Machine Translation:
its History, Current Status,
and Future Prospects

Jonathan Slocum

Siemens Communications Systems, Inc.
Linguistics Research Center
University of Texas
Austin, Texas

Abstract

Elements of the history, state of the art, and probable future of Machine Translation (MT) are discussed. The treatment is largely tutorial, based on the assumption that this audience is, for the most part, ignorant of matters pertaining to translation in general, and MT in particular. The paper covers some of the major MT R&D groups, the general techniques they employ(ed), and the roles they play(ed) in the development of the field. The conclusions concern the seeming permanence of the translation problem, and potential re-integration of MT with mainstream Computational Linguistics.

Introduction

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 1950's, interest and support was 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 quite heavily). During its second decade in the 1960's, 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 in 1966, condemning the field and, indirectly, its workers alike. The ALPAC report was criticized as narrow, biased, and short-sighted, but its recommendations were adopted (with the important exception of increased expenditures for long-term research in computational linguistics), and as a result MT projects were cancelled in the U.S. and elsewhere around the world. By 1973, the early part of the third decade of MT, only three government-funded projects were left in the U.S., and by late 1975 there were none. Paradoxically, MT systems were still being used by various government agencies here and abroad, because there was simply no alternative means of gathering information from foreign [Russian] sources so quickly; in addition, private companies were developing and selling MT systems based on the mid-60's technology so roundly castigated by ALPAC. Nevertheless the general disrepute of MT resulted in a remarkably quiet third decade.

We are now into the fourth decade of MT, and there is a resurgence of interest throughout the world — plus a growing number of MT and MAT (Machine-aided Translation) systems in use by governments, business and industry. Industrial firms are also beginning to fund M(A)T R&D projects of their own; thus it can no longer be said that only government funding keeps the field alive (indeed, in the U.S. 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, we will present a brief introduction to MT technology and a short historical perspective before considering the present status and state of the art, and then moving on to a discussion of the future prospects. For reasons of space and perspicuity, we shall concentrate on MT efforts in the U.S. and western Europe, though some other MT projects and less-sophisticated approaches will receive attention.

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 U.S., it is common to conceive of translation as simply that which a human translator does. It is generally believed that a college degree [or the equivalent] in a foreign language qualifies one to be a translator for just about any material whatsoever. Native speakers of foreign languages are considered to be that much more qualified. Thus, translation is not particularly respected as a profession in the U.S., and the pay is poor.

In Canada, in Europe, and generally around the world, this myopic attitude is not held. Where translation is a fact of life rather than an oddity, it is realized that any translator's competence is sharply restricted to a few domains (this is especially true of technical areas), and that native fluency in a foreign language does not bestow on one the ability to serve as a translator.

Thus, there are college-level and post-graduate schools that teach the theory (translatology) as well as the practice of translation; thus, a technical translator is trained in the few areas in which he will be doing translation.

Of special relevance to MT is the fact 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 "pre-publication stylistic editing".) Unrevised translations are always regarded as 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 post-edited. Such news generally comes as a surprise, if not a shock, to most people in the US.

It is easy to see, therefore, that the "fully-automatic high-quality machine translation" standard, imagined by most U.S. scholars to constitute minimum acceptability, must be radically redefined. 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, 71]. So an MT system does not have to 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. Academic debates about what constitutes "high-quality" and "fully-automatic" are considered irrelevant by the 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, like speed.

Machine Translation Technology

In order to appreciate the differences among translation systems (and their applications), 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.

MT systems are intended to perform translation without human intervention. This does not rule out pre-processing (assuming this is not for the

purpose of marking phrase boundaries and resolving part-of-speech and/or other ambiguities, etc.), nor post-editing (since this is normally done for human translations anyway). However, an MT system is solely responsible for the complete translation process from input of the source text to output of the target text without human assistance, using special programs, comprehensive dictionaries, and collections of linguistic rules (to the extent they exist, varying with the MT system). MT occupies the top range of positions on the scale of computer translation sophistication.

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. HAMT refers to a system wherein the computer is responsible for producing the translation per se, but may interact with a human monitor at many stages along the way -- for example, 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. MAHT refers to a system wherein 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 a MA(H)T system (the system does not need help, instead, 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 translator may not be working on-line), but usually is performed prior to human translation. Indeed the databank may not be accessible (to the translator) 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 the fact 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: there is already more to read than can be read in the time available, and having to labor through texts written in foreign languages -- when the probability is low that any given text is of real interest -- is not worth the effort. In the past, the lingua franca of science has been English; this is becoming less and less true 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, he could understand the content of the text well enough that a more careful translation would not be necessary.

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, U.S. companies have escaped this responsibility by requiring that the purchasers learn English; other exporters (German, for example) have never had this luxury. In the future, with the increase of nationalism, it is less likely that English documentation will be acceptable. Translation is becoming 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

Although literary translation is a case of information dissemination, there is little or no demand for literary translation by machine: relative to technical translation, there is no shortage of human translators capable of fulfilling this need, and in any case computers do not fare well at literary translation. By contrast, the demand for technical translation is staggering in sheer volume; 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 measures such as greater incentives for translators, however laudable that may be. The only hope for a solution to the technical translation problem lies with increased human productivity through computer technology: full-scale MT, less ambitious MAT, on-line terminology databanks, and word-processing all have their place. A serendipitous situation

involves style: in literary translation, emphasis is placed on style, perhaps at the expense of absolute fidelity to content (especially for poetry). In technical translation, emphasis is properly placed on fidelity, even at the expense of style. M(A)T systems lack style, but excel at terminology: they are best suited for technical translation.

Linguistic Techniques

There are 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 all eight combinations are realized in practice.) We shall characterize MT systems from these perspectives, in our discussions. In the past, "the use of semantics" was always used to distinguish MT systems; those which used semantics were labelled 'good', and those which did not were labelled 'bad'. Now all MT systems [are claimed to] make use of semantics, for obvious reasons, so this is no longer a distinguishing characteristic.

"Direct translation" is characteristic of a system (e.g., GAT) 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; for 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, is characteristic of a system (e.g., EUROTRA) wherein 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.

The "interlingua" approach is characteristic of a system (e.g., CETA) in which 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 is characteristic of a system (e.g., TAUM) in which the underlying representation of the 'meaning' of a grammatical unit (e.g., sentence) differs depending on the language it was derived from [or into which it is to be generated]; this implies the existence of a third translation stage which maps one language-specific meaning representation into another: this stage is 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" characterizes a system (e.g., SYSTRAN) 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 -- based on 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 part of speech and/or derivational history [thus meaning], but which are written alike) are a major problem, because a unified analysis of the sentence per se is not attempted. "Global scope" characterizes a system (e.g., METAL) in which 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, homographs do not typically constitute a significant problem because the amount of context taken into account is much greater than is the case with systems of "local scope."

Historical Perspective

There are several comprehensive treatments of MT projects [Bruderer, 77] and MT history [Hutchins, 78] available in the open literature. To illustrate some continuity in the field of MT, while remaining within reasonable space limits, our brief historical overview will be restricted to defunct systems/projects which gave rise to follow-on systems/projects of current interest. These are: Georgetown's GAT, Grenoble's CETA, Texas' METAL, Montreal's TAUM, and Brigham Young University's ALP system.

