Computers and the Humanities 19 (1985)
S/Paradigm Press, Inc.

Machine Translation

Jonathan Slocum

Machine Translation (MT) of natural human languages is not a subject about which most scholars feel neutral. This field has had a long, colorful career, and boasts no shortage of vociferous detractors and proponents alike. During its first decade in the 1950s, interest and support were fueled by visions of high-speed high-quality translation of arbitrary texts (especially those of interest to the military and intelligence communities who funded MT projects). During its second decade in the 1960s, disillusionment crept in as the number and difficulty of the linguistic problems became increasingly obvious, and as it was realized that the translation problem was not nearly so amenable to automated solution as had been thought. The climax came with the delivery of the National Academy of Sciences ALPAC report (1966), condemning the field and, indirectly, its workers alike. The ALPAC report was criticized as narrow, biased, and shortsighted, but its recommendations were adopted (except that expenditures for long-term research in computational linguistics were NOT increased, but instead decreased), and as a result most MT projects in the U.S. were cancelled. By 1973, the early part of the third decade of MT, only three government-funded projects were left in the United States, and by late 1975 there were none. (The University of Texas METAL project was subsequently funded, from 1977 to 1980). Paradoxically, MT systems were still being used by various government agencies here and abroad, because there was simply no alternative means of quickly gathering information from foreign (Russian) sources; in addition, private companies were developing and selling MT systems based on the mid-60s technology so roundly castigated by ALPAC. Nevertheless the general disrepute of MT resulted in a remarkably quiet third decade.

Jonathan Slocum is Director of the Natural Language Processing Project in the Artificial Intelligence Program at Microelectronics and Computer Technology Corporation in Austin, Texas.

We are now into the fourth decade of MT, and there is not only a resurgence of interest throughout the world but also a growing number of MT and MAT (Machine-aided Translation) systems in use by governments, business and industry. With industrial firms also beginning to fund M(A)T R&D projects of their own, it can no longer be said that only government funding keeps the field alive (indeed, in the United States there is no government funding, though the Japanese and European governments are heavily subsidizing MT R&D). In part this interest is due to more realistic expectations of what is possible in MT, and realization that MT can be very useful though imperfect, but it is also true that the capabilities of the newer MT systems lie well beyond what was possible just one decade ago.

In light of these events, it is worth reconsidering the potential of, and prospects for, Machine Translation. After opening with an explanation of how (human) translation is done where it is taken seriously (in order to clarify which arguments about MT are relevant and which are not), we will present a brief introduction to MT technology, then consider the state of the art, and finally discuss future prospects.

The Human Translation Context

When evaluating the feasibility or desirability of Machine Translation, one should consider the endeavor in light of the facts of human translation for like purposes. In the United States, it is common to conceive of translation as simply that which a human translator does. It is generally believed that a college major (or the equivalent) in a foreign language qualifies one to be a translator, for just about any material whatsoever. Not surprisingly, native speakers of foreign languages are considered to be that much more qualified; Worse, fluency in one foreign language is sometimes taken to be, ipso facto, proof of competence in related languages; an embarrassing incident involving a former U.S. President in Poland illustrates the folly of this position, yet it is still widely held. Thus, translation is not particularly respected as a profession in the U.S. (it is frequently perceived as the "last resort" of otherwise-unemployed liberal arts majors), and the pay is poor.

Elsewhere around the world, this myopic attitude is not held. Where translation is a fact of life rather than an oddity, it is realized that one translator does not equal another, that any technical translator's competence is sharply restricted to a few domains, that native fluency in a foreign language does not bestow on one the ability to serve as a translator, and that languages really are different and deserve treatment as such. Thus, in college-level and post-graduate schools that teach the theory (translatology) as well as the practice of translation, a technical translator is trained in the few technical areas in which he will be doing translation. True competence is expected only after years of experience beyond that training, and even then only in a few limited areas of translation expertise.

Of special relevance to MT is the tact that essentially all translations for dissemination (export) are revised by more highly qualified translators who necessarily refer back to the original text when post-editing the translation. (This is not like the "prepublication stylistic editing" that, for example, writers in the United States are familiar with.) Unrevised translations are always regarded as being inferior in quality, or at least suspect, and for many if not most purposes they are simply not acceptable. In the multi-national firm Siemens, even internal communications which are translated are also post-edited. Such news generally comes as a surprise, if not a shock, to most people in the United States.

