

Report of a Symposium on Japanese to English Machine Translation

Held at the National Academy of Sciences
in Washington, D.C.
on December 7, 1989

Office of Japan Affairs
Computer Science and Technology Board
National Research Council

NATIONAL ACADEMY PRESS
Washington, D.C. 1990

Contents

PREFACE

1. MACHINE TRANSLATION: FROM A TRANSLATION TO A COMMUNICATIONS AND INFORMATION CHALLENGE.....	1
CHANGING CONTEXT FOR MACHINE TRANSLATION.....	2
SETTING AND MEETING GOALS FOR MACHINE TRANSLATION DEVELOPMENT.....	3
2. THE COMMERCIAL CHALLENGES	8
CURRENT STATUS OF MACHINE TRANSLATION DEVELOPMENT	8
MARKET PROSPECTS	12
USER NEEDS	15
JAPANESE AND U.S. USERS: CONTRASTING NEEDS.....	18
3. THE TECHNICAL CHALLENGES: APPROACHES TO RESEARCH AND ASSESSMENT	20
DEVELOPMENTAL STRATEGIES AND PROBLEMS.....	21
EVALUATING MACHINE TRANSLATION SYSTEMS.....	25
4. THE RESEARCH AND DEVELOPMENT POLICY AGENDA.....	30
THE SKEPTICS.....	30
THE ADVOCATES	31
WHAT MIGHT BE DONE?	33
CONCLUSION	36

Preface

At the request of the United States Department of Commerce, the Office of Japan Affairs and the Computer Science and Technology Board of the National Research Council convened a Symposium on Japanese to English Machine Translation on December 7, 1989. With the initiative of a steering committee chaired by Dr. Roger Levien of the Xerox Corporation, the symposium was planned to provide a forum where participants from the private sector and other interested parties could discuss the current state of Japanese to English machine translation, market prospects, and user needs, as well as R&D policy options. Recognizing the potential contribution that machine translation and related technologies can make to the operations of companies doing business around the world, to researchers and policymakers who need to know about developments in Japanese science and technology, and to translators whose work could be augmented by the use of machine aids, the symposium was organized to share information and stimulate thinking about how machine translation and related technologies might address the needs of users in the United States.

This report was prepared by staff members from the Office of Japan Affairs in cooperation with the Computer Science and Technology Board to highlight major points in the symposium presentations and discussions. Participants in the symposium and other experts provided valuable comments and clarifications. This report, however, incorporates a range of views articulated at the symposium rather than a consensus of those presented.

Machine Translation: From a Translation to a Communications and Information Challenge

Developmental work on machine translation has been under way for more than 30 years. Some say that we have come a long way, while others question whether the goal is in sight. The reality is that perceptions of what machine translation is—its purpose and scope—are shifting. The paradigm that dominated thinking during most of this 30-year period of research and development (R&D) was the expectation that computer technologies could be developed to build machines capable of doing the work of human translators. Today, machine translation—defined as translation generated by a computer, with or without human intervention—embraces a broad spectrum of technologies. Included are machine translation systems that run on large mainframes and those that run on stand-alone personal computers, enhanced with automatic aids for the human translator.¹ Rather than eliminating human translators, machine translation and related technologies are now seen as ways of facilitating their work.²

The focus of the symposium organized by the Office of Japan Affairs and the Computer Science and Technology Board was machine translation from Japanese to English. Growing interest in machine translation between this language pair

¹ This report deals with computer-based translation, including machine translation systems and machine aids for translators.

² The concept includes machine translation systems that are used primarily for information scanning, where little or no post-editing is done. Such a system can in some circumstances serve as a mechanical translator for a monolingual researcher. Some of the commercial systems available today are beginning to see this kind of usage in limited circumstances.

reflects recognition of the importance of science and technology information produced in Japan. In order to understand the special challenges of Japanese to English machine translation, a comparatively new field of R&D, it is important to appreciate efforts in machine translation involving other languages.

CHANGING CONTEXT FOR MACHINE TRANSLATION

Machine translation technology development has taken on broader significance in an age of rapid international communication and intense market competition. Competition for global markets has intensified the need for companies to get their messages across to overseas customers who speak foreign languages. Companies doing business around the world must be able to speak the languages of their customers. Some large companies have targeted translation technologies as a component of their competitive strategies: IBM sees translation as a "gating" item³ for its marketing objectives in the 1990s; Xerox emphasizes the importance of machine translation in launching products simultaneously in multiple markets. As the importance of global markets grows and cycle times shorten for the introduction of new products, translation increasingly becomes an expensive bottleneck for international companies.

Another, related explanation for changes in perspectives on machine translation is the information explosion. More specifically, U.S. businessmen, researchers and product developers, and policymakers need a better understanding of what is going on in Japan, because that country has now joined the ranks of the world's technological leaders. The need for translations of Japanese technical documents is now apparent, but only a minute fraction of our science and technology community can speak and read Japanese. Growing recognition of the importance of technical information produced in Japan has stimulated interest in the role that machine translation might play in making it possible for Americans to access reports of new inventions, products, and financial developments in Japan.

Changes in thinking about machine translation also reflect the evolution of new concepts of how machine translation systems might be developed and used. Progress in natural language processing technology, the development of more powerful computers, the increasing availability of large, information-laden dictionary data sets, and advances in some aspects of linguistic theory suggest opportunities for R&D. Translations can be delivered through electronic mail and quickly incorporated in successive editions of technical manuals. Interest in "machine-aided translation" has been spurred by the development of workstations

³ Translation of documents is a prerequisite for entering overseas markets.

with dictionaries and other tools that can be used by professional translators. The expectation that translation machines will replace people has now been transformed into the view that these technologies are tools to enhance the efforts of professional translators, researchers, and secretaries.

Today the challenges of machine translation development illustrate the broader challenges of information technology research, development, and use. Machine translation technologies pose a range of theoretical, software, hardware, and even sociological problems that require integration of technologies and improved interactions among developers and users. For all these reasons, machine translation today is more than a linguistics problem. It is a communications and information challenge that demands a diverse range of expertise and resources.

SETTING AND MEETING GOALS FOR MACHINE TRANSLATION DEVELOPMENT

The dream that stimulated early R&D efforts was a machine that produces high-quality translations from a wide variety of source texts at low cost. Even the most ardent supporters of machine translation agree that three decades of effort have not produced the breakthroughs necessary to achieve this dream.

Why the dream remains unfulfilled is the subject of some disagreement. Viewed from one angle, the failure to achieve the goal is a result of giving up too early. Negative evaluations of machine translation in the 1960s were based on the argument that the understanding of text by computer was too difficult, rendering machine translation infeasible.⁴ The ALPAC report by the National Research Council concluded that the basic technology for machine translation had not been developed, and recommended a focus on long-term research in computational linguistics and improvement of translation methods. While the report made no recommendations with regard to funding for research and development on machine translation, the overall negative evaluation of the state-of-the-art is now seen by many as a major cause of the subsequent decline in funding for such research in the United States. Between 1960 and 1970, funding for R&D on machine translation declined from about \$10 million to \$1 million.

⁴ This argument was made by Bar Hillel. The ALPAC report written by the Automatic Language Processing Advisory Committee of the National Research Council, entitled *Language and Machines: Computers in Translation and Linguistics*, is widely seen as the most influential of the studies of machine translation in the 1960s. Published in 1966, the report concluded that the quality of machine translation was poor and cost savings had not been achieved. The report analyzed the products of machine translation at U.S. government agencies after development work had been under way for more than ten years.