GAT - Georgetown Automatic Translation

Georgetown University was the site of one of the earliest MT projects. Begun in 1952, and supported by the U.S. government, Georgetown's GAT system became operational in 1964 with its delivery to the Atomic Energy Commission at Oak Ridge National Laboratory, and to Europe's corresponding research facility EURATOM in Ispra, Italy. Both systems were used for many years to translate Russian physics texts into "English." The output quality was quite poor, by comparison with human translations, but for the intended purpose of quickly scanning documents to determine their content and interest, the GAT system was nevertheless superior to the only alternatives: slow and more expensive human translation or, worse, no translation at all. GAT was not replaced at EURATOM until 1976; at ORNL, it seems to have been used until around 1979 [Jordan et al., 76, 77].

The GAT strategy was "direct" and "local": simple word-for-word replacement, followed by a limited amount of transposition of words to result in something vaguely resembling English. Very soon, a 'word' came to be defined as a single word or a sequence of words forming an 'idiom'. There was no

true linguistic theory underlying the GAT design; and, given the state of the art in computer science, there was no underlying computational theory either. GAT was developed by being made to work for a given text, then being modified to account for the next text, and so on. The eventual result was a monolithic system of intractable complexity: after its delivery to ORNL and EURATOM, it underwent no significant modification. The fact that it was used for so long is nothing short of remarkable -- a lesson in what can be tolerated by users who desperately need translation services for which there is no viable alternative to even low-quality MT.

The termination of the Georgetown MT project in the mid-60's resulted in the incorporation of LATSEC by Peter Toma, one of the GAT workers. LATSEC soon developed the SYSTRAN system (based on GAT technology), which in 1970 replaced the IBM Mark II system at the USAF Foreign Technology Division (FTD) at Wright Patterson AFB, and in 1976 replaced GAT at EURATOM. SYSTRAN is still being used to translate Russian into English for information-acquisition purposes. We shall return to our discussion of SYSTRAN in the next major section.

CETA - Centre d'Études pour la Traduction Automatique

In 1961 a project was started at Grenoble University in France, to translate Russian into French. Unlike GAT, Grenoble began the CETA project with a clear linguistic theory -- having had a number of years in which to witness and learn from the events transpiring at Georgetown and elsewhere. In particular, it was resolved to achieve a dependency-structure analysis of every sentence (a "global" approach) rather than rely on intra-sentential heuristics to control limited word transposition (the "local" approach); with a unified analysis in hand, a reasonable synthesis effort could be mounted. The theoretical basis of CETA was "interlingua" (implying a language-independent, "neutral" meaning representation) at the grammatical level, but "transfer" (implying a mapping from one language-specific meaning representation to another) at the lexical [dictionary] level. The state of the art in computer science still being primitive, Grenoble was essentially forced to adopt IBM assembly language as the software basis of CETA [Hutchins, 78].

The CETA system was under development for ten years; during 1967-71 it was used to translate 400,000 words of Russian mathematics and physics texts into French. The major findings of this period were that the use of an interlingua erases all clues about how to express the translation; also, that it results in extremely poor or no translations of sentences for which complete analyses cannot be derived. The CETA workers learned that it is critically important in an operational system to retain surface clues about how to formulate the translation (Indo-European languages, for example, have many structural similarities, not to mention cognates, that one can

take advantage of), and to have "fail-soft" measures designed into the system. An interlingua does not allow this [easily, if at all], but the transfer approach does.

A change in hardware (thus software) in 1971 prompted the abandonment of the CETA system, immediately followed by the creation of a new project/system called GETA, based entirely on a fail-soft transfer design. The software was still, however, written in assembly language; this continued reliance on assembly language was soon to have deleterious effects, for reasons now obvious to anyone. We will return to our discussion of GETA, below.

METAL - MEchanical Translation and Analysis of Languages

Having had the same opportunity for hindsight, the University of Texas in 1961 used U.S. government funding to establish the Linguistics Research Center, and with it the METAL project, to investigate MT -- not from Russian, but from German into English. The LRC adopted Chomsky's transformational paradigm, which was quickly gaining popularity in linguistics circles, and within that framework employed a syntactic interlingua based on deep structures. It was soon discovered that transformational linguistics per se was not sufficiently well-developed to support an operational system, and certain compromises were made. The eventual result, in 1974, was an 80,000-line, 14-overlay FORTRAN program running on a dedicated CDC 6600. Indirect translation was performed in 14 steps of global analysis, transfer, and synthesis -- one for each of the 14 overlays -- and required prodigious amounts of CPU time and I/O from/to massive data files. U.S. government support for MT projects was winding down in any case, and the METAL project was shortly terminated.

Several years later, a small Government grant resurrected the project. The FORTRAN program was rewritten in LISP to run on a DEC-10; in the process, it was pared down to just three major stages (analysis, transfer, and synthesis) comprising about 4,000 lines of code which could be accommodated in three "overlays," and its computer resource requirements were reduced by a factor of ten. Though U.S. government interest once again languished, the Sprachendienst (Language Services) department of Siemens AG in Munich had begun supporting the project, and in 1980 Siemens AG became the sole sponsor.

TAUM - Traduction Automatique de l'Université de Montréal

In 1962 the University of Montreal established the TAUM project with Canadian government funding. This was probably the first MT project designed strictly around the transfer approach. As the software basis of the project, TAUM chose the PASCAL programming language on the CDC 6600. After an initial period of more-or-less open-ended research, the Canadian government began adopting specific goals for the TAUM system. A chance remark by a bored translator in the Canadian

Meteorological Center had led to a spin-off project: TAUM-METEO. Weather forecasters were already required to adhere to a prescribed manual of style and vocabulary in their English reports. Partly as a result of this, translation into French was so monotonous a task that human translator turnover in the weather service was extraordinarily high -- six months was the average tenure. TAUM was commissioned in 1975 to produce an operational English-French MT system for weather forecasts. A prototype was demonstrated in 1976, and by 1977 METEO was installed for production translation. We will discuss METEO in the next major section.

The next challenge was not long in coming: by a fixed date, TAUM had to be usable for the translation of a 90 million word set of aviation maintenance manuals from English into French (else the translation had to be started by human means, since the result was needed quickly). From this point on, TAUM concentrated on the aviation manuals exclusively. To alleviate problems with their purely syntactic analysis (especially considering the many multiple-noun compounds present in the aviation manuals), the group began in 1977 to incorporate partial semantic analysis in the TAUM-AVIATION system.

After a test in 1979, it became obvious that TAUM-AVIATION was not going to be production-ready in time for its intended use. The Canadian government organized a series of tests and evaluations to assess the status of the system. Among other things, it was discovered that the cost of writing each dictionary entry was remarkably high (3.75 man-hours, costing \$35-40), and that the system's runtime translation cost was also high (6 cents/word) considering the cost of human translation (8 cents/word), especially when the post-editing costs (10 cents/word for TAUM vs. 4 cents/word for human translations) were taken into account [Gervais, 1980]; TAUM was not yet cost-effective. Several other factors, especially the bad Canadian economic situation, combined with this to cause the cancellation of the TAUM project in 1981. There are recent signs of renewed interest in MT in Canada. State-of-the-art surveys have been commissioned [Pierre Isabelle, formerly of TAUM, personal communication], but no successor project has yet been established.