It is easy to see, therefore, that the "fully-automated high-quality machine translation" standard, imagined by most United States scholars to constitute minimum acceptability, must be radically defined. Indeed, the most famous MT critic of all eventually recanted his strong opposition to MT, admitting that these terms could only be defined by the users, according to their own standards, for each situation (Bar-Hillel, 1971). So an MT system does not have to read a book, then print and bind the result of its translation, in order to qualify as "fully automatic." "High quality" does not at all rule out post-editing, since the proscription of human revision would "prove" the infeasibility of high-quality Human Translation. In practice, academic debates about what constitutes "high-quality"and "fully-automatic" are considered irrelevant by users of Machine Translation (MT) and Machine-aided Translation (MAT) systems; what matters to them are two things: whether the systems can produce output of sufficient quality for the intended use (e.g., revision), and whether the operation as a whole is cost-effective or, rarely, justifiable on other grounds, such as speed.

Machine Translation Technology

In order to appreciate the technologies behind machine translation systems, it is necessary to understand, first, the broad categories into which they can be classified; second the different purposes for which translations (however produced) are used; third, the intended applications of these systems; and fourth, something about the linguistic techniques which MT systems employ in attacking the translation problem.

Categories of Systems

There are three broad categories of "computerized translation tools" (the differences hinging on how ambitious the system is intended to be): Machine Translation (MT), Machine-Aided Translation (MAT), and Terminology Databanks (TD).

MT systems are intended to perform translation without human intervention. This restriction does not rule out pre-processing (except for such decisions as marking phrase boundaries and resolving part-of-speech and/or other ambiguities), nor post-editing (since this is normally done for human translations anyway). However, an MT system is solely responsible, without human assistance, for the complete translation process from input of the source text to output of the target text, using special programs, comprehensive dictionaries, and collections of linguistic rules (to the extent that they exist, varying with the MT system). On the scale of computer translation sophistication, MT occupies—the top position.

MAT systems fall into two subgroups: human-assisted machine translation (HAMT)and machine-assisted human translation (MAHT). These occupy successively lower ranges on the scale of computer translation sophistication. In a HAMT system the computer is responsible for producing the translation per se, but may interact with a human monitor at many stages along the way, as by asking the human to disambiguate a word's part of speech or meaning, or to indicate where to attach a phrase, or to choose a translation for a word or phrase from among several candidates discovered in the system's dictionary. In a MAHT system, on the other hand, the human is responsible for producing the translation per se (on-line), but may interact with the system in certain prescribed situations, for example, requesting assistance in searching through a local dictionary/thesaurus, accessing a remote terminology databank, retrieving examples of the use of a word or phrase, or performing word processing functions like formatting. The existence of a pre-processing stage is unlikely in an MA(H)T system (the system does not need help, rather it is making help available), but post-editing is frequently appropriate.

Terminology Databanks (TD) are the least sophisticated systems because access frequently is not made during a translation task (the translation may not be working on-line), but usually is performed prior to human translation. Indeed, to the translator, the databank may not be accessible on-line at all, but may be limited to the production of printed subject-area glossaries. A TD offers access to technical terminology, but usually not to common words (the user already knows these). The chief advantage of a TD is not that it is automated (even with on-line access, words can be found just as quickly in a printed dictionary), but that it is up-to-date: technical terminology is constantly changing and published dictionaries are essentially obsolete by the time they are available. It is also possible for a TD to contain more entries because it can draw on a larger group of active contributors, its users.

The Purposes of Translation

The most immediate division of translation purposes involves information acquisition vs. dissemination. The classic example of the former purpose is intelligence-gathering: with masses of data to sift through, there is no time, money, or incentive to carefully translate every document by normal (i.e., human) means. Scientists more generally are faced with this dilemma: already more must be read than can be read in the time available. Having to labor through texts written in foreign languages—when the probability is low that any given text is of real interest—may not be worth the effort. In the past, English has been the lingua franca of science, but it is becoming less and less dominant for a variety of reasons, including the rise of nationalism and the spread of technology around the world. As a result, scientists who rely on English are having greater difficulty keeping up with work in their fields. If a very rapid and inexpensive means of translation were available, then—for texts within the reader's areas of expertise—even a low-quality translation might be sufficient for information acquisition. At worst, the reader could determine whether a more careful (and more expensive) translation effort might be justified. More likely, his understanding of the text would be good enough to preclude the need for a more careful translation.