Supporters of machine translation say that we would be closer to the dream if we had not given up so soon.⁵

Viewed from another perspective, however, the fact that the initial dream has not been fulfilled is no reason to dismiss the promise of machine translation. The machine translation "heaven" of high-quality, low-cost, general-purpose systems is still distant.⁶ The high-quality systems that exist today are in most cases special-purpose systems working in restricted domains, but none of these are for the translation of Japanese to English. (See Figure 1.) In addition, there are also a number of cost-effective systems that operate in broader domains. Among them, Systran is the only company to have developed a general-purpose commercial system that translates between Japanese and English, as well as 14 other language pairs. A dozen or more prototype machine translation systems have not been able to attain cost effectiveness after a decade of development.

The process of machine translating from Japanese to English and vice versa is comparatively difficult because of important differences in the structures of the two languages. Typical Japanese text consists of Chinese characters and two different styles of Japanese phonetic symbols. These characters and symbols are written without any spaces between individual words, and phrases are rarely separated by punctuation. Grammatically, Japanese differs from English in that it has no distinction of singular and plural nouns, there are no articles, and the subjects are often omitted. In general, the Japanese sentence puts the verb at the end, and the text preceding the verb is in no particular order and contains many compound clauses. The grammatical structure of Japanese may omit pronouns, subjects, and objects, so that the context must be understood in order to choose among alternative possible interpretations. Because of these and many other characteristics of the language, pre-editing is especially helpful to make the text more tractable for machine translation processing.

Most of the Japanese to English systems now in use in Japan succeed because they are limited to particular domains. In the eyes of many, these systems represent a significant step forward, even if they do not fulfill the initial dream.

⁵ In July 1989 the Japan Electronic Industry Development Association published a report entitled: *A Japanese View of Machine Translation in Light of the Considerations and Recommendations Reported by ALPAC, U.S.A.* This report argues that two major conclusions of the ALPAC report are no longer valid: the claim that there is no translation shortage is refuted by estimates of today's translation market in Japan, and numerous examples of successful machine translation are also cited in response to ALPAC's conclusion that it will have no practical use in the near future. The Japan Electronic Industry Development Association's Machine Translation System Research Committee, which prepared the report, was chaired by Dr. Makoto Nagao of Kyoto University and included representatives from Japanese corporations involved in machine translation development.

⁶ Some argue that the major explanation for the failure to reach the target is that there is a much bigger difference between general language and domain-specific language than has heretofore been suspected.

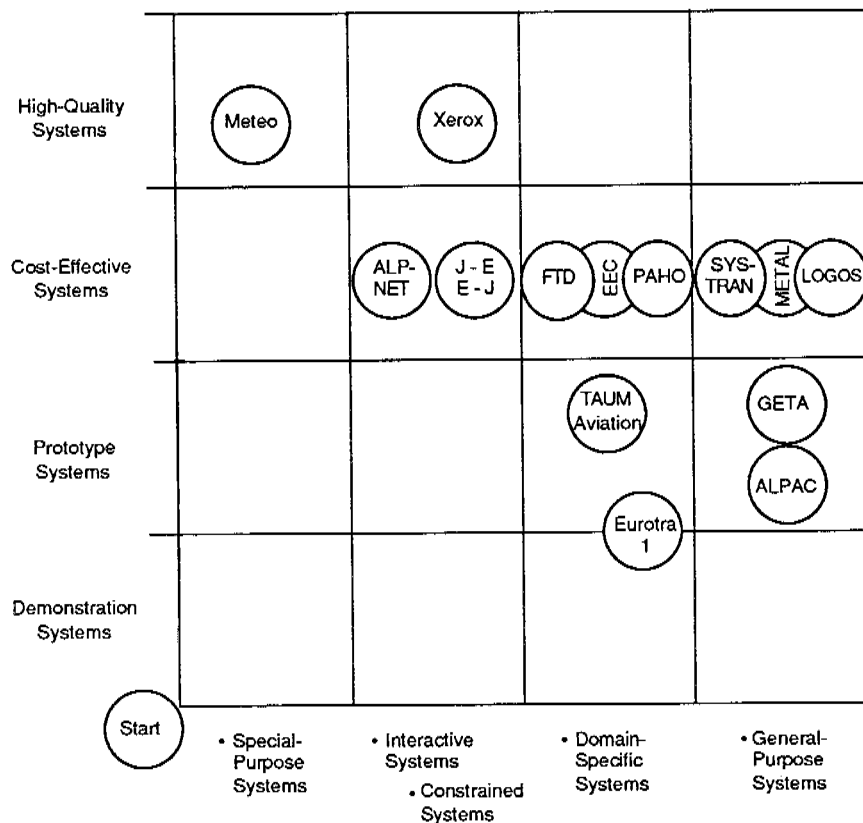


FIGURE 1 Machine translation quality/function matrix. SOURCE: Bernard Scott, Logos Corporation. See legend on page 6.

The argument may be rephrased as follows: Japanese to English machine translation systems can be used effectively for carefully targeted purposes and it will only be possible to improve these systems if we are willing to put resources into development and experimentation. When we also consider the growth of machine aids for human translators, such as dictionaries and other composition tools, it is clear that machine translation technologies have practical uses today, even if the dream of a fully automated, high-quality, low-cost, general-purpose system remains over the horizon.

Given this picture of promise and problems associated with machine translation, we need to examine the challenges that lie ahead for U.S. industry and the U.S. government. One set of challenges is commercial and relates to the fact that while a number of Japanese companies are working on Japanese to English machine translation systems, little similar work is going on in the United

LEGEND TO FIGURE 1

ALPAC	Russian-English machine translation systems being developed under U. S. Department of Defense funding prior to the publication of the ALPAC Report in 1966, which led to their discontinuation.
ALPNET	Interactive multilingual translation aid, developed in the United States. Previously known as ALPS for Automated Language Processing System.
EEC	European Economic Community.
EUROTRA I	Large-scale multilingual machine translation prototype effort sponsored by the European Community. Eurotra I is intended to lead to a full-scale industrialized version, Eurotra II, encompassing 72 language pairs.
FTD	Russian-English and German-English SYSTRAN-based machine translation systems of the Foreign Technology Division of the U.S. Air Force. FTD's installation in the late 1960's represented the first operational use of machine translation.
GETA	Machine translation system for various language pairs entailing French, developed at the University of Grenoble, France. (GETA - Groupe d'Etudes pour la Traduction Automatique.)
LOGOS	Multilingual machine translation system developed and marketed by Logos Corporation, a U.S. company. Language pairs are: German-English, -French, -Italian; English-French, -Spanish, -German, and -Italian.
METAL	German-English, English-German machine translation system originally developed at the Linguistic Research Center of the University of Texas and later under joint development with Siemens AG in Munich.
METEO	English-French machine translation system developed for the nationwide weather communications network of the Canadian Meteorological Center. METEO is a derivative of the TAUM system developed at the University of Montréal.
PAHO	English-Spanish/Spanish-English machine translation systems of the Pan American Health Organization designated as SPANAM and ENGSPAN, respectively.
SYSTRAN	Machine translation system for many language pairs, developed in the U.S. and controlled by French and Japanese interests. SYSTRAN is a much enhanced derivative of the early Georgetown Automatic Translation (GAT) system.
TAUM Aviation	An English-French adaptation of TAUM technology to a specific subject matter domain or language subset, developed at the University of Montréal. (TAUM - Traduction Automatique de l'Université de Montréal.)
J-E/E-J	Japanese-English/English-Japanese machine translation systems. There are about a dozen such systems in Japan, either on the market or being prepared for the market.
XEROX	English-multitarget machine translation system developed for XEROX Corporation by the developers of SYSTRAN, and used by XEROX to translate technical manuals written in controlled English.