ALP - Automated Language Processing

In 1971 a project was established at Brigham Young University to translate Mormon ecclesiastical texts from English into multiple languages -- starting with French, German, Portuguese and Spanish. The eventual aim was to produce a fully-automatic MT system based on Junction Grammar [Lytle et al., 75], but actual work proceeded on Machine-Aided Translation (MAT, where the system does not attempt to analyze sentences on its own, according to pre-programmed linguistic rules, but instead relies heavily on interaction with a human to effect the analysis [if one is even attempted] and complete the translation).

The BYU project never produced an operational system, and the Mormon Church, through the

University, began to dismantle the project. Around 1977, a group composed primarily of programmers left BYU to join Weidner Communications, Inc., and proceeded to develop the fully-automatic, direct Weidner MT system. Shortly thereafter, most of the remaining BYU project members left to form Automated Language Processing Systems (ALPS) and continue development of the BYU MAT system. Both of these systems are actively marketed today, and will be discussed in the next section. Some work continues at BYU, but at a very much reduced level and degree of aspiration (e.g., [Melby, 82]).

Current Production Systems

In this section we consider the major M(A)T systems being used and/or marketed today. Four of these originate from the "failures" described above, but four systems are essentially the result of successful (i.e., continuing) MT R&D projects. The full MT systems discussed below are the following: SYSTRAN, LOGOS, METEO, Weidner, and SPANAM; we will also discuss the MAT systems CULT and ALPS. Most of these systems have been installed for several customers (METEO, SPANAM, and CULT are the exceptions, with only one obvious "user" each). The oldest installation dates from 1970.

A "standard installation," if it can be said to exist, includes provision for pre-processing in some cases, translation (with much human intervention in the case of MAT systems), and some amount of post-editing. To MT system users, acceptability is a function of the amount of pre- and/or post-editing that must be done (which is also the greatest determinant of cost). Van Slype [82] reports that "acceptability to the human translator...appears negotiable when the quality of the MT system is such that the correction (i.e., post-editing) ratio is lower than 20% (1 correction every 5 words) and when the human translator can be associated with the upgrading of the MT system." It is worth noting that editing time has been observed to fall with practice: Pigott [82] reports that "...the more M.T. output a translator handles, the more proficient he becomes in making the best use of this new tool. In some cases he manages to double his output within a few months as he begins to recognize typical M.T. errors and devise more efficient ways of correcting them."

It is also important to realize that, though none of these systems produces output mistakable for human translation [at least not good human translation], their users have found sufficient reason to continue using them. Some users, indeed, are repeat customers. In short, MT & MAT systems cannot be argued not to work, for they are in fact being bought and used, and they save time and/or money for their users. Every user expresses a desire for improved quality and reduced cost, to be sure, but then the same is said about human translation. Thus, in the only valid sense of the idiom, MT & MAT have already "arrived." Future improvements in quality, and reductions in cost -- both certain to take place -- will serve to make M(A)T systems even more attractive.

SYSTRAN

SYSTRAN was one of the first MT systems to be marketed; the first installation replaced the IBM Mark II Russian-English system at the USAF FTD in 1970, and is still operational. Based on the GAT technology (SYSTRAN uses the same linguistic strategies, to the extent they can be argued to exist), SYSTRAN's software basis has been much improved by the introduction of modularity (separating the analysis and synthesis stages), by a recent shift away from simple "direct" translation (from the Source Language straight into the Target Language) toward the inclusion of something resembling an intermediate "transfer" stage, and by the allowance of manually-selected topical glossaries (essentially, dictionaries specific to [the subject area of] the text). The system is still ad hoc -- particularly in the assignment of semantic features [Pigott, 79]. The USAF FTD dictionaries number over a million entries; Bostad [82] reports that dictionary updating must be severely constrained, lest a change to one entry disrupt the activities of many others. (A study by Wilks [78] reported an improvement/degradation ratio [after dictionary updates] of 7:3, but Bostad implies a much more stable situation after the introduction of stringent [and expensive] quality-control measures.) NASA selected SYSTRAN in 1974 to translate materials relating to the Apollo-Soyuz collaboration, and EURATOM replaced GAT with SYSTRAN in 1976. Also by 1976, FTD was augmenting SYSTRAN with word-processing equipment to increase productivity (e.g., to eliminate the use of punch-cards).

In 1976 the Commission of the European Communities purchased an English-French version of SYSTRAN for evaluation and potential use. Unlike the FTD, NASA, and EURATOM installations, where the goal was information acquisition, the intended use by CEC was for information dissemination -- meaning that the output was to be carefully edited before human consumption. Van Slype [82] reports that "the English-French standard vocabulary delivered by Prof. Toma to the Commission was found to be almost entirely useless for the Commission environment." Early evaluations were negative (e.g., Van Slype [79]), but the existing and projected overload on CEC human translators was such that investigation continued in the hope that dictionary additions would improve the system to the point of usability. Additional versions of SYSTRAN were purchased (French-English in 1978, and English-Italian in 1979). The dream of acceptable quality for post-editing purposes was eventually realized: Pigott [82] reports that "...the enthusiasm demonstrated by [a few translators] seems to mark something of a turning point in [machine translation]." Currently, about 20 CEC translators in Luxembourg are using SYSTRAN on a Siemens 7740 computer for routine translation; one factor accounting for success is that the English and French dictionaries now consist of well over 100,000 entries in the very few technical areas for which SYSTRAN is being employed.

Also in 1976, General Motors of Canada acquired SYSTRAN for translation of various manuals (for vehicle service, diesel locomotives, and highway transit coaches) from English into French on an IBM mainframe. GM's English-French dictionary had been expanded to over 130,000 terms by 1981 [Sereda, 82]. Subsequently, GM purchased an English-Spanish version of SYSTRAN, and is now working to build the necessary [very large] dictionary. Sereda [82] reports a speed-up of 3-4 times in the productivity of his human translators (from about 1000 words per day); he also reveals that developing SYSTRAN dictionary entries costs the company approximately \$4 per term (word- or idiom-pair).

While other SYSTRAN users have applied the system to unrestricted texts (in selected subject areas), Xerox has developed a restricted input language ('Multinational Customized English') after consultation with LATSEC. That is, Xerox requires its English technical writers to adhere to a specialized vocabulary and a strict manual of style. SYSTRAN is then employed to translate the resulting documents into French, Italian, and Spanish; Xerox hopes to add German and Portuguese. Ruffino [82] reports "a five-to-one gain in translation time for most texts" with the range of gains being 2-10 times. This approach is not necessarily feasible for all organizations, but Xerox is willing to employ it and claims it also enhances source-text clarity.

Currently, SYSTRAN is being used in the CEC for the routine translation, followed by human post-editing, of around 1,000 pages of text per month in the couples English-French, French-English, and English-Italian [Wheeler, 83]. Given this relative success in the CEC environment, the Commission has recently ordered an English-German version as well as a French-German version. Judging by past experience, it will be quite some time before these are ready for production use, but when ready they will probably save the CEC translation bureau valuable time, if not real money as well.

LOGOS

Development of the LOGOS system was begun in 1964. The first installation, in 1971, was used by the U.S. Air Force to translate English maintenance manuals for military equipment into Vietnamese. Due to the termination of U.S. involvement in that war, and perhaps partly to a poor evaluation of LOGOS' cost-effectiveness [Sinaiko and Klare, 73], its use was ended after two years. As with SYSTRAN, the linguistic foundations of LOGOS are weak and inexplicit (they appear to involve dependency structures); and the analysis and synthesis rules, though separate, seem to be designed for particular source and target languages, limiting their extensibility.