The classic example of the latter purpose of translation is technology export: an industry in one country that desires to sell its products in another country must usually provide documentation in the purchaser's chosen language. In the past, United States companies have escaped this responsibility by requiring that the purchasers learn English; other exporters (German, for example) have never had this luxury. With the increase of nationalism, it is likely that English documentation will be less acceptable, and translation will become increasingly common as more companies look to foreign markets. More to the point, texts, for information dissemination (export), must be translated with a great deal of care: the translation must be "right" as well as clear. Qualified human technical translators are hard to find, expensive, and slow (translating somewhere around 4-6 pages/day, on the average). The information dissemination application is most responsible for the renewed interest in MT.

Intended Applications of M(A)T

In another type of information dissemination, literary translation, there is little or no demand for machine processing, because, relative to technical translation, there is no shortage of human translators capable of fulfilling this proportionally miniscule need, and in any case computers do-not do it well. By contrast, in sheer volume the demand for technical translation is staggering; moreover, the acquisition, maintenance, and consistent use of valid technical terminology is an enormous problem. Worse, in many technical fields there is a distinct shortage of qualified human translators, and it is obvious that the problem will never be alleviated by such measures as greater incentives for translators, however laudable they may be. The only hope for a solution to the technical translation problem lies with increased human productivity through the full range of computer technology: full-scale MT, less ambitious MAT, on-line terminology databanks, and word processing. A serendipitous situation involves style: in literary translation, emphasis is placed on style, perhaps at the expense of absolute fidelity to content (especially, for example, with poetry). In technical translation, on the other hand, emphasis is properly placed on fidelity, even at the expense of style. M(A)T systems, which lack style but excel at terminological accuracy and consistency, are best suited tor technical translation.

Linguistic Techniques

Of the several perspectives from which one can view MT techniques, we will use the following: direct vs. indirect, interlingua vs. transfer, and local vs. global scope. (Not ail eight combinations are realized in practice.) In the past, some MT systems were distinguished by "the use of semantics." Now, for obvious reasons, all MT systems (are claimed to) make use of semantics, so this is no longer a distinguishing characteristic.

"Direct translation" characterizes systems (e.g., GAT [Hutchins, 1978; Jordan et al., 1979]) designed from the start to translate out of one specific language and into another. Direct systems are limited to the minimum work necessary to effect that translation; tor example, disambiguation is performed only to the extent necessary for translation into that one target language, irrespective of what might be required for another language. "Indirect translation," on the other hand, characterizes systems (e.g., EUROTRA [King, 1981, 1982]) in which the analysis of the source language and the synthesis of the target language are totally independent processes; for example, disambiguation is performed to the extent necessary to determine the "meaning" (however represented) of the source language input, irrespective of which target language(s) that input might be translated into.

In the "interlingua" approach (e.g., CETA [Hutchins, 1978]) the representation of the "meaning" of the source language input is (intended to be) independent of any language, and this same representation is used to synthesize the target language output. The "linguistic universals" searched for and debated about by linguists and philosophers is the notion that underlies an interlingua. Thus, the representation of a given "unit of meaning" would be the same, no matter what language (or grammatical structure) that unit might be expressed in. The "transfer" approach (e.g., TAUM [Isabelle, 1984]) utilizes an underlying representation of the "meaning" of a grammatical unit (e.g., sentence) which differs according to the language it was derived from (or into which it is to be generated); this approach implies the existence of a third stage of translation which converts one language-specific meaning representation into another, a stage called Transfer. Thus, the overall transfer translation process is analysis followed by transfer and then synthesis. The "transfer" vs. "interlingua" difference is not applicable to all systems; in particular, "direct" MT systems use neither the transfer nor the interlingua approach, since they do not attempt to represent "meaning."

"Local scope" vs. "global scope" is not so much a difference of category as degree. "Local scope" designates systems (e.g., SYSTRAN [Toma, 1976]) in which words are the essential unit driving analysis, and in which that analysis is, in effect, performed by separate procedures for each word which try to determine—on the basis of the words to the left and/or right—the part of speech, possible idiomatic usage, and "sense" of the word keying the procedure. In such systems, for example, homographs (words which differ in pan of speech and/or jrivational history and thus in meaning, but which re written alike) are a major problem, because a nified analysis of the sentence per se is not attempted. In "global scope" systems (e.g., METAL Lehmannetal, 1981; Slocum, 1980, 1983,1984]) the meaning of.a word is determined by its context within a unified analysis of the sentence (or, rarely, paragraph). In such systems, by contrast, homo-graphs do not typically constitute a significant problem because the amount of context taken into account is much greater than with systems of "local cope."