States. A second set of challenges includes technical problems—issues relating to the choice of R&D focus and problems in evaluating system performance. The third set of challenges may prove to be the most pressing: defining R&D policy (either at the company or the U.S. government level). Without a clear consensus on what constitutes the “state-of-the-art,” or comprehensive data on market prospects and user needs, forging an appropriate policy response is not an easy task.

The sections that follow address each of these three sets of challenges in turn, highlighting areas of uncertainty and issues that deserve further study and debate.

The Commercial Challenges

CURRENT STATUS OF MACHINE TRANSLATION DEVELOPMENT

How, where, and by whom are machine translation technologies and machine aids for translation being developed? Who is using them and for what purposes? Before examining Japanese to English machine translation as it is being developed and used in the United States and Japan, it is helpful to review the broader context of machine translation development throughout the world.

Machine translation development has advanced rapidly in the past few years, and while numerous challenges remain, machine-assisted translation is no longer a dream, but is actively and increasingly used around the world. Companies and governments are developing and/or using machine translation technologies in the Americas, Europe, Japan, and in the newly industrializing countries of Asia. Machine translation and machine-aided translation are now being used in organizations as diverse as translation bureaus, multinational corporations, and government defense departments. These technologies help to scan information, collect intelligence, translate product manuals for export, and improve the efficiency of professional translators. The overview below provides a basis for understanding the organizational, political, financial, and societal reasons for differences in their development and use in the United States and Japan.

The major systems in use today are not, however, aiding translation from Japanese to English; most are performing what is considered by many the easier task of translating from one Western language to another. Many examples can be found in the public or government sector. The Pan American Health

Organization, for example, uses its internally developed system to translate health and agricultural information between English and Spanish. The U.S. government, including the Air Force's Foreign Technology Division, has long used Systran to translate Russian, German, and French into English for intelligence purposes. The government is currently supporting Systran's development of a Japanese language system. Similarly, the Canadian government uses the Logos system for translating English to French at the Departments of National Defense and State, among others, and uses METEO to translate weather forecasts. The Smart Translator is used to announce job openings. The European Community is supporting a large initiative aimed at developing a system to translate among all the official languages of the Community.

In the private sector, manufacturing companies are using machine translation technology to produce product manuals and translation companies are using it to improve the volume and efficiency of translation efforts. Xerox, for example, uses the Systran machine translation system to translate its photocopier product manuals for distribution throughout the world. A European computer maker, Nixdorf, uses the Logos system to achieve the same end. Translation companies such as Lexitech in Canada also use Logos, and ALPNET has developed a worldwide network of translation services based on its own internally developed interactive and machine-aided software. METAL has a number of clients in Europe and ATAMIRI is used by Wang for the translation of product manuals into several languages.

There are a few examples of machine translation technology used for Asian languages. Logos, one of the oldest U.S. machine translation companies, got its start in 1971 when it developed an English to Vietnamese system for the U.S. Defense Department. Currently, the Defense Department is funding machine translation from and into Korean.

It can be seen from the examples above that the world of machine translation spans the globe and includes efforts in many language pairs for a variety of uses. Some of the oldest, most successful, and well-established machine translation companies are American or are based on technology developed in the United States. Systran is perhaps the grandfather of U.S. machine translation companies and its efforts in language pairs other than those involving Japanese continue to be U.S.-based. The next oldest commercial machine translation system is that of Logos, also an American company. One of the best-known Japanese machine translation companies, Bravice International, which claims to have sold 4,500 software units, bought its technology in the United States through the acquisition of a U.S. firm. It is useful to consider this context as one examines the more specific case of Japanese to English machine translation.

Who, then, are the developers of machine or machine-aided translation between Japanese and English? The major commercial efforts in the United States are being conducted by a handful of companies. As mentioned above,

Systran, with U.S. government support, is supporting the Japanese-English combination; it should be noted, however, that Systran's commercial Japanese efforts have been Japanese-owned. Some large U.S. corporations, including IBM and DEC, are pursuing internal research in this field, but their Japanese to English efforts are located in Japan.⁷ IBM's English to Japanese machine translation system, developed for internal use in the corporation's Tokyo facilities, is currently in operational use for the translation of computer manuals.

There are some small-scale efforts, experiments in assembling and using machine translation technologies for very limited domains. Examples are Smart Communications and a small ongoing effort at the Veterans Administration Medical Center in Baltimore that uses public domain software to translate a narrow range of medical documents between Japanese and English. Efforts in machine-aided translation development are currently underway at ALPNET and LinguaTech, both of which are based in Utah and originated at Brigham Young University. Logos, which continues to invest heavily in machine translation development, has conducted research in Japanese but has never undertaken development of a Japanese system.

There is also considerable theoretical research under way among computational linguists at U.S. universities, such as work at Carnegie Mellon University's Center for Machine Translation, and the Linguistics Research Center at the University of Texas, and the sub-language approach pursued cooperatively by Hunter, Monmouth, and N.Y.U. If basic research on natural language processing and computational linguistics is taken into account, the United States still maintains a significant research effort.⁸ What is lacking in the United States is a strong development effort on Japanese to English machine translation.

The world of Japanese to English machine translation development in Japan offers a sharp contrast to that in the United States. Every major Japanese computer or electronics firm has invested considerable effort in machine translation research and development and many claim to have introduced workable systems. Without evaluating quality, it is nonetheless significant that there are at least twenty operational⁹ systems in Japan that translate from English to Japanese. Operational systems in Japan that translate from Japanese to English

⁷ At its Tokyo Research Lab, IBM is working on a Japanese to English machine translation system for use in translating newspaper editorials and economic materials, which is now in a research prototype stage.

⁸ It should be noted, however, that research in computational linguistics that is not related to machine translation will not necessarily contribute to machine translation. On the other hand, the case can be made that work on machine translation serves as a test bed and stimulus for other kinds of natural language processing investigation.

⁹ There is considerable ambiguity about what constitutes an "operational" or "usable" system. The sheer number of Japanese developers who even claim to have an operational system is nonetheless an indication of the comparatively strong interest and resources devoted to machine translation in Japan.

are available from NEC, Fujitsu, Oki Electric, Bravice International, Sharp, Toshiba, Hitachi, and Sanyo Electric to mention a few of the companies. Japanese to English systems are also under development at NTT, Mitsubishi Electric, KDD, and Toshiba.

In addition, the Japanese government has supported an important effort in machine translation development. This effort, which involves the Ministry of International Trade and Industry's (MITI) Electrotechnical Lab, the Science and Technology Agency's (STA) Japan Information Center of Science and Technology (JICST), and Kyoto University, was started in 1982. The Electronic Dictionary Research Project conducted by MITI in connection with the Fifth Generation Computer Project aims at the development of a detailed dictionary with more than 200,000 words and multiple usages.¹⁰ Supported by the Ministry of Post and Telecommunications, research is underway at ATR on speech translation telephony.¹¹ Technology transfer to industry has been made possible throughout these projects via industrial participation. The projects feed into the effort at JICST for translating science and technology abstracts. These projects are developmental vehicles that spin off nationwide results; their continuity and commercial emphasis build capability in the companies.