LOGOS continued to attract customers. In 1978, Siemens AG began funding the development of a LOGOS German-English system for telecommunications manuals. After three years LOGOS delivered a "production" system, but it was not found suitable for use (due in part to poor quality of the

translations, and in part to the economic situation within Siemens which had resulted in a much-reduced demand for translation, hence no immediate need for an MT system). Eventually LOGOS forged an agreement with the Wang computer company which allowed LOGOS to implement the German-English system (formerly restricted to large IBM mainframes) on Wang office computers. This system is being marketed today, and has recently been purchased by the Commission of the European Communities. Development of other language pairs has been mentioned from time to time.

METEO

TAUM-METEO is the world's only example of a truly fully-automatic MT system. Developed as a spin-off of the TAUM technology, as discussed earlier, it was fully integrated into the Canadian Meteorological Center's (CMC's) nation-wide weather communications network by 1977. METEO scans the network traffic for English weather reports, translates them "directly" into French, and sends the translations back out over the communications network automatically. Rather than relying on post-editors to discover and correct errors, METEO detects its own errors and passes the offending input to human editors; output deemed "correct" by METEO is dispatched without human intervention, or even overview.

TAUM-METEO was probably also the first MT system where translators were involved in all phases of the design/development/refinement; indeed, a CMC translator instigated the entire project. Since the restrictions on input to METEO were already in place before the project started (i.e., METEO imposed no new restrictions on weather forecasters), METEO cannot quite be classed with the TIUS and Xerox SYSTRAN systems which rely on restrictions geared to the characteristics of those MT systems. But METEO is not extensible.

One of the more remarkable side-effects of the METEO installation is that the translator turn-over rate within the CMC went from 6 months, prior to METEO, to several years, once the CMC translators began to trust METEO's operational decisions and not review its output [Brian Harris, personal communication]. METEO's input constitutes over 11,000 words/day, or 3.5 million words/year. Of this, it correctly translates 80%, shuttling the other ("more interesting") 20% to the human CMC translators; almost all of these "analysis failures" are attributable to violations of the CMC language restrictions, though some are due to the inability of the system to handle certain constructions. METEO's computational requirements total about 15 CPU minutes per day on a CDC 7600 [Thouin, 82]. By 1981, it appeared that the built-in limitations of METEO's theoretical basis had been reached, and further improvement was not possible.

Weidner Communications Systems, Inc.

Weidner was established in 1977 by Bruce Weidner, who hired a group of MT workers (predominantly programmers) from the fading BYU project. Weidner

delivered a production English-French system to Mitel in Canada in 1980, and a beta-test English-Spanish system to the Siemens Corporation (USA) in the same year. In 1981 Mitel took delivery on Weidner's English-Spanish and English-German systems, and Bravice (a translation service bureau in Japan) purchased the Weidner English-Spanish and Spanish-English systems. To date, there are about 22 installations of the Weidner MT system around the world. The Weidner system, though "fully automatic" during translation, is marketed as a "machine aid" to translation (perhaps to avoid the stigma usually attached to MT). It is highly interactive for other purposes (the lexical pre-analysis of texts, the construction of dictionaries, etc.), and integrates word-processing software with external devices (e.g., the Xerox 9700 laser printer at Mitel) for enhanced overall document production. Thus, the Weidner system accepts a formatted source document (actually, one containing formatting/typesetting codes) and produces a formatted translation. This is an important feature to users, since almost everyone is interested in producing formatted translations from formatted source texts.

Given the way this system is tightly integrated with modern word-processing technology, it is difficult to assess the degree to which the translation component itself enhances translator productivity, vs. the degree to which simple automation of formerly manual (or poorly automated) processes accounts for the productivity gains. The "direct" translation component itself is not particularly sophisticated. For example analysis is "local," being restricted to the noun phrase or verb phrase level -- so that context available only at higher levels can never be taken into account. Translation is performed in four independent stages: idiom search, homograph disambiguation, structural analysis, and transfer. These stages do not interact with each other, which creates more problems; for example, an apparent idiom in a text is always treated idiomatically -- never literally, no matter what its context (since no other contextual information is available until later). Hundt [82] comments that "idioms are an extremely important part of the translation procedure." It is particularly interesting that he continues: "...machine assisted translation is for the most part word replacement..." Then, "It is not worthwhile discussing the various problems of the [Weidner] system in great depth because in the first place they are much too numerous..." Yet even though the Weidner translations are of low quality, users nevertheless report economic satisfaction with the results. Hundt continues "...the Weidner system indeed works as an aid..." and, "800 words an hour as a final figure [for translation throughput] is not unrealistic." This level of performance was not attainable with previous [human] methods, and some users report the use of Weidner to be cost-effective, as well as faster, in their environments.

In 1982, Weidner delivered English-German and German-English systems to ITT in Great Britain; but there were some financial problems (a third of the

employees were laid off that year) until a controlling interest was purchased by a Japanese company: Bravice, one of Weidner's customers, owned by a group of wealthy Japanese investors. Weidner continues to market MT systems, and is presently working to develop Japanese MT systems. A prototype Japanese-English system has recently been installed at Bravice, and work continues on an English-Japanese system. In addition, Weidner has implemented its system on the IBM Personal Computer, in order to reduce its former dependence on the PDP-11.

SPANAM

Following a promising feasibility study, the Pan American Health Organization in Washington, D.C. decided in 1975 to undertake work on a machine translation system, utilizing many of the same techniques developed for GAT; consultants were hired from nearby Georgetown University, the home of GAT. The official PAHO languages are English, French, Portuguese, and Spanish; Spanish-English was chosen as the initial language pair, due to the belief that "This combination requires fewer parsing strategies in order to produce manageable output [and other reasons relating to expending effort on software rather than linguistic rules]" [Vasconcellos, 83]. Actual work started in 1976, and the first prototype was running in 1979, using punched card input on an IBM mainframe. With the subsequent integration of a word processing system, production use could be seriously considered.

After further upgrading, the system in 1980 was offered as a service to potential users. Later that year, in its first major test, SPANAM reduced manpower requirements for a certain translation effort by 45%, resulting in a monetary savings of 61% [Vasconcellos, 83]. Since then it has been used to translate well over a million words of text, averaging about 4,000 words per day per post-editor. (Significantly, SPANAM's in-house developers seem to be the only revisors of its output.) The post-editors have amassed "a bag of tricks" for speeding the revision work, and special string functions have also been built into the word processor for handling SPANAM's English output.

Sketchy details imply that the linguistic technology underlying SPANAM is essentially that of GAT; the rules may even still be built into the programs. The software technology has been updated considerably in that the programs are modular (in the newest version). The total lack of sophistication by modern Computational Linguistics standards is evidenced by the offhand remark that "The maximum length of an idiom [allowed in the dictionary] was increased from five words to twenty-five" in 1980 [Vasconcellos, 83]. Also, the system adopts the "direct" translation strategy, and fails to attempt a "global" analysis of the sentence, settling for "local" analysis of limited phrases. The SPANAM dictionary currently numbers 55,000 entries. A follow-on project to develop ENGSPAN, underway since 1981, has produced some test translations.

