Machine Translation vs. Translator Aids: A False Dichotomy

Alan K. Melby
Brigham Young University

BACKGROUND ON COMPUTERS AND TRANSLATION

Advances in Computers

The past 30 years have been a period of rapid development for computers and their uses. In computer hardware, we have seen a progression from vacuum tubes to transistors to integrated circuits to microprocessors. In computer uses, we have seen an expansion beyond military and business applications that includes computer-aided design and manufacturing, bibliographic data retrieval, concordance making, education, and, of course, games. Where does translation fit into this picture?

Machine Translation: An Old and Unsuccessful Effort

Actually, the use of computers for translation is one of the oldest and least successful computer applications. It is one of the oldest because it was suggested by Warren Weaver over 30 years ago in a well-known memo (see reprint of memo in Weaver 1955). The reasons for its lack of success are complex, but I suggest that at least some of the failures are a result of unfortunate assumptions. Early work in computers and translation assumed that the only acceptable goal was fully automatic high-quality translation and, with equal misfortune, assumed that any interaction between the computer and the human translator was temporary and undesirable. These assumptions are a part of the machine translation tradition which continues today.

Translator Aids: An Opposing Approach

The translator aid tradition arose in reaction to these assumptions. This new tradition has seen considerably more success. Translator aids include terminology access, document access, and word processing, and they are in daily use, particularly in Canada and Europe. An important factor in their success is, of course, the fact that translator aids do not translate. They are strictly tools to be used at the discretion of human translators.

Until recently, there was no middle ground between machine translation and translator aids. Translator aid projects have been designed to help human translators

become more productive and effective, while machine translation projects have been designed to eliminate human translators. Designs so diametrically opposed in goals are, surprisingly, not incompatible. But is this dichotomy necessary?

Recent Motions Toward Integration

In this paper I will argue that the dichotomy is unnecessary and that it can be dissolved if one adopts a modified viewpoint of machine translation and incorporates it into a system of translator aids.

In the past few years, some machine translation systems have been marketed as machine-assisted translation systems. The computer does the translation, perhaps asking questions of the human along the way, and then presents a rough translation to the human translator for revision. For nearly a decade, I was a member of a team working on one of these machine-assisted translation systems. These systems do not replace human translators, but rather increase their productivity. So are they really translator aids? Not yet. This is because of the way the translator perceives the machine. The translator is told to postedit and, perhaps, to answer questions. The human translator often feels like a tool for the machine. Happily, however, the scene is changing. Some commercial teams which market machine translation systems are now also offering new tools for translators. It is important to translators whether they are the tool or the machine is the tool. When machine translation is viewed as one of many tools for a human translator to choose from, the dichotomy between machine translation and translator aids disappears.

INTECRATING TRANSLATOR AIDS AND MACHINE TRANSLATION

A Success Story: The METEO System

An example of a translation system involving machine translation yet breaking out of the old tradition is the METEO system developed by the University of Montreal. This system daily translates many thousands of words of Canadian weather fore-:asts from English to French (Chandioux 1978). Until it was installed, the Canadian government had considerable difficulty retaining weather forecast translators because of the boring nature of their work. Now, a computer program identifies the straightforward sentences (about 80%) and translates them automatically. In the remaining sentences, the computer system detected some problem, such as an unknown word or grammatical construction which would probably cause a flaw in he machine output. These sentences are presented to the human for translation, Note that the human is not asked to postedit the machine translation of these sentences. The human docs not even sec the French translation, but rather translates the original English. From the translator's point of view, the machine is a tool for filtering out the boring sentences. The translators are relieved of much drudgery and are significantly happier in their work.

In contrast, consider a traditional machine translation system in which the transator is asked to postedit the raw machine translation of an entire text, compared

with the METEO system, in which the human translates those sentences which the computer cannot handle. To the translator, the difference is dramatic. In the traditional system, the translator is a 'garbage collector', removing the botched up words along the sentences of the machine translation. In the METEO system, however, the translator is an executive, handling the difficult cases that are beyond the capacity of the limited intellect of the computer. The success of the METEO system shows that machine translation can be effective when used in the proper environment.

The BYU Translator Aid System

The remainder of this paper will describe a system design which integrates word processing, terminology access, and machine translation. The system is a tool for the translator, yet includes a machine translation component which, hopefully, extends some principles found in the METEO system to a more general environment. The system is being developed at Brigham Young University (BYU) as an academic research project. It is called the BYU Translator Aid System. The design is the result of years of discussion with colleagues. Special credit for inspiration should be given to Martin Kay (see Kay 1980).

