

COMPUTER-ASSISTED TRANSLATION SYSTEMS:
The Standard Design and A Multi-level Design

Alan K. Melby
Linguistics Department
Brigham Young University
Provo, Utah 84602 USA

ABSTRACT

The standard design for a computer-assisted translation system consists of data entry of source text, machine translation, and revision of raw machine translation. This paper discusses this standard design and presents an alternative multi-level design consisting of integrated word processing, terminology aids, preprocessing aids and a link to an off-line machine translation system. Advantages of the new design are discussed.

I THE STANDARD DESIGN FOR A COMPUTER-ASSISTED TRANSLATION SYSTEM.

The standard design for a computer-assisted translation system consists of three phases: (A) data entry of the source text, (B) machine translation of the text, and (C) human revision of the raw machine translation. Most machine translation projects of the past thirty years have used this design without questioning its validity, yet it may not be optimal. This section will discuss this design and some possible objections to it.

The data entry phase may be trivial if the source text is available in machine-readable form already or can be optically scanned, or it may involve considerable overhead if the text must be entered on a keyboard and proofread.

The actual machine translation is usually of the whole text. That is, the system is generally designed to produce some output for each sentence of the source text. Of course, some sentences will not receive a full analysis and so there will be a considerable variation in the quality of the output from sentence to sentence. Also, there may be several possible translations for a given word within the same grammatical category and subject matter so that the system must choose one of the translations arbitrarily. That choice may of course be appropriate or inappropriate. It is well-known that for these and other reasons, a machine translation of a whole text is usually of rather uneven quality. There is an alternative to translating the whole text -- namely, "selective translation," a notion which will be discussed further later on.

Revision of the raw machine translation by a human translator seems at first to be an attractive way to compensate for whatever errors may occur in the raw machine translation. However, revision is effective only if the raw translation is already nearly acceptable. Brinkmann (1980) concluded that even if only 20% of the text needs revision, it is better to translate from scratch instead of revising.

The author worked on a system with this standard design for a whole decade (from 1970 to 1980). This design can, of course, work very well. The author's major objection to this design is that it must be almost perfect or it is nearly useless. In other words, the system does not become progressively more useful as the output improves from being 50% correct to 60% to 70% to 80% to 90%. Instead, the system is nearly useless as the output improves and passes some threshold of quality. Then, all of a sudden, the system becomes very useful. It would, of course, be preferable to work with a design which allows the system to become progressively more useful.

Here is a summary of objections to the standard design:

WHY COMPUTATIONAL LINGUISTS DO NOT LIKE IT:
Because even if the algorithms start out "clean", they must be kludged to make sure that something comes out for every sentence that goes in.

WHY TRANSLATORS DO NOT LIKE IT:
Because they feel that they are tools of the system instead of artists using a tool.

WHY SPONSORS DO NOT LIKE IT:
Because the system has to be worked on for a long time and be almost perfect before it can be determined whether or not any useful result will be obtained.

II AN ALTERNATIVE DESIGN

There has been for some time a real alternative to the standard design -- namely, translator aids. These translator aids have been principally terminology aids of various kinds and some use of standard word processing. These aids have been found to be clearly useful. However, they have not attracted the attention of computational linguists because they do not involve any really interesting or challenging linguistic processing. This is not to say that they are

trivial. It is, in fact, quite difficult to perfect a reliable, user-friendly word processor or a secure, easy to use automated dictionary. But the challenge is more in the area of computer science and engineering than in computational linguistics.

Until now, there has not been much real integration of work in machine translation and translator aids. This paper is a proposal for a system design which allows just such an integration. The proposed system consists of two pieces of hardware: (1) a translator work station (probably a single-user micro-computer) and (2) a "selective" machine translation system (probably running on a mainframe). The translator work station is a three-level system of aids. All three levels look much the same to the translator. At each level, the translator works at a keyboard and video display. The display is divided into two major windows. The bottom window contains the current segment of translated text. It is a work area, and nothing goes in it except what the translator puts there. The upper window contains various aids such as dictionary entries segments of source text, or suggested translations.

To the translator, the difference between the various levels is simply the nature of the aids that appear in the upper window; and the translator in all cases produces the translation a segment at a time in the lower window. Internally, however, the three levels are vastly different.

Level 1 is the lowest level of aid to the translator. At this level, there is no need for data entry of the source text. The translator can sit down with a source text on paper and begin translating immediately. The system at this level includes word processing of the target text, access to a terminology file, and access to an expansion code file to speed up use of commonly encountered terms.

Level 2 is an intermediate level at which the source text must be available in machine readable form. It can be entered remotely and supplied to the translator (e.g. on a diskette) or it can be entered at the translator work station. Level 2 provides all the aids available at level 1 and two additional aids -- (a) preprocessing of the source text to search for unusual or misspelled terms, etc., and (b) dynamic processing of the source text as it is translated. The translator sees in the upper window the current segment of text to be translated and suggested translations of selected words and phrases found by automatically identifying the words of the current segment of source text and looking them up in the bilingual dictionary that can be accessed manually in level 1.