CULT - Chinese University Language Translator

CULT is perhaps the most successful of the Machine-aided Translation systems. Development began at the Chinese University of Hong Kong around 1968. CULT translates Chinese mathematics and physics journals (published in Beijing) into English through a highly-interactive process [or, at least, with a lot of human intervention]. The goal was to eliminate post-editing of the results by allowing a large amount of pre-editing of the input, and a certain [unknown] degree of human intervention during translation. Although published details [Loh, 76, 78, 79] are not unambiguous, it is clear that humans intervene by marking sentence and phrase boundaries in the input, and by indicating word senses where necessary, among other things. (What is not clear is whether this is strictly a pre-editing task, or an interactive task.) CULT runs on the ICL 1904A computer.

Beginning in 1975, the CULT system was applied to the task of translating the Acta Mathematica Sinica into English; in 1976, this was joined by the Acta Physica Sinica. This production translation practice continues to this day. Originally the Chinese character transcription problem was solved by use of the standard telegraph codes invented a century ago, and the input data was punched on cards. But in 1978 the system was updated by the addition of word-processing equipment for on-line data entry and pre/post-editing.

It is not clear how general the techniques behind CULT are -- whether, for example, it could be applied to the translation of other texts -- nor how cost-effective it is in operation. Other factors may justify its continued use. It is also unclear whether R&D is continuing, or whether CULT, like METEO, is unsuited to design modification beyond a certain point already reached. In the absence of answers to these questions, and perhaps despite them, CULT does appear to be an MAT success story: the amount of post-editing said to be required is trivial -- limited to the re-introduction of certain untranslatable formulas, figures, etc., into the translated output. At some point, other translator intervention is required, but it seems to be limited to the manual inflection of verbs and nouns for tense and number, and perhaps the introduction of a few function words such as English determiners.

ALPS - Automated Language Processing Systems

ALPS was incorporated by another group of Brigham Young University workers, around 1979; while the group forming Weidner was composed mostly of the programmers interested in producing a fully-automatic MT system, the group forming ALPS (reusing the old BYU acronym) was composed mostly of linguists interested in producing machine aids for human translators (dictionary look-up and substitution, etc.) [Melby and Tenney, personal communication]. Thus the ALPS system is interactive in all respects, and does not seriously pretend to perform translation at all; rather, ALPS provides the translator with a set of software

tools to automate many of the tasks encountered in everyday translation experience. ALPS adopted the tools originally developed at BYU -- and hence, the language pairs the BYU system had supported: English into French, German, Portuguese, and Spanish. Since then, other languages (e.g., Arabic) have been announced, but their commercial status is unclear.

The ALPS system is intended to work on any of three "levels" -- providing capabilities from simple dictionary lookup on demand to word-for-word (actually, term-for-term) translation and substitution into the target text. The central tool provided by ALPS is a menu-driven word-processing system coupled to the on-line dictionary. One of the first ALPS customers seems to have been Agnew TechTran -- a commercial translation bureau which acquired the ALPS system for in-house use. Recently, another change of ownership and consequent shake-up at Weidner communication Systems, Inc., has allowed ALPS to hire a large group of former Weidner workers, leading to speculation that ALPS might itself be intending to enter the MT arena.

Current Research and Development

In addition to the organizations marketing or using existing M(A)T systems, there are several groups engaged in on-going R&D in this area. Operational (i.e., marketed or used) systems have not yet resulted from these efforts, but deliveries are foreseen at various times in the future. We will discuss the major Japanese MT efforts briefly (as if they were unified, in a sense, though for the most part they are actually separate), and then the major U.S. and European MT systems at greater length.

MT R&D in Japan

In 1982 Japan electrified the technological world by widely publicizing their new Fifth Generation project and establishing the Institute for New Generation Computer Technology (ICOT) as its base. Its goal is to leapfrog Western technology and place Japan at the forefront of the digital electronics world in the 1990's. MITI (Japan's Ministry of International Trade and Industry) is the motivating force behind this project, and intends that the goal be achieved through the development and application of highly innovative techniques in both computer architecture and Artificial Intelligence.

Of the research areas to be addressed by the ICOT scientists and engineers, Machine Translation plays a prominent role. Among the western Artificial Intelligentsia, the inclusion of MT seems out of place: AI researchers have been trying (successfully) to ignore all MT work in the two decades since the ALPAC debacle, and almost universally believe that success is impossible in the foreseeable future -- in ignorance of the successful, cost-effective applications already in place. To the Japanese leadership, however, the inclusion of MT is no accident. Foreign language training aside, translation into Japanese is still

one of the primary means by which Japanese researchers acquire information about what their Western competitors are doing, and how they are doing it. Translation out of Japanese is necessary before Japan can export products to its foreign markets, because the customers demand that the manuals and other documentation not be written only in Japanese. The Japanese correctly view translation as necessary to their technological survival, but have found it extremely difficult to accomplish by human means. Accordingly, their government has sponsored MT research for several decades. There has been no rift between AI and MT researchers in Japan, as there has been in the West -- especially in the U.S. MT may even be seen as the key to Japan's acquisition of enough Western technology to train their scientists and engineers, and thus accomplish their Fifth Generation project goals.

Nomura [82] numbers the MT R&D groups in Japan at more than eighteen. (By contrast, there might be a dozen significant MT groups in all of the U.S. and Europe, including commercial vendors.) Several of the Japanese projects are quite large. (By contrast, only one MT project in the western world [EUROTRA] even appears as large, but most of the 80 individuals involved work on EUROTRA only a fraction of their time.) Most of the Japanese projects are engaged in research as much as development. (Most Western projects are engaged in development.) Japanese progress in MT has not come fast: until a few years ago, their hardware technology was inferior; so was their software competence, but this situation has been changing rapidly. Another obstacle has been the great differences between Japanese and Western languages -- especially English, which is of greatest interest to them -- and the relative paucity of knowledge about these differences. The Japanese are working to eliminate this ignorance: progress has been made, and production-quality systems already exist for some applications. None of the Japanese MT systems are "direct," and all engage in "global" analysis; most are based on a transfer approach, but a few groups are pursuing the interlingua approach.

MT research has been pursued at Kyoto University since 1968. There are now two MT projects at Kyoto (one for near-term application, one for long-term research). The former has developed a practical system for translating English titles of scientific and technical papers into Japanese [Nagao, 80, 82], and is working on other applications of English-Japanese [Tsujii, 82] as well as Japanese-English [Nagao, 81]. The other group at Kyoto is working on an English-Japanese translation system based on formal semantics (Cresswell's simplified version of Montague Grammar [Nishida et al., 82, 83]). Kyushu University has been the home of MT research since 1955, with projects by Tamachi and Shudo [74]. The University of Osaka Prefecture and Fukuoka University also host MT projects.

However, most Japanese MT research (like other research) is performed in the industrial laboratories. Fujitsu [Sawai et al., 82], Hitachi, Toshiba [Amano, 82], and NEC [Muraki & Ichiyama,

82], among others, support large projects generally concentrating on the translation of computer manuals. Nippon Telegraph and Telephone is working on a system to translate scientific and technical articles from Japanese into English and vice versa [Nomura et al., 82], and is looking into the future as far as simultaneous machine translation of telephone conversations [Nomura, personal communication].