The State of the Art

Human languages are, by nature, different, so much so that the illusory goal of abstract perfec-tion in translation—once imagined (and still im-agined by some) to be achievable—can be comfor-tably ruled out of possibility, whether attempted by machine or man. Even the abstract notion of "quality" is undefinable, hence immeasurable. In its place, we must substitute the notion of evalua-tion of translation according to its purpose, judg-ed by the consumer. One must therefore accept the truth that the notion of quality is inherently sub-jective. Certainly there will be translations hailed by most if not all as "good," and correspondingly here will be translations almost universally labeled as "bad." Most translations, however, will surely all in between these extremes, and each user must render his own judgment according to his needs.

In corporate circles, however, there is and has always been an operational definition of "good" vs. "bad" translation: a good translation is what senior translators are willing to expose to outside scrutiny (not that they are fully satisfied, for they never are); and a bad one is what they are not willing to release. These experienced translators—usually post-editors—impose a judgment which the corporate body is willing to accept at face value; such judgment is the very purpose for having senior translators. Arrived at subjectively, according to the purpose for which the translation is intended, it still comes as close to being an objective assessment as the world is likely to see. In a post-editing context, a "good" original translation is one worth revising, i.e., one which the editor will endeavor to change, rather than reject or replace with his own translation.

Therefore, any rational position on the state of the art in MT and MAT must respect the operational decisions about the quality of MT and MAT as judged by the present users. These systems are _all, of course, based on old technology ("ancient," by the standards of Artificial Intelligence researchers); but by the time systems employing today's AI technology hit the market, they too will be "antiquated" by the research laboratory standards of their time. Such is the nature of technology. We will therefore distinguish, in our assessment, between what is available and/or used now ("old;" yet operationally current, technology), and what is around the next corner (techniques working in research labs today), and what is farther down the road (experimental approaches).

Production Systems

Production M(A)T systems are based on old technology; some, for example, still (or until very recently did) employ punch-cards and print(ed) out translations in all upper-case. Few if any attempt a comprehensive "global" analysis at the sentence level (trade secrets make this hard to discern), and none go beyond that to the paragraph level. None use a significant amount of semantic information (though all claim to use some). Most if not all perform as *idiots savants,* making use of enormous amounts of very unsophisticated pragmatic information and brute-force computation to determine the proper word-for-word or idiom-for-idiom translation followed by local rearrangement of word order, leaving the translation chaotic, even if understandable.

But they work! Some of them do. anyway, well enough that their customers find reason to invest enormous amounts of time and capital developing the necessary massive dictionaries specialized to their applications. Translation time is certainly reduced. Translator frustration is increased or decreased, as the case may be (it seems that personality differences, among other things, have a large bearing on this aspect). Some translators resist their introduction with varying degrees of success, but then there are those who still resist the introduction of typewriters, to say nothing of word processors. Yet most are thinking about accepting the place of computers in translation, and a few actually look forward to relief from much of the drudgery they now face. While current MT systems seem to take some getting used to, further productivity increases are realized as time goes by, and they are usually accepted, eventually, as a boon to the bored translator. Of the new products embodying old technology constantly being introduced, most are found not viable, and quickly disappear from the market. But those which have been around for years must be economically justifiable to their users, or else, presumably, they would no longer exist.

Development Systems

Systems being developed for near-term introduction employ Computational Linguistics (CL) techniques of the late 1970s, if not the 80s. Essentially all are full MT, not MAT, systems. As Hutchins (1982) notes, "there is now considerable agreement on the basic strategy, i.e., a 'transfer' system with some semantic analysis and some interlingual features in order to simplify transfer components." These systems employ one of a variety of sophisticated parsing/transducing techniques, typically based on charts, whether the grammar is expressed via phrase-structure rules (e.g., METAL) or (strings of) trees (e.g., GETA, EUROTRA); they operate at the sentence level, or higher, and make significant use of semantic features. Proper linguistic theories, whether elegant or not quite, and heuristic software strategies take the place of simple word substitutions and brute-force programming. If the analysis attempt succeeds, the translation stands a fair chance of being acceptable to the revisor; if analysis fails, then fail-soft measures are likely to produce something equivalent to the output of a current production MT system.