Who, then, is using Japanese to English machine translation? If the typical pattern is for translation to be done in the country of the target language, then Japanese to English machine translation should be done most efficiently in the United States, rather than Japan. At the present time, however, this is not the case. Most Japanese to English machine translation is conducted in Japan. The fact that many Japanese computer makers have developed their own machine translation capability reflects their orientation toward product exports and their need to control the quality of translated manuals. It also reflects an understanding of the importance of machine translation technology and its possible spin-offs to the information industry as a whole.

Japanese-developed Japanese to English systems are not widely used in the United States. Lack of compatibility between hardware and software is one of the impediments. Microelectronics and Computer Technology Corporation (MCC), a private U.S. microelectronics and computer science consortium, is a relatively new user of Japanese to English machine translation (MCC is using a Japanese-developed system) and has begun to use the technology in a particularly forward looking manner that will be discussed in more detail below. The University of Wisconsin's Biotechnology Center planned to use a Bravice system

¹⁰ The 10-year EDR project that began three years ago includes an English dictionary, a Japanese dictionary, and a neutral dictionary that connects the Japanese and English dictionaries.

¹¹ See Hitoshi Iida, "Advanced Dialogue Translation Techniques," ATR Interpreting Telephony Research Laboratories, ATR Symposium on Basic Research for Telephone Interpretation, December 11-12, 1989.

on an experimental basis to scan and translate Japanese language databases on biotechnology, but has been unable to do so because of the high cost of the necessary hardware. Currently, there is no U.S.-developed Japanese to English machine translation system on the market. The U.S. defense and intelligence community will, in all likelihood, be the first major user of Systran's Japanese to English system when it is fully operational.

MARKET PROSPECTS

Today, the volume of machine-translated documents remains comparatively small. In contrast, some Japanese believe that the annual market for all translated materials in 1988 was about 800 million yen, and that the quantity of translation will double over the 1990-1992 period.¹² One developer even estimates that by the year 2000 there will be 500,000 to 2 million machine translation systems in use throughout the world,¹³ assuming that substantial improvements are made in intermediate systems that can run on small personal computers. According to this vision, international businessmen will need small computers with built-in machine translation systems.

At the same time, Japanese experts note that large Japanese companies working on machine translation do not believe that this business will yield great profits, at least in the short term. They do, however, see machine translation as a mechanism for learning more about natural language processing technology in general, which they judge to be a key technology in the next century.

Despite the interest in machine translation technology, profits in the United States and Europe are very slim, if there are any at all. While there are companies developing systems for internal use, the independent developers are (as noted above) few in number. This contrast with the situation in Japan may be explained in a number of ways. Critics of machine translation argue that the products are based on dated technology. They argue that translators harbor serious doubts about machine translation on quality grounds, even if translators have a hard time quantifying the concept of quality. Translation, particularly in Europe, is poorly integrated with office and publishing computer environments where the potential benefits of machine translation could be substantial.

From a Japanese viewpoint, however, it is possible to create demand. Viewed from this angle, the more machine translation systems are made available and put into practical use, the more the demand will grow. Japanese industry and government are willing to plow large investments into Japanese to English

¹² See Japan Electronic Industry Association, *A Japanese View of Machine Translation*. . . , op. cit., p. 5 and Appendix 5.

¹³ Many experts see this as a very high estimate.

machine translation. It is estimated that a four- to five-year effort involving 50 to 70 people is needed to develop a general-purpose mainframe system for delivery, at a cost of \$13 million for the entire period.¹⁴ This is, however, only the beginning of the investment that is necessary. After delivery of the system, considerable resources must be invested in maintaining and improving it in response to user complaints and needs. In short, the required investments are so large that most companies find it impossible to recover costs by selling only hundreds of systems. This is a primary reason for Japanese government support not only of researchers, but also of commercial developers, although there are other important reasons such as reducing the language barrier between Japan and other countries and disseminating Japanese technical information worldwide.

Given that the major commercial machine translation systems have been mainframe-based, the traditional markets for machine translation have been limited to translation bureaus, multinational corporations, and intelligence and information gatherers, particularly in government. Developers have identified two major targets for development. One is the large-scale, general-purpose, mainframe-based system for use by big companies and governments. The goal of this type of machine translation system using large-scale hardware is high volume translation of documents, sometimes for mass distribution. Even those involved in development work on general-purpose, publication quality machine translation systems say that such high-quality, general-purpose systems are 10 to 20 years over the horizon.

The second is a small-scale machine translation system for use by small to medium size companies and even individual researchers. The purpose of this type of smaller scale system is to translate for specialized applications. Many Japanese commercial developers are focusing on small-scale system development. There is a third kind of development that is intermediate between the two systems mentioned above.¹⁵ From a commercial perspective, it is this kind of intermediate system development that is seen by the Japanese as the only feasible target for at least the next decade. Such systems are in use today for information scanning in restricted domains.

In developing and enhancing machine translation systems, developers stand to benefit from close interaction with users. Japanese developers rarely consult with users on the details of systems design, but they do seek out users' views on what features are needed in pre- and post-editing and on general issues of man-machine interface. Japanese developers prefer to interact intensively with a limited number of users so that they can respond effectively to their needs. A

¹⁴ Makoto Nagao, written response to "stimulation questions" prepared for the Symposium on Japanese to English Machine Translation.

¹⁵ Systems like those at Xerox run on microcomputers and are specialized, but cannot be considered "small-scale" since they translate more than 10 million words annually.

conscious strategy of selling a limited number of systems is often pursued in order to facilitate this process. New users benefit later from this accumulated feedback embodied in system configurations developed and perfected for other users. Although Japanese observers complain that mechanisms for exchange of information between developers and users are inadequate, in comparison to the situation in the United States there has been closer interaction among the research, development, and user communities in Japan.¹⁶

One of the limiting factors on the machine translation market is the fact that most Japanese systems operate on only one type of hardware. In fact, some large mainframe computer makers pursue machine translation development as a strategy for marketing their hardware to large companies. Bravice produces the only Japanese to English machine translation system that can be used on a variety of hardware. It runs on most small microcomputers sold in Japan. A practical barrier to widespread usage of small-scale Japanese-developed machine translation systems is the limited availability of hardware that supports Japanese characters (*kanji*, *hiragana*, *katakana*) outside of Japan.

International competition in machine translation is not mature. Given the barriers to hardware interoperability, the most prominent examples of competition occur among large-scale communication systems. In 1988, a number of computer companies competed for contracts associated with the development of a large communication system for the Korean and U.S. armies and some machine translation system developers received contracts.

Reflecting the practical limitations on machine translation technology today, there are a number of efforts that focus on combining machine translation technology and human translation. ALPNET has developed a strategy based on the assessment that machine translation is not the solution to all user needs, and that it is not an all-or-nothing alternative. Technology tools now in existence (such as multilingual word processors, dictionary look-up systems, character recognition and word processing systems) can be effectively used by human translators at much less cost than fully automated machine translation systems. New "linguistic engine" options can be added to a user's existing applications, thereby increasing ease of use without forcing the customer to learn new user interfaces to complex systems. The experiences of developers like ALPNET illustrate the fact that successful application of machine translation and related technologies depends on an understanding of user needs, establishment of expectation levels, matching the tools to the job, and the training of skilled professional people to use the tools.

Another possible approach is to redesign systems using new technology and locate smaller, modular systems in places such as schools where translators are

¹⁶ Exceptions to this general statement about interactions in the United States include relationships between developers and users such as Systran and the U.S. Air Force or between Logos and AT&T.

being trained. Also, a broader appreciation of machine translation and related technologies to cover a wide range of communications problems could enhance the use of currently existing tools in solving practical problems. For example, many potential users in Europe are not professional translators but secretaries working for large multinational corporations who must write correspondence in foreign languages. Market prospects could be broadened by integrating linguistic tools into the office environment, the publishing business, and even the engineering development environment.