Level 3 requires a separate machine translation system and an interface to it. Instead of supplying just the source text to the translator work station, the work station receives (on diskette or through a network) the source text and (for each segment of source text) either a machine

translation of the segment or an indication of the reason for failure of the machine translation system on that segment. This explains the notion of "selective" machine translation referred to previously. A selective machine translation system does not attempt to translate every segment of text. It contains a formal model of language which may or may not accept a given segment of source text. If a given segment fails in analysis, transfer, or generation, a reason is given. If no failure occurs, a machine translation of that segment is produced and a problem record is attached to the segment indicating difficulties encountered, such as arbitrary choices made. Level 3 provides to the translator all the aids of levels 1 & 2. In addition, the translator has the option of specifying a maximum acceptable problem level. When a segment of source text is displayed, if the machine translation of that segment has a problem level which is low enough, the machine translation of that segment will be displayed below the source text instead of the level 2 suggestions. The translator can examine the machine translation of a given segment and, if it is judged to be good enough by the translator, the translator can pull it down into the bottom window with a single keystroke and revise it as needed. Note that writing a selective machine translation system need not mean starting from scratch. It should be possible to take any existing machine translation system and modify it to be a selective translation system. Note that the translator work station can provide valuable feedback to the machine translation development team by recording which segments of machine translation were seen by the translator and whether they were used and if so how revised.

The standard design for a machine translation system and the alternative multi-level design just described use essentially the same components. They both involve data entry of the source text (although the data entry is needed only at levels 2 and 3 in the multi-level design). They both involve machine translation (although the machine translation is needed only at level 3 in the multi-level design). And they both involve interaction with a human translator. In the standard design, this interaction consists of human revision of the raw machine translation. In the multi-level design, this interaction consists of human translation in which the human uses word processing, terminology lookup, and suggested translations from the computer. At one extreme (level 1), the multi-level system involves no machine translation at all, and the system is little more than an integrated word processor and terminology file. At the other extreme (level 3), the multi-level system could act much the same as the standard design. If every sentence of the source text received a machine translation with a high quality estimate, then the translation could conceivably be produced by the translator choosing to pull each segment of translated text into the translation work area and revise it as needed. The difference between the two designs becomes apparent only when the raw machine translation is not almost perfect. In that case, which is of course common, the multi-level system continues

to produce translations with the human translator translating more segments using level 1 and level 2 aids instead of level 3 aids; the translation process continues with some loss of speed but no major difficulty. When the same raw machine translation is placed in a standard design context, the translator is expected to revise it in spite of the problems, and according to the author's experience, the translators tend to become frustrated and unhappy with their work. Both designs use the same components but put them together differently. See Figure 1.

Here is a summary of the arguments for a multi-level design:

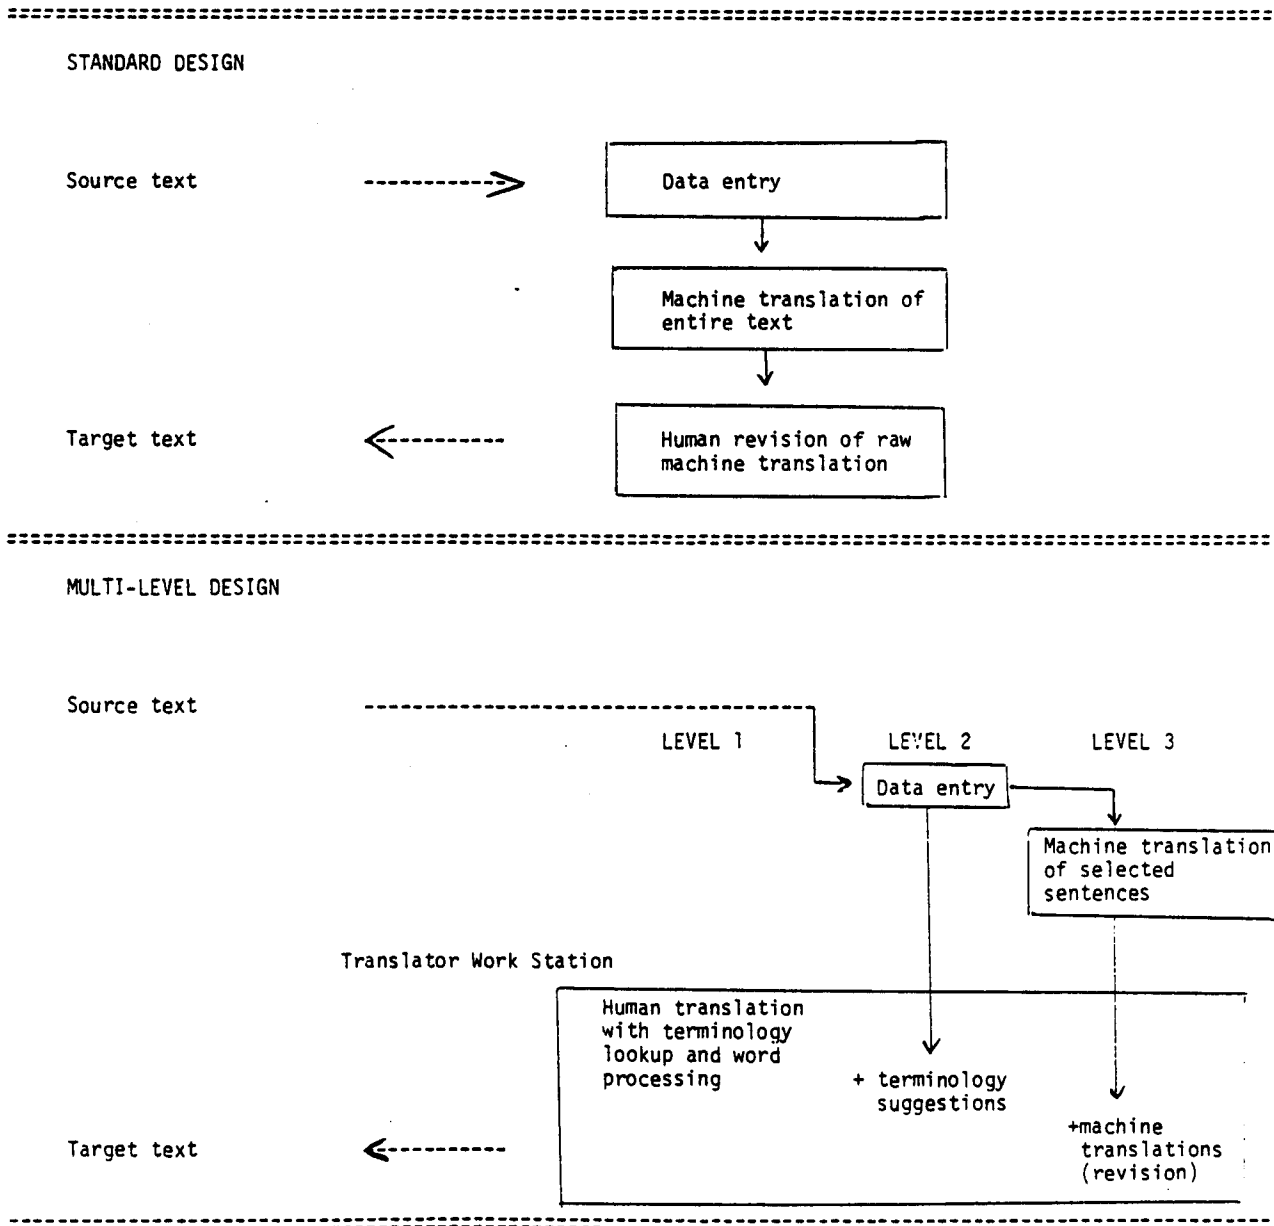
WHY COMPUTATIONAL LINGUISTS LIKE IT:

Because they can set up a "clean" formal model and keep it clean, because there is no pressure to produce a translation for every sentence that goes in.

WHY TRANSLATORS LIKE IT:

Because the system is truly a tool for the translator. The translator is never pressured to

Figure 1 Two Designs



revise the machine output. Of course, if the raw machine translation of a sentence is very good and needs only a minor change or two, the translator will naturally pull it down and revise it because that is so much faster and easier than translating from scratch.

WHY SPONSORS LIKE IT:

Because the system is useful after a modest investment in level 1. Then level 2 is added and the system becomes more useful. While the system is being used at levels 1 and 2, level 3 is developed and the machine translation system becomes a useful component of the multi-level system when only a small fraction of the source sentences receive a good machine translation. Thus, there is a measurable result obtained from each increment of investment.

III IMPLEMENTATION EXPERIENCE AND PLANS

The multi-level design grew out of a Naval Research Laboratory workshop the summer of 1981, a paper on translator aids by Martin Kay (1980); and user reaction to a translator aid system (called a "Suggestion Box" aid) was tested on a seminar of translators fall 1981. The current implementation is on a Z-80 based micro-computer. The next implementation will be on a 16-bit micro-computer with foreign language display capabilities.

The author is now looking for a research machine translation system to use in level 3, e.g. ARIANE-78 (See Boitet 1982). Further papers will discuss the successes and disappointments of a multi-level translation system.

IV REFERENCES

ANDREYEWski, ALEXANDER. 1981. Translation: Aids, Robots, and Automation. META Translators' Journal 26:1.57-66.

BOITET, Ch., Guillaume, and M. Quezel-Ambrunaz. 1982. Implementation and Conversational Environment, of ARIANE 78.4. COLING82, ed. by Jan Horecky, 19-28. Amsterdam: North Holland.

BRINKMANN, Karl-Heinz. 1980. Terminology Data Banks as a Basis for High-Quality Translation. COLING80, Proceedings. Tokyo.

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

KAY, MARTIN. 1980. The Proper Place of Men and Machines in Language Translation, Palo Alto Research Center: Xerox Corporation.

MELBY, ALAN K. 1982. Multi-level Translation aids in a Distributed System, COLING82, ed. by Jan Horecky, 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