These systems work well enough in experimental settings to give their sponsors and waiting customers (to say nothing of their implementors) reason to hope for near-term success in application. Their technology is based on some of the latest techniques which appear to be workable in immediate large-scale application. Since most "pure AI" techniques do not fall in this category, serious AI researchers look down on these development systems (to say nothing of production systems) as old, uninteresting—and probably useless. Some likely are. But others, though "old," will soon find an application niche, and will begin displacing any of the current production systems which try to compete. (Since the present crop of development systems all seem to be aimed at the "information dissemination" application, the current production systems that are aimed at the "information acquisition" market may survive for some time.) The major hurdle is time: time to write and debug the grammars (a very hard task), and time to develop lexicons with roughly ten thousand general vocabulary items, and the few tens of thousands of technical terms required per subject area. Some development projects have invested the necessary rime, and stand ready to deliver commercial applications (e.g., GETA, METAL). In 1984, four Japanese companies announced commercial MT systems.

Research Systems

The biggest problem associated with MT research systems is their scarcity (nonexistence in the United States). If current CL and AI researchers were seriously interested in multiple languages—even if not for translation per se—this would not necessarily be a bad situation. But in the United States they certainly are not, and in Europe, AI research has not yet reached the level achieved in the United States. Western business and industry **arc** naturally more concerned with near-term payoff, and some track development systems; very few support MT development directly, and none yet support pure MT research at a significant level. (The Dutch firm Philips may, indeed, have the only long-term research project in the West, but it is a very small one.) Some European governments (e.g., Germany and France) fund significant R&D projects, but Japan is making by far the world's largest investment in MT research. The United States government, which otherwise supports the best overall AI and (English) Computational Linguistics research in the world, is not involved.

Where pure MT research projects do exist, they tend to concentrate on the problems of deep meaning representations, striving to pursue the goal of a true AI system, which would presumably include language-independent meaning representations of great depth and complexity. Translation here is seen as just one application of such a system: the system "understands" natural language input, then "generates" natural language output; if the languages happen to be different, then "translation" has been performed via paraphrase. Translation could thus be viewed as one of the ultimate tests of an Artificial Intelligence: if a system "translates correctly" then to some extent it can be argued to lave "understood correctly," and in any case will ell us much about what translation is all about. In this role, MT research holds out its greatest promise as a once-again scientifically respectable discipline. The first requirement, however, is the existence of research groups interested in, and funded for, the study of multiple languages and translation among hem within the framework of AI research. At the present time only Japan, and to a somewhat lesser extent western Europe, can boast such groups.

Future Prospects

The world has changed in the two decades since ALPAC. The need and demand for technical translation has increased dramatically, and the supply of qualified human technical translators has not kept pace. (Indeed, it is debatable whether even in 1966 there existed a sufficient supply of qualified technical translators, contrary to ALPAC's claims.) For whatever reasons, the classic "law of supply and demand" has not worked in this instance; the shortage is real, all over the world, and nothing is yet serving to stem this worsening situation. Outside of dramatic productivity increases via computer automation, nothing seems capable of doing so. in the EEC, for example, the already overwhelming load of technical translation is projected to rise sixfold within five years.

In technical translation, the future promises greater acceptance by translators of the role of machine aids, running the gamut from word processing systems and on-line term banks to MT systems. Correspondingly, M(A)T systems will experience greater success in the marketplace. As these systems continue to drive down the cost of translation, the demand and capacity for that service will grow even more than it would otherwise, and many "new" needs for translation, not presently economically justifiable, will surface. If MT systems are to continue to improve so as to further reduce the burden on human translators, there will be a greater need and demand for continuing MT R&D efforts.

Conclusions

The translation problem will not go away, and human solutions (short of full automation) do not now, and never will, suffice. MT systems have already scored successes among the user community, a trend that can hardly fail to continue as users demand further improvements and greater speed, and MT system vendors respond. Of course, the need for research is great, but some current and future applications will continue to succeed on economic grounds alone—and to the user community, this is virtually the only measure of success or failure.

It is important to note that translation systems are not going to "fall out" of AI efforts which are not seriously contending with multiple languages from the start. There are two reasons for this. First, English is not a representative language. Relatively speaking, it is not even a very hard language from the standpoint of Computational Linguistics. Surely in part because of the nearly total concentration of CL workers on English, other languages, such as Japanese, Chinese, Russian, and even German, seem more difficult to deal with using existing CL techniques. Developing translation capability will require similar concentration by CL workers on other languages; nothing less will suffice.