The system provides three integrated levels of aid to the translator. The software is being implemented on a single-user microcomputer consisting of a processor, video screen with 24 lines of 80 characters, keyboard, two diskette drives, and printer. To go beyond a demonstration version to a practical system, it is necessary to add a hard disk with a storage capacity of at least 5,000,000 characters. At level one, the translator can sit down with a text in the source language and begin translating immediately. To use level two, the source text must be available in a segmented form on diskette or over a network. A segment is a piece of text of up to about 100 words. It may be one long sentence or several short sentences. The source text may be available in machine readable form, as a by-product of publication in the source language, or it may be obtained using optical scanners. To use level three, the source text must be processed off-line by a machine translation system which keeps track of problems encountered. The machine translation must be segmented to correspond to the segments of the source text and each problem encountered must be associated with some segment.

If several translators are working together, they should share disk files by means of a network. They should not be multiple users of a single processor, as processors are inexpensive while people are not. Multiple processors in a network provide consistent response time to each user for local tasks, such as word processing, regardless of what oilier users are doing.

The Operation of the System

We will explore the operation of the system at the various levels. At each point in the translation process, the screen is divided into windows. For example, the first line on the screen might be a title, lines 3 through 11 might form the upper window, and lines 12 through 24 might form the lower window. At another point the screen may be divided up into three windows. A special case is the final revision step.

After the translator aid system has been used to produce a draft translation and the translator is satisfied with the content of the translation, the text file is converted to a standard word processing file and the translator uses a standard word processor to make final revisions and to adjust the format as desired. In the final revision step, the whole screen consists of a single window and the text is no longer segmented. Thus, the translator produces a translation in two steps: first, the draft translation step, using one or more levels of aid and switching freely among available levels, and second, the final revision step, using a standard word processor. We will now look at the three levels of aid in the draft translation step.

Level One Word Processing. At level one, the system can look much like a word processor. The screen consists of two major windows. The bottom window (consisting of 10 lines) is the translator's work area. Nothing goes into the bottom window except what the translator puts there. At some point, the translator decides to move on to a new segment by pressing the appropriate function key. The segment moves into the upper window and the bottom window is filled with the next segment or, if there are no more segments, the window is blank and ready for a new segment of translation. The translator can freely move forwards and backwards through the text by segment, or move to the top or bottom of the file or to a specific segment. Editing can always be done on the segment of text which is in the bottom window. To edit something in the upper window, the translator simply scrolls down a segment at a time until the text in question is in the bottom window, thus becoming the current or active segment. A standard word processor scrolls by line or by screen. The one being discussed scrolls by segment. While a standard word processor allows changes to be made to the text anywhere on the screen, this one allows changes to be made in the current segment, which is in a window consisting of approximately the bottom half of the screen. I have worked with standard word processors and with this unconventional one and do not find it difficult to adjust from one to the other. The reason for this segment orientation in the draft translation step is to make the three levels of aid look very much the same to the translator. At level one, the source text is on paper and the translator places it next the computer display. At levels two and three, the source text appears on the screen one segment at a time, and the screen is divided into three major windows: the current segment of source text, aids corresponding to that segment, and the current segment of target text. Inside the computer, each segment of target text is linked back to its corresponding source segment. This means that as the translator scrolls back and forth through his or her translation, the appropriate source segment can be automatically displayed in the upper window. If the text were not segmented, the translator would have to scroll separately through the source and target texts. At level one, since the source text is not on disk, the system would not have to work with segments, but this would be at the cost of decreasing the similarity among the three levels. I have attached a very high priority to making the three levels look very much alike to the translator. It may take years of experimentation before it becomes clear whether this was a good decision.

Level One Terminology Access. A terminology access option is still available at level one. At the touch of a function key, the upper portion of the screen becomes a window into a bilingual terminology file. The translator enters the term to be looked up and a bilingual dictionary entry for that term is retrieved from disk and placed in the upper window. If there is more than one window of information on the term, the translator can move forward or backward through the information associated with that term. If the term is not found, the term can be added. Existing entries can also be modified. A number of terminology banks already exist, such as those of the Canadian government and the European Community, and these term files can be accessed on video terminals. The translator aid system would provide the ability to access these remote files, but the primary access would be to its own local term files. These local term files could be built up item by item by the translator as new translations are identified, or they could be purchased from or traded with other translators or from term banks. The local storage option provides much faster lookup and avoids long distance telephone charges. Local storage also provides the possibility of local printouts of glossaries or subsets of glossaries.

An important feature of this terminology access is that the translator need not leave the word processor to access the terms. In fact, the current segment of translation is visible during the entire terminology access procedure, and a simple keystroke returns the system to word processing mode with the cursor exactly where it was before the term was looked up. Upon returning to word processing mode, the dictionary entry is still visible in the upper window and the translator can pull down the desired translation and insert it into the current segment of translation with two or three keystrokes.