If we consider particular applications of machine-aided translation technologies to specific consumer products, demand may indeed grow significantly over the next few years. A composition aid for technical writers, for example, could improve understanding of cultural nuances, not to mention grammar and syntax. Another product application that illustrates the point is a hand-held electronic phrase book for use by foreign travelers. Such applications, it should be noted, offer promise, but they remain distant from the machine translation heaven of general-purpose, high-quality systems.

USER NEEDS

People who want to read translated material care little whether the translation is done by people or machines or a combination. They simply want reliable translation that is cheap and fast. For some of these users, such as those who must scan very large volumes of information in order to follow trends in research and development or identify documents for full translation, the speed of translation is important. Turnaround time is, in the experience of the users of the Air Force's system, more important than quality, defined in terms of naturalness of expression and precision in conveying meaning.¹⁷

Other large volume users, such as companies that require translation of their operating manuals, require precise and understandable translation—but in fairly limited domains. The classic example of such a user in the United States is Xerox. Xerox has been using machine translation for more than a decade to translate technical documents. These are large volume projects which involve high reproduction and updating rates. For large companies that sell a wide range of products in global markets, translation costs represent a significant component of new product expenditures.

The Xerox experience deserves further mention. Xerox has developed an approach that integrates desktop editing programs and a Systran machine translation system to produce service manuals, training programs, and operator

¹⁷ Some argue that quality standards must include a mixture of turnaround time, accuracy, and readability, the relative importance of each varying with the needs of the particular user.

manuals for use in Europe, Latin America, and North America. This approach, which has evolved over time, involves the training of technical writers and pre-editors in a set of simple writing tools (Multinational Customized English) developed by Xerox. Dictionary building is constantly under way as technical writers and translators from different parts of the world, joined through the Xerox worldwide network, submit new words and phrases.

At Xerox, machine translation has shown concrete results. Significant improvements have been seen in productivity per finished page of translated text. Producing more than 40,000 pages annually of documentation translated through this process, Xerox now finds that translation is no longer a barrier to product launch. Machine translation enhances the company's ability to introduce new products almost simultaneously all over the world.

Xerox found that management initially had high expectations for what machine translation could do, while linguists were skeptical. Over time and with growing experience, the perceptions of these two groups have begun to converge. (See Figure 2.) By gradually developing a system that translators are able to see as a tool in their work, and one that produces high-volume product in restricted technical domains, machine translation has been fully integrated with business operations.

In contrast to large users, there are countless potential users of small machine translation systems who have different kinds of needs. In many instances, these are researchers (scientists and technical personnel) in the United States who need to know about developments in Japanese science and technology. For many of them, readability and smoothness of translation is less important than timeliness and access to a narrowly targeted range of technical documents of interest.




FIGURE 2 Expectations versus reality of computer translation. SOURCE: Maria Russo.

MCC is now experimenting with machine translation to monitor overseas technology developments and provide translated materials to member firms. For an investment of \$50,000, MCC has assembled a machine translation system that includes an optical character reader, a Japanese-made personal computer that is also used as a JICST terminal, a Japanese to English machine translation software package, and a U.S.-developed workstation. MCC encountered problems in connecting the computer systems developed in Japan and the United States and in getting adequate vendor support for the machine translation system developed in Japan. Fifty percent of the sentences in the trial output have been judged to be accurate. For a modest investment (\$50,000 to \$60,000), the machine translation experiment at this U.S. industry consortium is considered to be worth the effort.¹⁸ MCC's International Liaison Office expects to augment its already successfully established functions of monitoring developments in Japan and providing technical support to its researchers by using machine translation first to translate titles and, eventually, abstracts and short papers. MCC has, in addition, recently initiated a research effort on knowledge-based natural language processing targeted toward machine translation technology development.

Beyond corporations, individual professional translators and researchers also qualify as potential users of machine translation systems and machine aids for translators. Companies such as LinguaTech are developing workstations for translators who can work at home using a microcomputer with a high-capacity disk drive, a printer, and a port for telecommunications. A translator can receive source texts and glossaries with specific subject matter along telephone lines. A variety of options are available to the translator, including composition tools, a bilingual corpus that permits text retrievals when necessary, multilevel terminology files, and optional use of machine translated texts. To use a baking analogy, machine translation is one ingredient in a range of elements that the translator (baker) can combine in myriad ways to accomplish his work. The baker may wish to use the heavy-duty bread mixer (machine translated text) for some tasks, but he also keeps his drawer of spoons (manual tools).¹⁹

Potential users, particularly individuals or small companies, may find it hard to learn about machine translation. In Japan, potential users learn about machine translation through newspapers and television as well as through demonstrations at computer company service centers. Only limited information is obtained through these channels. The potential Japanese user is, however, in a much better

¹⁸ In judging the value of machine translation, it must be remembered that, unlike the case with European languages with cognates and familiar script, Japanese is totally unintelligible to the potential English-speaking user. Some, therefore, believe that capturing even 50% of the Japanese text can be a useful step forward in communication.

¹⁹ This analogy is developed by Alan K. Melby in his article "Recipe for a Translator Work Station," *Multilingua*, 3-4 (1984), pp. 225-228.

position than his U.S. counterpart. For the U.S. translation community, demonstrations and meetings of professional organizations provide channels for transmitting information that augments the articles appearing in specialist journals. For other users who have no direct contact with the translation community (such as researchers and people in small business), acquiring information about machine translation and making judgments about purchase and use are even more difficult. We shall return to this point again in the final section on R&D policy.

JAPANESE AND U.S. USERS: CONTRASTING NEEDS

As we think about what type of Japanese to English machine translation system would help users in the United States, it is important to keep in mind that their needs are different from those of most users in Japan. The profile of the typical user of Japanese to English machine translation in Japan is a large corporation engaged in global exports. This user can "control" the source document because it is in his native language and pre-editing is easier in this case. This user can customize the system to suit his or her needs and tolerate marginal machine translation because there is a direct cost justification for this effort.

The profile of the typical user of Japanese to English machine translation in the United States is someone who needs expanded access to Japanese-language technical information but who has no fluency in Japanese. This user has no control over the source text and little ability to pre-edit. For monolingual users, the "raw" output must be more reliable and accurate than for the typical Japanese user. For the user in the United States, there are many problems associated with gaining knowledge of and access to Japanese databases, inputting text, and post-editing requirements. In view of the broad-based needs of the U.S. user, marginal machine translation is in many cases not useful. (See Figure 3.) However, there are many possible uses of machine translated text now available, particularly if some post-editing is done, for scanning information.

Japanese to English machine translation systems developed in Japan do not, in the opinion of some leading U.S. developers, fit the needs of many potential high-volume U.S. users. The requirements are for high-quality, broad-domain systems which will be based on new technology.

What are our choices? Should we wait for Japan to develop high-quality Japanese to English machine translation systems in the year 2000 and later? After all, Japan has a head start and Japanese is their language not ours. Or should we develop Japanese to English machine translation systems here in the United States? Are there other alternatives? Answering these questions requires an understanding of the technical challenges facing those engaged in R&D on machine translation.

