost of post-editing raw output may also be prohibitive. Claims about MAT increasing speed and translator productivity are easily made and should be scrutinized objectively.

Human translators and machines — dividing the work

In most MAT applications, translators and machines must somehow work together to get acceptable output. The most interesting classification of such systems, from a user's point of view, hinges on how the translators and machines divide the work.

Five people came to the party

Figure 6. The system divides the text segments which are displayed one-at-a-time in a window

The simplest division of labour is illustrated by the Taum-Meteo system in Canada, which translates weather reports from English into French (see Figure 5). As the weather reports come in electronically from around the country, they are sorted by the machine (essentially an automatic translation manager) into easy and hard text.

The easy text, about 80% of the total, is translated automatically by the machine, and the output is used without post-editing. The remaining 20%, harder and more interesting, is sent directly to human translators.

The success of this system is partly due to the limited subject matter (the weather), the resultingly limited vocabulary and the restricted telegraphic style of the reports. Using Taum-Meteo, the Canadian government is happier, and so are the translators, who used to last about six months before the tedium of translating weather reports drove them to other jobs.

Sometimes the raw output of an automatic system can be improved by pre-editing the text to be translated, solving potential problems at an early stage. There are some constructions in English which are harder than others for machines to handle. These include conjunctions, prepositional phrases and noun-noun compounds as in *small car factory* and *aluminium cylinder head gasket*. The Cult system at the Chinese University of Hong Kong, which is used to translate Chinese mathematical papers into English, requires extensive pre-

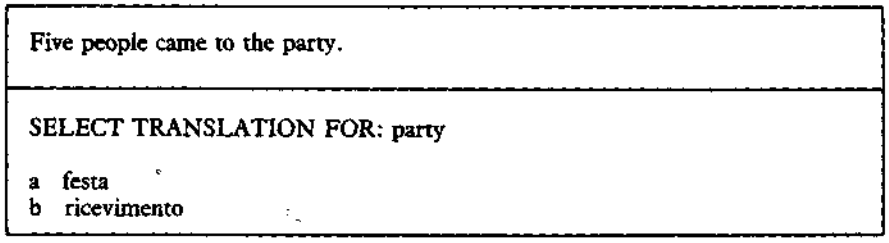

Figure 7. The user is presented with choices for translating words

editing of the Chinese input. Syntactic constituent boundaries, ambiguous constructions and some word categories must be marked by hand before the automatic translation starts (see Figure 5).

The Titus system used by the Textile Institute of France is used only for translating textile abstracts for purposes of information gathering. Every abstract that comes in has to be extensively rewritten into a standard form to ease the translation process (see Figure 5).

If an organization translates only its own internally produced documents, then it can require its writers to use a specially-defined, customized language which avoids problem constructions. The effect is similar to but less drastic than rewriting or pre-editing. The hope in pre-editing and controlled writing is that the raw output will be better, requiring less human post-editing.

By far the most common approach to MAT today is to let the machine translate normal technical text fully automatically, producing a raw translation which is then turned over to human post-editors (see Figure 3). If the output is good, and it can be in limited applications, then this post-editing step is relatively painless. If the output is poor, post-editing can be a frustrating and demoralizing chore.

The new goal in most modern MAT systems is not to do away with post-editing entirely — even human translations are post-edited to some extent — but to keep it down to an acceptable, cost-effective amount. Commercial systems designed for automatic batch translation and human post-editing include Systran, Weidner and Logos.

Both the post-editing and pre-editing strategies just mentioned assume an automatic machine translation stage where the computer system produces a raw translation without human assistance. The most important point to understand here is that the ALP Systems approach to MAT is different. The ALP Systems Transactive product relies on interactive translation, involving close teamwork between the system and the translator during the translation process itself. In addition, it does not translate whole documents in batch mode — rather it presents its results sentence-by-sentence for immediate human review.

Using this system is an interesting exercise in human/machine interaction. The process starts with the translator designating the text to be translated. Next the translator consults a library of dictionaries and designates those general and technical dictionaries appropriate for the text to be translated. In a process called Dixtraction, a new document dictionary containing only those words needed for the text is extracted from the chosen dictionaries.

The ALP Systems program then divides up the text into segments, usually one sentence long, which are displayed one-at-a-time in a screen window (see Figure 6). All the words are looked up automatically in the document dictionary.

Many potential syntactic and lexical ambiguities are resolved by the system itself, but when it comes to clear an ambiguity in the source text, the program simply asks for help to the user, presenting him with questions designed as simple menus. The ambiguities commonly found in a source text are usually of three kinds: vocabulary, grammar and syntax. Here the user is being asked to select an Italian translation for *party* — there are two equivalents in the dictionary: *festa* and *ricevimento* (see Figure 7). By typing 'a', the user selects *festa* and processing continues.