A special format, monolingual term file is also available at level one. This file, called the expansion code file, allows the translator to define abbreviations which are automatically expanded at the touch of a special function key. This feature is useful whenever there are words or phrases which occur repeatedly, such as names of parts in a machine or titles of organizations. The translator decides which terms occur frequently enough to warrant an expansion code and devises a short code for those terms. For example, the code UI might be defined by the translator to be expanded as USER INTERFACE. The term is inserted into the translation by typing UI and then pressing the expansion key. This feature is already available in a number of word processing systems and has been shown to be useful. However, two cautions are in order. The first is that the access must be very fast to be useful. The delay between the instant the translator presses the expansion key and the point where the system has expanded the code so that the translator can go on typing text should be well under a second. The second caution is against eliminating the expansion key. Some systems automatically check for expansion codes every time the space bar is pressed, potentially causing surprises for the translator. As a rule, nothing should happen in the current segment window unless the translator specifically requests it. This policy is consistent with the effort to design a system in which the translator feels in control, is treated with respect, and is provided with an array of tools from which to choose.

Other Options at Level One. Other options which should be available at level one are communication with other translators and retrieval of relevant documents. Communication with other translators may be accomplished via electronic mail or electronic bulletin boards. Retrieval of relevant documents could be provided by access to data bases of translated materials or target language publications. If a new term is not in the printed dictionaries and not yet in the term banks, the translator may have to do some basic research. One approach is to find a recently translated document containing the term in question. Another approach is to find a recent article in the target language dealing with the topic at hand. By studying such articles, the translator may determine by context that a certain term is the translation for which it is looking.

Level Two Suggested Terms. At level two, the source text to be translated is available in segmented, machine readable form, which adds a new type of screen to the word processing, terminology access, and communications screens already available at level one. This new screen consists of three windows. The top window contains a segment of source text to be translated. The middle window contains suggested translations for highlighted terms in the segment. The bottom window is the translator's work area for the current segment, as in the other level. The system finds suggested translations of terms (words or phrases) simply by splitting the current source segment into words, finding the root of each, and looking them up in the local term file. When several entries are found to begin with the same word, the system chooses the term which matches the longest stretch of words in the source segment. For example, if the source segment contains the phrase *fountain pen* and there are two entries in the term file, one for *fountain* and another for *fountain pen*, the system will choose *fountain pen* because it is a longer match. Those terms in the source segment which are found in the term file are highlighted and numbered. The translations found in the term file are placed in the middle window with numbers corresponding to the numbers in the source text. The translator is free to use the suggestions or to ignore them. To use a suggestion the translator simply types the number of the suggestion and presses a function key which pulls down the suggestion and inserts it into the current segment.

Suppose that the current segment of source text contains the French expressions *grand public* and *démarcheurs à domicile* and that the system correctly suggests the translations *general public* and *door-to-door salesmen.* The translator might type in something such as: *Sales to the 1$ may be done through 2$ or by mail.* The dollar signs indicate that some special key was pressed at that point. Pressing that key tells the computer to use the preceding number as a suggestion number. The suggestion with that number is instantly retrieved from the suggestion window and inserted into the text at the cursor position. Thus, two or three keystrokes suffice to use a suggestion, no matter how long the suggestion might be. As the translator marks additional entries in the term files, the number of suggestions that appear per segment increase. Its usefulness is also influenced by the system's effectiveness at finding root forms and reconstructing appropriate inflected forms. Of course, if a

suggestion is often found to be unusable, it can be modified or the entry can be marked so that it does not appear automatically. It is important to note here that nothing is imposed on the translator. Suggested translations of terms will never appear in the target text unless the translator explicitly requests that they be inserted. The translator may also choose to switch back to a level one screen at the touch of a function key.

Level two is intended to be more convenient to the translator than level one. At level two, the translator has the current segment of source text on the screen along with suggested translations of terms, and as the translator scrolls back and forth through the translation, the appropriate segment of source text automatically appears in the upper window. As mentioned previously, this avoids having to scroll separately through the source and target texts and is the main motivation for using segments throughout the system. Another benefit of level two is that it can facilitate communication among several translators working on the same project. Given that a term file with suggestions is shared by several translators, they will automatically be made aware of standard translations that have been established for technical terms. This could help the team produce translations with more consistency in terminology. Consistency of usage in technical documents is important as lack of it can lead to confused readers. Moreover, another benefit of level two is that having the source text in machine readable form allows preprocessing of the source text, including a spelling check and a check for repeated phrases that may need standard translations.