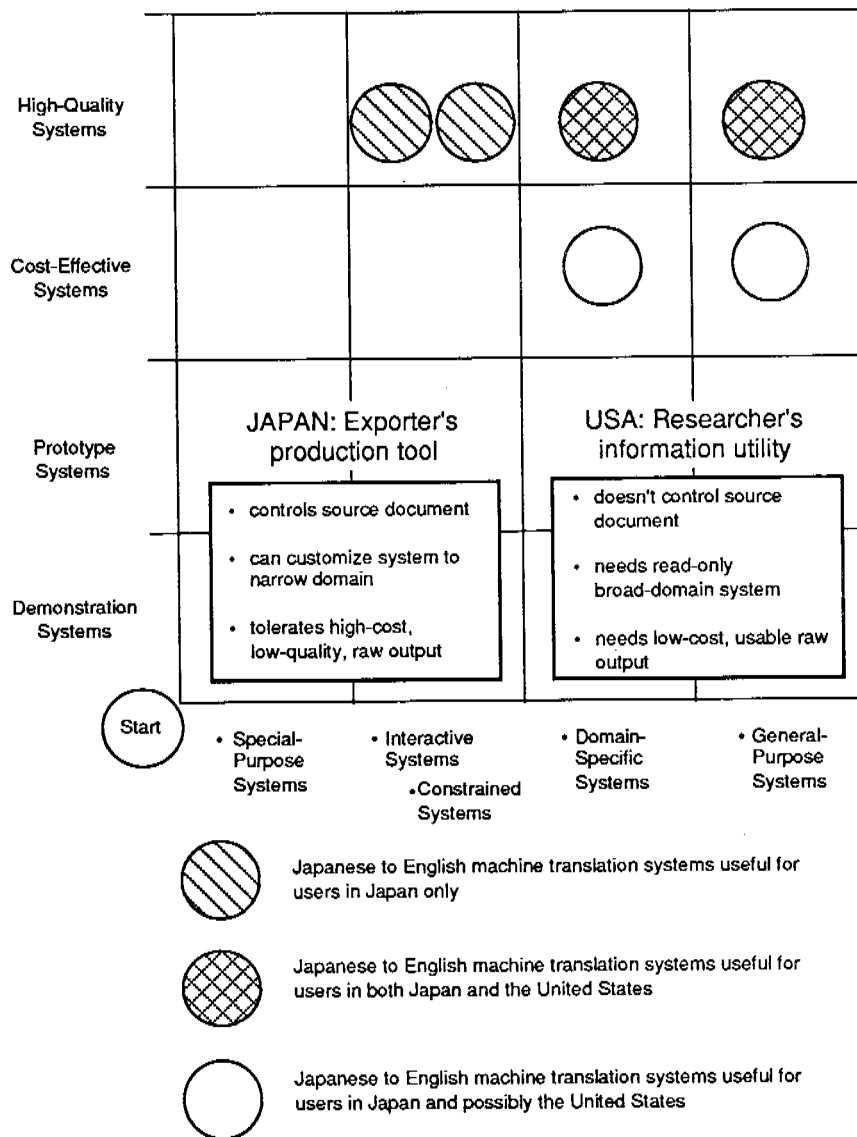


FIGURE 3 User profiles for Japanese-English machine translation. SOURCE: Bernard Scott, Logos Corporation.

3

The Technical Challenges: Approaches to Research and Assessment

Dramatic improvements in computer hardware and software have contributed to progress in machine translation. System performance, which can be measured in “raw LIPS” or logical inferences per second, is now doubling every one to two years.²⁰ Despite, or perhaps because of, these rapid advances in computer technology the barriers in linguistic theory and other areas have become ever more apparent. The result is a receding horizon—as strides are made in R&D it is clear that much more needs to be done.

Over more than 30 years of research and development, work on machine translation has taken three approaches to the process of translation. But there are a variety of machine translation systems in use today and new advances in technology have ushered in systems using intermediate language representations and artificial intelligence that enable the computer to “learn” a language. The discussion that follows briefly reviews the current strategies for developing machine translation systems, key problems for research and development, and issues in the evaluation of machine translation systems.

²⁰ Logical inferences per second is a measure of the speed at which the system: 1) recognizes a pattern match between an element of the input text and a previously stored pattern, and 2) applies the rule that goes with that pattern to translate the element.

DEVELOPMENTAL STRATEGIES AND PROBLEMS

Regardless of which languages are translated, there are now three primary translation strategies for machine translation.²¹ The direct translation method deals only with single language pairs and translates words directly from one language to another. Used for most of the earliest systems, this method involves very little or no linguistic analysis and produces very rough translations. Although this strategy is not designed to handle translations of complete documents, it has been used for machine translation of large databases, tables of contents, and titles of technical publications. The International Liaison Office of MCC is using this strategy for Japanese to English translation in order to develop databases on scientific developments in Japan for its customers who require an overview of available materials in particular fields. It is expected that the users will select documents for full translation by other means.

The transfer system, which operates in three stages, is the most widely used strategy for machine translation. The source language is first analyzed and converted into representations that can be transposed into sentence structures through semantic analysis. In the second stage, the source language representations are converted to the target language transfer representations. The final process synthesizes the transfer language representation into the text of the target language. This method is the one most widely used in Japanese to English machine translation systems, although the transfer system works best when the language pairs are closely related. In some mainframe systems such as experimental systems ATHENE/N developed by Hitachi and FAI by Mitsubishi, the transfer system is enhanced by the use of case analysis.²²

The pivot method is the third translation strategy. It is based on the ideas that language is a universal human experience and that a universal interlingua can be developed, which can be understood by a machine. This method is designed to convert the source language into the interlingua which is then converted into the target language. The interlingua today is still largely a theoretical concept, although the Logos system makes use of an interlingua in a hybrid interlingual/transfer architecture. Researchers working on the interlingua expect that the application of artificial intelligence will permit significant advances to be made.²³

²¹ This overview of machine translation strategies is drawn from a paper by Wayne Kiyosaki, "Machine Translation: Time for a Reappraisal," forthcoming.

²² "Japanese Machine Translation Systems Described," Tokyo NIKKEI ELECTRONICS in Japanese, February 1986, pp.137-168.

²³ For a detailed analysis of the strengths and weaknesses of these three strategies, see W. John Hutchins, "Recent Developments in Machine Translation," *New Directions in Machine Translation*, Conference Proceedings, Budapest, August 1988.

The systems that have been developed use various approaches; they can be compared to tools in a tool chest. No one tool is always best but in some cases one tool may be better than another.²⁴

Using primarily direct translation or translation by the transfer system, there are three possible ways in which human intervention can occur. Pre-editing can involve two kinds of operations. In one case, a text is revised to eliminate structural or lexical ambiguities before being translated by a computer. In the past this approach was not widely used, due to the difficulty in anticipating structures or words that will be difficult for a computer to handle. More recently, with the introduction of text-critiquing software, the potential ambiguities can be brought to the attention of a human translator automatically.

In a second approach to pre-editing, the input text is produced especially for the machine. In some cases it is a new version of an existing text, in others an entirely new text. Multinational Customized English is an example of a restricted English developed by Xerox for use on its Systran system. In some cases, pre-editing is almost as difficult as traditional translation. The efficacy of pre-editing depends to a great extent on the human editor knowing the limitations of the machine translation system.

In the case of interactive editing, the computer calls on the editor to make choices among various alternatives in order to resolve ambiguities. It is also possible to combine post-editing with interactive editing. However, this can make the process costly. The first interactive systems were introduced in the early 1980s and the interactive editing method seems to have gained wider acceptance in recent years.