The Japanese industrialists are not confining their attention to work at home. Several AI/MT groups in the U.S. (e.g., SRI, U. Texas) have been approached by Japanese companies desiring to fund MT R&D projects. More than that, some U.S. MT vendors (SYSTRAN and Weidner, at least) have recently sold partial interests to Japanese investors. Various Japanese corporations (e.g., NTT and Hitachi) and trade groups (e.g., JEIDA [Japan Electronic Industry Development Association]) have sent teams to visit MT projects around the world and assess the state of the art. University researchers have been given sabbaticals to work at Western MT centers (e.g., Shudo at Texas, Tsujii at Grenoble). Other representatives have indicated Japan's desire to participate in the CEC's EUROTRA project [Margaret King, personal communication]. Japan evidences a long-term, growing commitment to acquire and develop MT technology. The Japanese leadership is convinced that success in MT is vital to their future.

METAL

Of the major MT R&D groups around the world, it would appear that the new METAL project at the Linguistics Research Center of the University of Texas is closest to delivering a product. The METAL German-English system passed tests in a production-style setting in late 1982, mid-83, and early 1984, and the system has been installed at the sponsor's site in Germany for further testing and final development of a translator interface. The METAL dictionaries are being expanded for maximum possible coverage of selected technical areas in anticipation of production use in 1984. Commercial introduction is also a possibility. Work on other language pairs has begun: English-German is now underway, and Spanish and Chinese are in the target language design stage.

One of the particular strengths of the METAL system is its accommodation of a variety of linguistic theories/strategies. The German analysis component is based on a context-free phrase-structure grammar, augmented by procedures with facilities for, among other things, arbitrary transformations. The English analysis component, on the other hand, employs a modified GPSG approach and makes no use of transformations. Analysis is completely separated from transfer, and the system is multi-lingual in that a given constituent structure analysis can be used for transfer and synthesis into multiple target languages. Experimental translation of English into Chinese (in addition to German) will soon be underway; translation from both English and German into Spanish is expected to begin in the immediate future.

The transfer component of METAL includes two transformation packages, one used by transfer grammar rules and the other by transfer dictionary entries; these co-operate during transfer, which is effected during a top-down exploration of the [highest-scoring] tree produced in the analysis phase. The strategy for the top-down pass is controlled by the linguist who writes the transfer rules; these in turn are paired 1-1 with the grammar rules used to perform the original analysis, so that there is no need to search through a general transfer grammar to find applicable rules (potentially allowing application of the wrong ones). As implied above, structural and lexical transfer are performed in the same pass, so that each may influence the operation of the other; in particular, transfer dictionary entries may specify the syntactic and/or semantic contexts in which they are valid. If no analysis is achieved for a given input, the longest phrases which together span that input are selected for independent transfer and synthesis, so that every input (a sentence, or perhaps a phrase) results in some translation.

In addition to producing a translation system per se, the Texas group has developed software packages for text processing (so as to format the output translations like the original input documents), data base management (of dictionary entries and grammar rules), rule validation (to eliminate most errors in dictionary entries and grammar rules), dictionary construction (to enhance human efficiency in coding lexical entries), etc. Aside from the word-processing front-end (being developed by Siemens, the project sponsor), the METAL group is developing a complete system, rather than a basic machine translation engine that leaves much drudgery for its human developers/users. Lehmann et al. [81], Bennett [82], and Slocum [83, 84] present more details about the METAL system.

GETA

As discussed earlier, the Groupe d'Études pour la Traduction Automatique was formed when Grenoble abandoned the CETA system. In reaction to the failures of the interlingua approach, GETA adopted the transfer approach. In addition, the former software design was largely discarded, and a new software package supporting a new style of processing was substituted. The core of GETA is composed of three programs: one converts strings into trees (for, e.g., word analysis), one converts trees into trees (for, e.g., syntactic analysis and transfer), and the third converts trees into strings (for, e.g., word synthesis). The overall translation process is composed of a sequence of stages, wherein each stage employs one of these three programs.

One of the features of GETA that sets it apart from other MT systems is the insistence on the part of the designers that no stage be more powerful than is minimally necessary for its proper function. Thus, rather than supplying the linguist with programming tools capable of performing any operation whatever (e.g., the arbitrarily powerful Q-systems of TAUM), GETA supplies at each stage

only the minimum capability necessary to effect the desired linguistic operation, and no more. This reduces the likelihood that the linguist will become overly ambitious and create unnecessary problems, and also enabled the programmers to produce software that runs more rapidly than would be possible with a more general scheme.

A "grammar" in GETA is actually a network of subgrammars; that is, a grammar is a graph specifying alternative sequences of applications of the subgrammars and optional choices of which subgrammars are to be applied (at all). The top-level grammar is therefore a "control graph" over the subgrammars which actually effect the linguistic operations -- analysis, transfer, etc. GETA is sufficiently general to allow implementation of any linguistic theory, or even multiple theories at once (in separate subgrammars) if such is desired. Thus, in principle, GETA is completely open-ended and could accommodate arbitrary semantic processing and reference to 'world models' of any description.

In practice, however, the story is more complicated. In order to increase the computational flexibility, as is required to take advantage of substantially new linguistic theories, especially 'world models', the underlying software would have to be changed in many various ways. Unfortunately, it is written in IBM assembly language, making modification extremely difficult. Worse, the programmers who wrote the software have long since left the GETA project, and the current staff is unable to safely attempt significant modification. As a result, there has been no substantive change to the GETA software since 1975, and the GETA group has been unable to experiment with any new computational strategies. Back-up, for example, is a known problem [Tsujii, personal communication]: if the GETA system "pursues a wrong path" through the control graph of subgrammars, it can undo some of its work by backing up past whole graphs, discarding the results produced by entire subgrammars; but within a subgrammar, there is no possibility of backing up and reversing the effects of individual rule applications. The GETA workers would like to experiment with such a facility, but are unable to change the software to allow this. Until GETA receives enough funding that new programmers can be hired to rewrite the software in a high-level language, facilitating present and future redesign, the GETA group is "stuck" with the current software, now 10 years old and showing clear signs of age, to say nothing of non-transportability.

GETA seems not to have been pressed to produce an application early on, and the staff was relatively "free" to pursue research interests. Until GETA can be updated, and in the process freed from dependence on IBM mainframes, it may never be a viable system. The project staff are actively seeking funding for such a project. Meanwhile, the French government has launched an application effort through the GETA group.

The University of the Saar at Saarbruecken, West Germany, hosts one of the larger MT projects in Europe, established in the late 1960's. After the failure of a project intended to modify GAT for Russian-German translation, a new system was designed along somewhat similar lines to translate Russian into German after "global" sentence analysis into dependency tree structures, using the transfer approach. Unlike most other MT projects, the Saarbruecken group was left relatively free to pursue research interests, rather than forced to produce applications, and was also funded at a level sufficient to permit significant on-going experimentation and modification. As a result, SUSY tended to track external developments in MT and AI more closely than other projects. For example, Saarbruecken helped establish the co-operative MT group LEIBNIZ (along with Grenoble and others) in 1974, and adopted design ideas from the GETA system. Until 1975, SUSY was based on a strict transfer approach; since 1976, however, it has evolved, becoming more abstract as linguistic problems mandating "deeper" analysis have forced the transfer representations to assume some of the generality of an interlingua. Also as a result of such research freedom, there was apparently no sustained attempt to develop coverage for specific applications.