Level Three Suggestion Translations of Larger Units. Level three does not take anything away from levels one and two, rather it does make one huge addition: it includes a modified machine translation system. Eventually, the machine translation component may be completely incorporated into a high-powered translator work station. But for the moment, I anticipate a link with an existing off-line machine translation system. I have a tentative agreement with the University of Grenoble (France) to integrate their machine translation system, so I will use it as an example. Their system runs on an IBM mainframe, accepts a source text, divides it into segments in approximately the way prescribed for level two, and translates the text a segment at a time. To be used in the BYU Translator Aid System, the Grenoble system will have to be modified to keep track of problems encountered during translation, such as missing words in the dictionary, unknown syntactic constructions, and questionable translation choices. The output of the system must be a series of segment clusters. In each cluster will be the source segment, the machine translation of it, and the problem record for it. This information is presented to the translator work station on diskette or over a network, along with the segmented source text. At level two, only the segmented source text is presented to the work station. At level three, each segment of source text is accompanied by a problem record and perhaps a machine translation of the segment. The translator work station is always set to a particular tolerance level. When a given source segment is displayed, the translator work station checks the problem record in the

segment cluster against the translator's tolerance level. If the problem severity does not exceed the tolerance level, then the machine translation for the segment will be displayed in the middle window of the screen. The translator can examine it and, if it is acceptable as is or with minor changes, can pull it down into the bottom window (the work area) and edit it. Or, if the machine translation is not acceptable, the translator can ignore it or switch back to a level two suggested term screen or to a level one screen. If a number of translated segments are unacceptable, the translator may choose to lower the tolerance level, thus blocking more segments. On some documents, the machine translation may be so poor that the translator finds it more efficient to turn off level three entirely. The translator can even turn off level two and proceed at level one with word processing and only those dictionary entries explicitly requested.

It should be obvious at this point that the design takes pains to ensure that the translator is in control, even at level three. The role of the translator is neither temporary nor undesirable. The translator is viewed as a competent professional who can be trusted to use the level of aids most appropriate for a given text and segment within a text. No piece of machine translation is ever forced upon the translator and cannot appear in the target text unless the translator requests it to be there. It should also be clear that to the translator there is little difference between the levels, despite the radical underlying differences. The translator can freely switch among various windows in the upper portion of the screen: previous text, terminology access, communications, suggested terms, and suggested segment translations. To the translator, all the screens are part of one integrated system. It just happens that on some documents, level two and level three aids are not available, and on other documents, level two aids are available and level three aids are available only on some segments. This integration of levels, retaining respect for the translator, resolves the conflict between machine translation and translator aids. Stripped of the assumption that every segment must be translated by machine, machine translation becomes more friendly and more *useable.* The idea may be simple, but it has come into being only recently after nearly 30 years of hitting heads against fully automatic walls.

A possible objection to the approach proposed in this paper is that it limits work in machine translation by assuming that it will not work. Actually, it should encourage work in machine translation by allowing it to be useful, even if it is not of consistently high quality. If the machine translation component produces a perfectly acceptable translation for each segment of a text, the translator can simply pull down each segment of machine translation and pass it through as fast as the machine translation can be read. On the other hand, if the machine translation is somewhat less than perfect, the good segments can be pulled down and the translation process still proceeds smoothly. No limit is imposed on the sophistication of the machine translation component.

A final note is in order to explain the status of the BYU Translator Aid System. Last year we implemented a limited version of level two. We also have a separate version of a portion of level one. We will soon begin writing a new version on a 16-

bit microcomputer. This new version will include level one and two and at least a simulation of level three. True implementation of level three depends on the progress of negotiations to obtain access to a machine translation component. Undoubtedly, there will be surprises and design changes during the implementation and testing of the system. Our hope is to produce a *useable* demonstration system in the next few years or to spread ideas to the right people so that someone else will develop the system. In either case, I am very hopeful that the dignity and professional status of translators will be increased by treating machine translation as a tool, rather than as a replacement for translators, thus dissolving a dichotomy which is false because it is unnecessary and unproductive.

REFERENCES

ANDREYEWSKI, ALEXANDER. 1981. Translation: Aids, robots, and automation. META Translators' Journal 26:1.57-66.

CHANDIOUX, JOHN. 1978. METEO. TAUM: The University of Montreal.

KAY, MARTIN. 1980. The proper place of men and machines in language translation. Palo Alto, CA: Xerox Corporation.

MELBY, ALAN K. 1982. Multi-level translation aids in a distributed system. COLING82, ed. by Jan Horecký, 215-220. Amsterdam: North Holland.

WEAVER, WARREN. 1955. Translation. Machine translation of languages, ed. by W. N. Locke and A. D. Booth, 15-23. New York: Wiley.