Of the three editing options, post-editing is clearly the most widely used. Usually a professional translator, the post-editor corrects the machine's output. This is more efficient when done directly on the screen using appropriate word processing software. If the post-editor writes the corrections on hard copy and they are then entered into the computer, the process is much slower. Some estimate that an experienced post-editor can produce 4,000 to 8,000 words a day and in some cases as many as 10,000.²⁵

Since human intervention is costly, the goal of some developmental efforts is fully automatic operation of a machine translation system.²⁶ When the application involves merely gleaning the "gist" of the text, some of the large, general-purpose systems are used on a fully automatic basis. If a more careful translation is needed, output can be post-edited. Such systems include general-

²⁴ Observation by Jaime Carbonell, Carnegie Mellon University.

²⁵ This discussion of editing options is based on work by Muriel Vasconcellos.

²⁶ The Center for Machine Translation at Carnegie Mellon University is working to improve the quality of machine translation output through the incorporation of knowledge bases, especially for applications in limited domains.

purpose systems that are able to handle a wide variety of source texts and special-purpose systems designed to translate a special type of source text such as weather reports or abstracts of technical articles in particular fields.

In addition to full machine translation systems, there are many related technologies that are used as translation aids. These include on-line dictionaries, grammar checkers, and libraries of phrases that are regularly used by human translators. These systems are updated and developed as post-editors contribute to the on-line dictionaries and users give input to improve the way that the machine translation system does the actual translation.

Research and technical challenges particularly relevant to Japanese to English machine translation include the problem of inputting the text, which includes Chinese characters as well as two phonetic scripts. Optical character readers will help to solve the problem of text input, but there are still many difficulties associated with input of Japanese text because of different character fonts and the placement of charts, graphs, and tables in the text. Optical character readers are now being coupled with machine translation systems in Japan, but the extent to which they increase savings over manual input is not clear.

A major research and development question relates to the problem of pre-editing. The better the source text (the clearer the expression and the shorter the sentences), the better the resulting machine translated text and the less post-editing needed. But, as mentioned above, pre-editing is time consuming and tedious work that requires special skills.

While significant advances have been made in computational linguistics, there remain problems that must be overcome in order to build linguistic theory and develop more sophisticated machine translation systems. This set of challenges could be approached in a step-by-step fashion, as some Japanese experts suggest. Research in the following areas is needed: the introduction of priority information in order to disambiguate several possible sentence structures and words; the development of learning mechanisms that produce preference values for the disambiguations; the establishment of grammatical rules that consider many more than two elements simultaneously; improved capabilities for dealing with such problems as anaphora resolution, ellipsis, and the analysis of sentence fragments.

Although machine translation strategies and system types are more or less universal, the ways in which the researchers in the United States and Japan approach these subjects are quite different. As one observer put it: Americans write papers; Japanese build machines.²⁷ The Japanese approach has been more

²⁷ This distinction should be considered carefully. Some question the notion that Japanese researchers are not theory-oriented: one leading Japanese researcher believes that (instead) they focus on second and third approximations required for machine processing of natural language—less beautiful and less academic but useful theory. On the other hand, one critic of Japan's machine translation says that the machines that the Japanese build do not really do the job and (therefore) their approach is not practical.

pragmatic and oriented toward experimenting with systems. This involves a problem-solving approach to linguistic analysis. The parts of the language that do not fit neatly into linguistic theory models are approached by combining different theories or by accumulating individual facts to deal with specific problem areas.

In contrast, the U.S. research community has concentrated more on machine translation theory than on applications. Much of the researcher's time is devoted to writing papers and developing new models of natural language. As a result, critics argue that U.S. researchers construct models that are elegant but not amenable to practical use. At the same time, we should remember that the strides that have been made in basic computational linguistics, a research approach recommended by the ALPAC report, make today's machine translation systems possible.

The theoretical work that has been done in the United States and other countries, including Japan, has made machine translation developers and users aware of the research challenges that are present. These include the need for a bilingual text corpus and the development of automatic comparison algorithms for this corpus. The automatic collection of special terminology words and construction of a thesaurus of these terms would improve many machine translation systems. Standardizing dictionary theory and practice, proper analysis of broken utterances, improved grammar checking devices, and automated approaches to the detection and resolution of ambiguities are other important research themes. Even linguistic and cognitive studies of pre- and post-editors' behavior have been suggested as avenues to improved machine translation.

All of this requires an increase in the number of researchers working on machine translation as well as more basic research themes. Some practical steps might be taken to make experimental tools for natural language processing and machine translation easily available to researchers. These might include the construction of a portable software package for natural language processing, and its distribution to interested researchers; establishment of core grammars for English and Japanese that are linguistically sound, and their distribution to interested researchers; the construction of a text database that includes bilingual text data for use in natural language processing.

In the United States, where the thrust of research has been in more theoretical areas, there is a need to improve interactions with those who take an "engineering" and applications-oriented approach if commercialization of useful systems is the objective. As noted above, interaction with users is essential to system development. These and other questions central to R&D policy in the United States will be explored more fully in Section 4.

EVALUATING MACHINE TRANSLATION SYSTEMS

Corporations involved in development, researchers working on fundamental technologies, potential users, and government policymakers all need to know how good machine translation systems are in order to make choices. Unfortunately, there is no generally accepted method for evaluating the quality and accuracy of translations by people, or by machines.

Japanese developers of machine translation systems often say that the systems are 80% acceptable. This general score is, however, more an intuitive judgment than the result of systematic research. It was pointed out that if 20% of the cookies in the cookie jar are poisoned, no one will want to eat any of them.

Overall assessments of machine translation are less useful than evaluations of specific systems because the evaluation depends very much on the needs of a specific user. Japanese developers note that in some cases a reasonably accurate or even a rough translation may be appropriate, while in other cases where high levels of accuracy are essential, machine translation is unacceptable. A researcher who needs to comb through a vast mountain of information may find rough translations of abstracts very useful in tracking overall trends in research or in selecting articles for full translation. Nothing less than absolute accuracy in translation will satisfy a lawyer working on a legal brief or a politician whose words are quoted by the media. The machine translation systems now in operation, particularly the prototype Japanese to English systems, have been developed to translate technical documents, manuals, and information in restricted domains.

Participants in the Japanese machine translation project supported by the Science and Technology Agency of Japan developed an approach to evaluation using two independent indicators: intelligibility (the extent to which the translation can be understood by a native speaker of the target language) and accuracy (the degree to which the translated text conveys the meaning of the original).²⁸ Samples of machine translated sentences were evaluated by the researchers as roughly 80% acceptable. This overall evaluation was based on the result that 80% of the sentences were given a score of at least 3 in intelligibility and accuracy.²⁹ It is estimated that 20 to 30% of the output sentences in Japanese to English machine translation systems are unacceptable, and in those cases post-editing cannot be carried out effectively.

²⁸ See Makoto Nagao, Junichi Tsuji, and Junichi Nakamura, "Machine Translation from Japanese into English," *Proceedings of the IEEE*, vol. 74, July 1986.

²⁹ A score of 3 in intelligibility was given to sentences whose meaning was clear, but where the evaluator was not sure of some word and grammar usage. A score of 3 in accuracy was given to sentences where the content of the input sentence was generally conveyed in the output sentence, but where there were problems with tense, voice, etc. *Ibid.*, p. 1006.