Intended as a multi-lingual system involving English, French, German and Russian, work on SUSY has concentrated on translation into German from Russian and, recently, English. Thus, the extent to which SUSY may be capable of multilingual translation has not yet been ascertained. Then, too, some aspects of the software are surprisingly primitive: only very recently, for example, did the morphological analysis program become nondeterministic (i.e., general enough to permit lexical ambiguity). The strongest limiting factor in the further development of SUSY seems to be related to the initial inspiration behind the project: SUSY adopted a primitive approach in which the linguistic rules were organized into independent strata, and were incorporated directly into the software [Maas, 84]. As a consequence, the rules were virtually unreadable, and their interactions, eventually, became almost impossible to manage. In terms of application potential, therefore, SUSY seems to have failed. A second-generation project, SUSY-II, begun in 1981, may fare better.

EUROTRA

EUROTRA is the largest MT project in the Western world. It is the first serious attempt to produce a true multi-lingual system, in this case intended for all seven European Economic Community languages. The justification for the project is simple, inescapable economics: over a third of the entire administrative budget of the EEC for 1982 was needed to pay the translation division (average individual income: \$43,000/year), which still could not keep up with the demands placed on it; technical translation costs the EEC \$.20 per word for each of six translations (from the seventh

original language), and doubles the cost of the technology documented; with the addition of Spain and Portugal later this decade, the translation staff would have to double for the current demand level (unless highly productive machine aids were already in place) [Perusse, 83]. The high cost of writing SYSTRAN dictionary entries is presently justifiable for reasons of speed in translation, but this situation is not viable in the long term. The EEC must have superior quality MT at lower cost for dictionary work. Human translation alone will never suffice.

EUROTRA is a true multi-national development project. There is no central laboratory where the work will take place, but instead designated University representatives of each member country will produce the analysis and synthesis modules for their native language; only the transfer modules will be built by a "central" group -- and the transfer modules are designed to be as small as possible, consisting of little more than lexical substitution [King, 82]. Software development will be almost entirely separated from the linguistic rule development; indeed, the production software, though designed by the EUROTRA members, will be written by whichever commercial software house wins the contract in bidding competition. Several co-ordinating committees are working with the various language and emphasis groups to insure co-operation.

The linguistic basis of EUROTRA is nothing novel. The basic structures for representing 'meaning' are dependency trees, marked with feature-value pairs partly at the discretion of the language groups writing the grammars (anything a group wants, it can add), and partly controlled by mutual agreement among the language groups (a certain set of feature-value combinations has been agreed to constitute minimum information; all are constrained to produce this set when analyzing sentences in their language, and all may expect it to be present when synthesizing sentences in their language) [King, 81, 82].

The software basis of EUROTRA will not be novel either, though the design is not yet complete. The basic rule interpreter will be "a general re-write system with a control language over grammars/processes" [King, personal communication]. As in GETA, the linguistic rules will be bundled into packets of subgrammars, and the linguists will be provided with a means of controlling which packets of rules are applied, and when; the individual rules will be non-destructive re-write rules, so that the application of any given rule may create new structure, but will never erase any old information (no back-up).

EUROTRA will engage in straightforward development using state-of-the-art but "proven" techniques. The charter requires delivery of a small representative prototype system by late 1987, and a prototype covering one technical area by late 1988. EUROTRA is required to translate among the native languages of all member countries which sign the "contract of association" by early mid-84; thus, not all seven EEC languages will necessarily be

represented, but by law at least four languages must be represented if the project is to continue.

The State of the Art

Human languages are, by nature, different. So much so, that the illusory goal of abstract perfection in translation -- once and still imagined by some to be achievable -- can be comfortably ruled out of the realm of possible existence, 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 evaluation of translation according to its purpose, judged by the consumer. One must therefore accept the truth that the notion of quality is inherently subjective. Certainly there will be translations hailed by most if not all as "good," and correspondingly there will be translations almost universally labelled "bad." Most translations, however, will surely fall in between these extremes, and each user must render his own judgement 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 judgement which the corporate body is willing to accept at face value: after all, such judgement is the very purpose for having senior translators. It is arrived at subjectively, based on the purpose for which the translation is intended, but 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 original translation.

Therefore, any rational position on the state of the art in MT & MAT must respect the operational decisions about the quality of MT & MAT as judged by the present users. These systems are all, of course, based on old technology ("ancient," by the standards of AI 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). Some translators resist their introduction -- there are those who still resist the introduction of typewriters, to say nothing of word processors -- with varying degrees of success. But 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. Current MT systems seem to take some getting used to, and further productivity increases are realized as time goes by; they are usually accepted, eventually, as a boon to the bored translator. New products embodying old technology are constantly 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 -- else, presumably, they would no longer exist.

Development Systems

Systems being developed for near-term introduction employ Computational Linguistics (CL) techniques of the late 1970's, if not the 80's. Essentially all are full MT, not MAT, systems. As Hutchins [82] 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 substitution 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. Most "pure AI" techniques do not fall in this category; thus, 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 time, and stand ready to deliver commercial applications (e.g., GETA, METAL).

Research Systems

The biggest problem associated with MT research systems is their scarcity (nonexistence, in the U.S.). 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 U.S. they certainly are not, and in Europe, CL and AI research has not yet reached the level achieved in the U.S. Western business and industry are 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.) Some European governments fund significant R&D projects (e.g., Germany and France), but Japan is making by far the world's largest investment in MT research. The U.S. government, which otherwise supports the best overall AI and [English] CL 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 have "understood correctly," and in any case will tell 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 them 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 there existed a sufficient supply of qualified technical translators even in 1966, contrary to ALPAC's claims.) The classic "law of supply and demand" has not worked in this instance, for whatever reasons: the shortage is real, all over the world; nothing is yet serving to stem this worsening situation; and nothing seems capable of doing so outside of dramatic productivity increases via computer automation. In the EEC, for example, the already overwhelming load of technical translation is projected to rise sixfold within five years.

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 -- in technical translation. 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 translation will grow even more than it would otherwise: 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, and the trend 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: Japanese, Chinese, Russian, and even German, for example, seem more difficult to deal with using existing CL techniques -- surely in part due to the nearly total concentration of CL workers on English. Developing translation ability 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. On the one hand, there is the serious question of whether, in for example the case of an article on front-line research in semiconductor switching theory, or nuclear physics, a translator really does "fully comprehend" the content of the article he is translating. One would suspect not. (Johnson [83] makes a point of claiming 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 in order to provide the basic multi-lingual analysis and synthesis tools. Translators must be consulted, for they are the experts in translation. None of this will happen by accident; it must result from design.