Second, it would seem that translation is not by any means a simple matter of understanding the source text, then reproducing it in the target language, even though some translators (and virtually every layman) will say this is so. One may question, in the example of an article on front-line research in semiconductor switching theory, or nuclear physics, whether a translator "fully comprehends" the content of the article he is translating. One would suspect not. (Johnson [1983] claims that he has produced translations, judged good by informed peers, in technical areas where his expertise is deficient, and his understanding incomplete.) On the other hand, it is also true that translation schools expend a great deal of effort teaching techniques for low-level lexical and syntactic manipulation—a curious fact to contrast with the usual "full comprehension" claim. In any event, every qualified translator will agree that there is much more to translation than simple analysis/synthesis (an almost *prima facie* proof of the necessity for Transfer).

What this means is that the development of translation as an application of Computational Linguistics will require substantial research in its own right in addition to the work necessary to provide the basic multi-lingual analysis and synthesis tools. Translators must be consulted, for they are the experts in translation. (They also stand to gain the most: automation of the monotonous aspects of translation will result in increased job interest and prestige, as translators are promoted to the ranks of editors.) None of this will happen by accident; it must result from design.

REFERENCES

ALPAC. Languages and Machines: Computers in Translation and Linguistics. A report by the Automatic Language Processing Advisory Committee (ALPAC], Division of Behavioral Sciences, National Academy of Sciences, National Research Council, Publication 1416, Washington, D.C., 1966.

Bar-Hillel, Y., "Some Reflections of the Present Outlook for High-Quality Machine Translation." in W.P. Lehmann and R. Stachowitz (eds.). Feasibility Study on Fully Automatic High Quality Translation. Final technical report RADC-TR-71-295. Linguistics Research Center, University of Texas at Austin, December 1971.

Hutchins, W.J., "Progress in Documentation: Machine Translation and Machine-Aided Translation;' *Journal of Documentation* 34, 2, June 1978, pp. U9-159.

Hutchins, W.J., "The Evolution of Machine Translation Systems," in V. Lawson (ed.), *Practical Experience of* Machine *Translation.* North-Holland, Amsterdam, 1982, pp 21-37,

Isabelle, P., "Machine Translation at the TAUM Group" presented at the ISSCO Tutorial on Machine Translation,. Lugano, Switzerland, 2-6 April 1984.

Johnson, R.L., "Parsing—an MT Perspective," in K.S. Jones an Y.
Wilks (eds.), *Automatic Natural Language Parsing.* Ell Horwood,
Ltd. Chichester, Great Britain, 1983.

Jordan, S.R., A.F.R. Brown, and F.C.. Hutton, "Computerized Rus-sian Translation at ORNL" in *Proceedings of the ASIS Annual Meeting, San Francisco, 1976,* p. 163: also in *ASIS Journal 28,* 1, 1977, pp. 26-33.

King, M., "EUROTRA—A European System for Machine Translation',' *Lebende Sprachen* 26, 1981, pp. 12-I4.

King, M., "EUROTRA: An Attempt to Achieve Multilingual MT" in V. Lawson (ed.), *Practical Experience of Machine Translation.* North-Holland, Amsterdam 1982, pp. 139-147.

King, M., and S. Perschke, "EUROTRA and its Objectives" *Multilingua,* 1,1, 1982, pp. 27-32.

Lehmann, W.P., W.S. Bennett, J. Slocum, H. Smith, S.M.V. Pfluger, and S.A. Eveland. The METAL System. Final technical report RADC-TR-80-374, Linguistics Research Center, University of Texas at Austin, January 1981. Available as NTIS report AO-97896.

Slocum, J., "An Experiment in Machine Translation," *Proceedings of the Eighteenth Annual Meeting of the Association for Computational Linguistics, Philadelphia, 19-22 June 1980,* pp. 163-167.

Slocum, J., "A Status Report on the LRC Machine Translation System" *Proceedings of the ACL-NRL Conference on Applied Natural Language Processing,* Santa Monica. California, 1-3 February 1983, pp. 166-173.

Slocum, J., "METAL: The LRC Machine Translation System" presented at the ISSCO Tutorial on Machine Translation, Lugano, Switzerland, 2-6 April 1984.

Toma, P., "An Operational Machine Translation System," in R.W. Brislin (ed.). *Translation: Applications and Research.* Gardner Press, New York. 1976. pp. 247-259