After the segment has been clarified according to the user's answers, the system then displays its own first-draft translation of the segment in a second window (see Figure 8). Note that the selection of *festa*, which has feminine gendre, has had a ripple effect, causing the choice of the feminine article *la* — with contractions, 'to the party' comes out as *alla festa*. If *ricevimento*, a masculine word, had been chosen, 'to the party' would have come out as *al ricevimento*. If a fully automatic system chooses the wrong translation equivalent, it can spread throughout a document. Fixing it during post-editing is not as simple as searching for and replacing the offending word with the correct one, the ripple effect can also require changes to related articles, contracted articles, adjectives, participles, verbs and other words.

In the ALP Systems program, any necessary fine-tuning of the output is done immediately, sentence-by-sentence. The user can easily delete, move, add or amend words using a

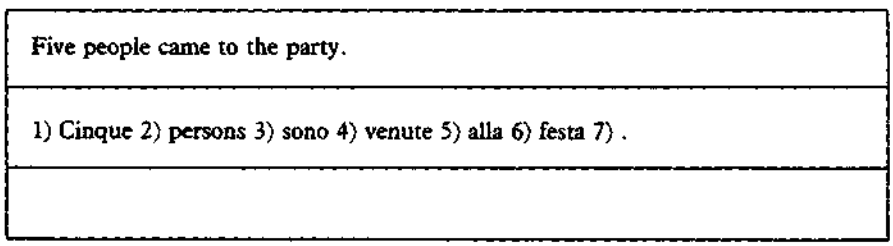

N.B. If *ricevimento* had been chosen:
1) Cinque 2) persons 3) sono 4) venute 5) al 6) ricevimento 7) .

Figure 8. The system displays its own first-draft translation in a second window

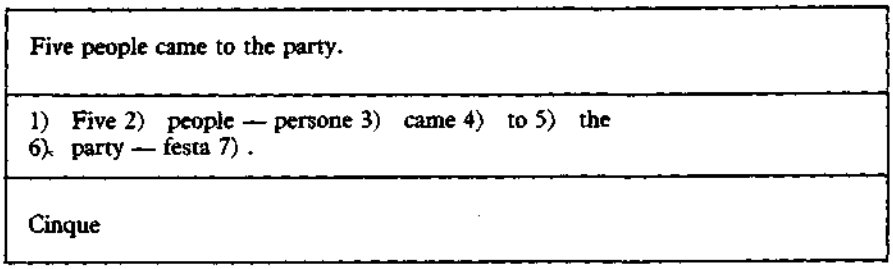

Figure 9. With Autoterm, words are automatically converted to their base forms and looked up in specified dictionaries

built-in multilingual text editor. The bottom window serves as a work area if needed. Otherwise, the user can simply hit a function key and move on to the next segment. Both the source and target texts remain displayed in their respective windows, allowing direct comparison at all times. This is one of the greatest advantages of the sentence-by-sentence mode.

Another important benefit of the sentence-by-sentence mode over the batch mode is the ability to make changes to the document dictionary at any time during the translation process. Whole entries and equivalents can be added, existing inappropriate equivalents can be suppressed and coding errors can be rectified before the problems spread throughout the text. For example, if the translator determines that *festa*, rather than *ricevimento*, will be the appropriate translation for 'party' in the rest of a document, he or she can enter the dictionary and suppress *ricevimento*. In the example just shown, *festa* would then be chosen automatically, without any interaction. As dictionaries are fine-tuned, translators can see immediate improvements in the speed and quality of the output.

This system provides maximum assistance to the translator, semi-automatically providing a first-draft translation of each segment. It is different from most other commercial MAT systems in that it does not produce a raw machine translation fully automatically. Rather, because it is interactive, its output has already been reviewed by a human translator and can be treated as a first-draft human translation.

Transactive systems are currently available for the language pairs: English-French, English-German, English-Spanish, English-Italian and French-English.

Another example of a translation aid is Autoterm. With this program, the user specifies the technical dictionaries to be consulted.

Autoterm divides the source text automatically into segments, which are displayed one-at-a-time in a screen window (see Figure 9). The words are then automatically converted to their base forms, if necessary, and are looked up.

Words and whole phrases found in the dictionary are displayed in a second window with their target equivalents.

The translator then creates the translation in a third window, pulling in any of the target terms he or she needs with a keystroke. In this simple example, the translator has typed *Cinque* and can pull down *persons* by simply typing the tag number '2'. In a real example, dense with technical vocabulary and long multiword phrases, Autoterm can speed up translation considerably.

Autoterm provides a different type of assistance from Transactive in that it simply automates terminology look-up and accelerates typing. The translation itself remains completely under the control of the translator. Translators often appreciate this personal control, and the feelings and insights of translators have all too often been ignored within the whole field of MAT. In addition, everyone in MAT freely admits that many kinds of text are awkward and impractical to translate using automatic or semi-automatic systems — it may be necessary for example, to perform wholesale syntactic revisions of machine-translated documents to achieve natural-sounding translations. Autoterm thus fills a significant but seldom addressed need for automation in the overall translation process.