References

- Amano, S. Machine Translation Project at Toshiba Corporation. Technical note. Toshiba Corporation, R&D Center, Information Systems Laboratory, Kawasaki, Japan, November 1982.
- Bar-Hillel, Y., "Some Reflections on 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.
- Bennett, W. S. The Linguistic Component of METAL. Working paper LRC-82-2, Linguistics Research Center, University of Texas at Austin, July 1982.
- Bostad, D. A., "Quality Control Procedures in Modification of the Air Force Russian-English MT System," in V. Lawson (ed.), *Practical Experience of Machine Translation*. North-Holland, Amsterdam, 1982, pp. 129-133.
- Bruderer, H. F., "The Present State of Machine and Machine-Assisted Translation," in Commission of the European Communities, *Third European Congress on Information Systems and Networks: Overcoming the Language Barrier*, vol. 1. Verlag Dokumentation, Munich, 1977, pp. 529-556.
- Gervais, A., et DG de la Planification, de l'Évaluation et de la Vérification. Rapport final d'évaluation du système pilote de traduction automatique TAUM-AVIATION. Canada, Secrétariat d'État, June 1980.
- Hundt, M. G., "Working with the Weidner Machine-Aided Translation System," in V. Lawson (ed.), *Practical Experience of Machine Translation*. North-Holland, Amsterdam, 1982, pp. 45-51.
- Hutchins, W. J., "Progress in Documentation: Machine Translation and Machine-Aided Translation," *Journal of Documentation* 34, 2, June 1978, pp. 119-159.
- 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.
- Johnson, R. L., "Parsing - an MT Perspective," in K. S. Jones and Y. Wilks (eds.), *Automatic Natural Language Parsing*. Ellis Horwood, Ltd., Chichester, Great Britain, 1983.
- Jordan, S. R., A. F. R. Brown, and F. C. Hutton, "Computerized Russian 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., "Design Characteristics of a Machine Translation System," *Proceedings of the Seventh IJCAI*, Vancouver, B.C., Canada, Aug. 1981, vol. 1, pp. 43-46.
- 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.
- 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. NTIS report AO-97896.
- Loh, S.-C., "Machine Translation: Past, Present, and Future," *ALLC Bulletin* 4, 2, March 1976, pp. 105-114.
- Loh, S.-C., L. Kong, and H.-S. Hung, "Machine Translation of Chinese Mathematical Articles," *ALLC Bulletin* 6, 2, 1978, pp. 111-120.
- Loh, S.-C., and L. Kong, "An Interactive On-Line Machine Translation System (Chinese into English)," in B. M. Snell (ed.), *Translating and the Computer*. North-Holland, Amsterdam, 1979, pp. 135-148.
- Lytle, E. G., D. Packard, D. Gibb, A. K. Melby, and F. H. Billings, "Junction Grammar as a Base for Natural Language Processing," *AJCL* 3, 1975, microfiche 26, pp. 1-77.
- Maas, H.-D., "The MT system SUSY," presented at the ISSCO Tutorial on Machine Translation, Lugano, Switzerland, 2-6 April 1984.

- Melby, A. K. "Multi-level Translation Aids in a Distributed System," Ninth ICCL [COLING 82], Prague, Czechoslovakia, July 1982, pp. 215-220.
- Muraki, K., and S. Ichiyama. An Overview of Machine Translation Project at NEC Corporation. Technical note. NEC Corporation, C & C Systems Research Laboratories, 1982.
- Nagao, M., J. Tsujii, K. Mitamura, H. Hirakawa, and M. Kume, "A Machine Translation System from Japanese into English: Another Perspective of MT Systems," Proceedings of the Eighth ICCL [COLING 80], Tokyo, 1980, pp. 414-423.
- Nagao, M., et al. On English Generation for a Japanese-English Translation System. Technical Report on Natural Language Processing 25. Information Processing of Japan, 1981.
- Nagao, M., J. Tsujii, K. Yada, and T. Kakimoto, "An English Japanese Machine Translation System of the Titles of Scientific and Engineering Papers," Proceedings of the Ninth ICCL [COLING 82], Prague, 5-10 July 1982, pp. 245-252.
- Nishida, F., and S. Takamatsu, "Japanese-English Translation Through Internal Expressions," Proceedings of the Ninth ICCL [COLING 82], Prague, 5-10 July 1982, pp. 271-276.
- Nishida, T., and S. Doshita, "An English-Japanese Machine Translation System Based on Formal Semantics of Natural Language," Proceedings of the Ninth ICCL [COLING 82], Prague, 5-10 July 1982, pp. 277-282.
- Nishida, T., and S. Doshita. An Application of Montague Grammar to English-Japanese Machine Translation. Proceedings of the ACL-NRL Conference on Applied Natural Language Processing, Santa Monica, California, February 1983, pp. 156-165.
- Nomura, H., and A. Shimazu. Machine Translation in Japan. Technical note. Nippon Telegraph and Telephone Public Corporation, Musashino Electrical Communication Laboratory, Tokyo, November 1982.
- Nomura, H., A. Shimazu, H. Iida, Y. Katagiri, Y. Saito, S. Naito, K. Ogura, A. Yokoo, and M. Mikami. Introduction to LUTE (Language Understander, Translator & Editor). Technical note, Musashino Electrical Communication Laboratory, Research Division, Nippon Telegraph and Telephone Public Corporation, Tokyo, November 1982.
- Perusse, D., "Machine Translation," ATA Chronicle 12, 8, 1983, pp. 6-8.
- Pigott, I. M., "Theoretical Options and Practical Limitations of Using Semantics to Solve Problems of Natural Language Analysis and Machine Translation," in M. MacCafferty and K. Gray (eds.), The Analysis of Meaning: Informatics 5. ASLIB, London, 1979, pp. 239-268.
- Pigott, I. M., "The Importance of Feedback from Translators in the Development of High-Quality Machine Translation," in V. Lawson (ed.), Practical Experience of Machine Translation. North-Holland, Amsterdam, 1982, pp. 61-73.
- Ruffino, J. R., "Coping with Machine Translation," in V. Lawson (ed.), Practical Experience of Machine Translation. North-Holland, Amsterdam, 1982, pp. 57-60.
- Sawai, S., H. Fukushima, M. Sugimoto, and N. Ukai, "Knowledge Representation and Machine Translation," Proceedings of the Ninth ICCL [COLING 82], Prague, 5-10 July 1982, pp. 351-356.
- Sereda, S. P., "Practical Experience of Machine Translation," in V. Lawson (ed.), Practical Experience of Machine Translation. North-Holland, Amsterdam, 1982, pp. 119-123.
- Shudo, K., "On Machine Translation from Japanese into English for a Technical Field," Information Processing in Japan 14, 1974, pp. 44-50.
- Sinaiko, H. W., and G. R. Klare, "Further Experiments in Language Translation: A Second Evaluation of the Readability of Computer Translations," ITL 19, 1973, pp. 29-52.
- 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.
- Thouin, B., "The Meteo System," in V. Lawson (ed.), Practical Experience of Machine Translation. North-Holland, Amsterdam, 1982, pp. 39-44.
- Tsujii, J., "The Transfer Phase in an English-Japanese Translation System," Proceedings of the Ninth ICCL [COLING 82], Prague, 5-10 July 1982, pp. 383-390.
- Van Slype, G., "Évaluation du système de traduction automatique SYSTRAN anglais-français, version 1978, de la Commission des communautés Européennes," Babel 25, 3, 1979, pp. 157-162.
- Van Slype, G., "Economic Aspects of Machine Translation," in V. Lawson (ed.), Practical Experience of Machine Translation. North-Holland, Amsterdam, 1982, pp. 79-93.
- Vasconcellos, M., "Management of the Machine Translation Environment," Proceedings of the ASLIB Conference on Translating and the Computer, London, November 1983.
- Wheeler, P., "The Errant Avocado (Approaches to Ambiguity in SYSTRAN Translation)," Newsletter 13, Natural Language Translations Specialist Group, BCS, February 1983.
- Wilks, Y., and LATSEC, Inc. Comparative Translation Quality Analysis. Final report on contract F-33657-77-C-0695. 1978.