Autoterm is currently available for the source languages English, French, German, Spanish and Italian. Target dictionaries can be easily constructed for any language using the roman alphabet.

The company recognizes that a powerful, friendly, multilingual text editor is the most basic translator aid of all, and it should form the core of any higher-level translation aid. Multilingual text editor does not mean just any word processor, but one which is designed from the beginning to handle multiple languages naturally, providing familiar keyboard layouts and all necessary punctuation, accents and letters.

Important features include the ability to reset the keyboard (with a simple keystroke), say from the English QWERTY layout to French AZERTY or German QWERTZ, and the ability to type multiple languages in the same file. The translation aids are built around such a multilingual text editor, which can be used not only to edit texts but to compose them.

Although more and more documents are being created on word processors and text editors, and so are already in machine-readable form when handed over for translation, most documents to be translated still arrive typed or printed. Before they can be translated using any automatic or semi-automatic MAT system, they must be converted to a machine-readable form, usually by manual retyping. The cost of this step, which is often underplayed in discussion of MAT, can cancel out the benefits, especially for shorter texts. Using a system such as the ALP Systems multilingual text editor, a translator can consult a hard-copy, source-language document and type the translation directly.

Conclusions

- The translation and document production processes are highly complex. Many translation-related activities are candidates for automation.
- Machine-assisted translation is the approach in which the human translator works *with* the computer (interactive approach).
- The international business community has accepted machine-assisted translation tools as a realistic solution to their multilingual document production problems.

In addition to providing all the characters and punctuation marks required for virtually any European language, the ALP Systems multilingual text editor allows the user to type a word on the screen for automatic look-up in an ALP Systems online dictionary. As in Autoterm, the translation equivalents of the word, if found, are displayed on the screen and can be pulled into the text with a simple keystroke.

Considerations in evaluating automated translation systems

No single system can meet the varied needs of translation automation. Translation is a highly complex process, and many steps are candidates for automation. In evaluating any MAT system or translator aid, one needs to look at the *overall* process, considering *overall* speed and productivity, bottom-line cost-effectiveness, output quality, the type of text to be translated and the audience for the translations.

Even with improvements in translator productivity, it is possible for MAT to cost as much as or more than normal human translation, because MAT carries its own costs. But even in such cases, increased speed of machine translation can still tip the balance in favour of automation. For example, national laws and practical necessity are forcing more and more multinational companies to translate their documentation before marketing their products in foreign countries. The introduction of a new product can be seriously delayed while waiting for the documentation to be translated, and even a slight delay in a competitive market can be financially disastrous.

The most neglected aspect of MAT is perhaps the most important: the needs, attitudes and sensitivities of human translators. MAT researchrs often ignore the fact that human beings have to use their systems. No matter how much automation is introduced into a translation shop, human beings are still going to have to run the machines, and human beings are still going to be responsible for the quality of the translations produced. The cooperation of human translators is therefore essential to making any MAT installation work. If translators are saddled with intimidating and inappropriate technology, or if the machines become their masters rather than their slaves, then the technology will fail — the translators will see to it.

Given that human translators have to work together with MAT systems for most purposes, the design of the man/machine interaction becomes very important. For years, machine translation designers have had the arrogance to tell translators what they need and how they ought to do their work. They have ignored the cries of frustration from translators forced against their will and better judgement to become servants to a machine, burdened with unreasonable amounts of pre-editing, post-editing or interaction.

They have offended translators, who often have an arts background and who distrust computers, by suggesting that the machine can replace them. They have ignored the fact that translators differ in temperament and method, and so will prefer differing amounts of assistance, if any, from a computer.

At various times and with various tasks, some translators will find that typing their own translations with a multilingual text editor is the most automation that they can handle and benefit from. Some will consult a term-bank or machine dictionary more than others. Many will appreciate having technical words looked up automatically and displayed on the screen. Some will appreciate having the computer provide a first draft translation for them to work with.

Conclusion

In conclusion, automated translation is already being used in government agencies, the military, translation houses, and general industry. Dependence on MAT will no doubt explode in the next few years, and several software companies, including ALP Systems, are already moving ahead to claim a portion of the market.

Except in those cases where raw output is acceptable for simple information gathering, MAT always involves people and machines working together. The attitude of the people often determines the success of any installation, and the quality of the human/machine interface is crucial. The most interesting classification of MAT systems, from a user's point of view, is based on how human translators and computers divide the work.

Finally, the cost-effectiveness of a MAT system must be judged in light of the overall translation process, which is highly complex. Many steps in the translation process are candidates for automation, and different kinds of automation (including different levels of assistance) are appropriate for different translation tasks. ALP Systems offers a wide range of translator aids to meet the wide range of translator needs. □

Automated Language Processing Systems Research and Development, 190 West 800 North, Provo, UT 84604, USA.