While no commercially available system can do it, some Japanese to English systems now in use by researchers in Japan reportedly can identify inaccurate text. Leaders in Japanese to English machine translation research, however, note that no accurate data are available to judge particular systems and that the assessments of accuracy and intelligibility are not based on rigorous testing.³⁰

Nor are there unambiguous cost evaluations of machine translation systems, although developers contend that the time taken and cost are generally less than with pure human translation. Here, again, the conclusions drawn about the relative cost of machine translation depend on the type of text and the purpose of the user. According to Japanese expert reports, the best Japanese machine translation systems are cost effective. In one example, a page of text can be translated in 40 minutes when post-editing is done on hard copy, while human translation requires about 43 minutes per page. The charge for machine translation is about 75% the amount for human translation in this particular instance.³¹

The more a user uses a machine translation system, the more efficient the work. It takes at least one year and usually two years for a user to become really familiar with a system and for cost efficiencies to become apparent. (See Figures 4 and 5.) It appears that a significant volume of text must be translated in order to achieve such "learning curve" benefits. The more carefully selected the text (with short sentences and well tuned content consistent with the parameters of the system), the more apparent the cost efficiencies over time. (See Figures 6 and 7.)

Unfortunately, evaluations of machine translation systems currently depend on subjective judgments as to what constitutes acceptable levels of cost and accuracy. In many respects, beauty is in the eye of the beholder. What may be unacceptable text to one user may be usable to another. A major obstacle to the development of machine translation systems is the reluctance of some involved in development to provide detailed information about performance characteristics and to exchange information about their experiences. Developers anxious to convince potential funders of research and users of the systems have oversold their systems, resulting in frustration. Potential users are well advised to conduct systematic comparisons of system performance on sample texts of their own selection that are typical of the application envisaged. In order to facilitate research and development, it will be

³⁰ It should be noted that the ratings are carried out by the developers and reflect evaluations of carefully "tuned" texts appropriate to the system.

³¹ See Japan Electronic Industry Development Association, *A Japanese View of Machine Translation*. . . , op. cit., p. 12. This utilization example involves machine translation of a technical text "tuned" to the system. See also Appendix 9 of the report, *Examples of Machine Translation Use in Japan*. One participant in the symposium reports that better results for machine translation as compared to human translation from Japanese to English were recently reported at a conference in Munich.




FIGURE 4 Developer's effort to improve. SOURCE: Data collected by a major Japanese firm involved in machine translation development.




FIGURE 5 User's effort to improve. SOURCE: Data collected by a major Japanese firm involved in machine translation development.

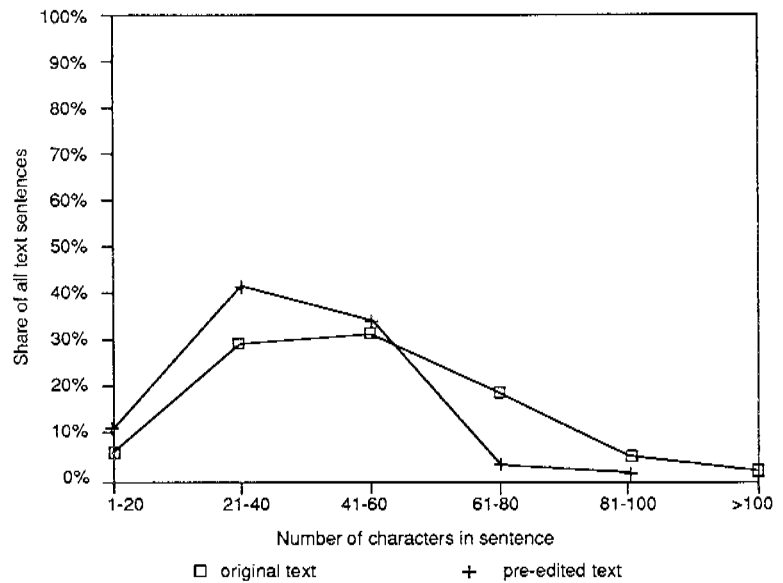


FIGURE 6 Length of sentences in text. SOURCE: Data collected by a major Japanese firm involved in machine translation development.

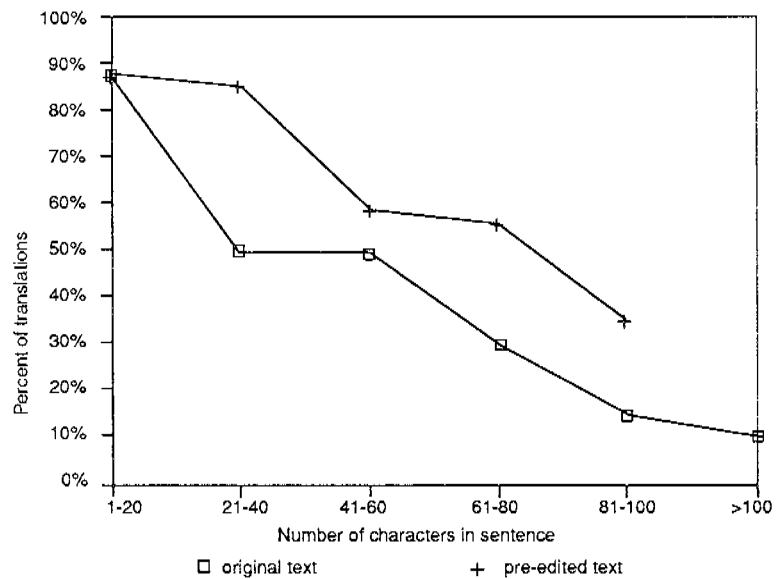


FIGURE 7 Accuracy of translation (by length of sentence). SOURCE: Data collected by a major Japanese firm involved in machine translation development.

necessary to improve techniques for evaluating system performance and timely exchange of information about new developments.³²

³² One participant in the symposium expressed doubt, based on experience of the past 20 years, that reliable methodologies for evaluating machine translation systems can be developed. A comparison of parsers under controlled conditions was suggested as a possibility.

4

The Research and Development Policy Agenda

The discussion above highlighted the special needs of U.S. users of Japanese to English machine translation for high-quality, general-purpose systems and the research and technical challenges in developing such systems. What should be done and who should do it? The sections that follow review a range of perspectives on these issues discussed at the symposium and lay out alternative approaches to the formulation of a policy agenda.

THE SKEPTICS

Despite the apparent growing need for translations of Japanese technical documents, and significant efforts by members of Congress and others to expand U.S. government efforts in this area, skeptics question whether there is real demand for such information. A number of commercial ventures to provide translated Japanese technical literature have failed and many question whether the U.S. research and business communities will overcome a “not invented here” syndrome—the assumption that technical information from abroad is second rate. At the same time many suspect that the information being made available is out-of-date and of little technical interest.

A number of observers question whether the United States government should attempt to mount a major initiative on machine translation. The case against doing something, particularly at the U.S. government level, is based on uncertainty about the need to target this set of technologies for special attention.

Despite a decline in U.S. funding for machine translation, the United States maintains a solid research base in machine translation. Skeptics, moreover, question whether Japan has really taken the lead. Instead, they say that the Japanese have rushed prematurely to market machine translation systems that need much more work and would not be attractive to the typical U.S. user. Therefore, the argument goes, there is no justification for a U.S. government effort. If the demand is there and the technical problems are amenable to resolution, they say, U.S. industry should take the lead rather than calling on the government to support costly and uncertain research.³³

Furthermore, support for R&D on machine translation must be seen in broader context. Some believe other technologies deserve more attention and may have quicker commercial payoffs. In an era of tight federal budgets, how can the case be made that machine translation deserves special attention?

Nor is machine translation a panacea for widening U.S. access to Japanese scientific and technological information. It is one piece of a large picture. Efforts to expand Japanese language study, harmonize patent systems, and reward young scientists and engineers who spend time in Japanese laboratories are also important, sometimes competing themes in the national debate. Professional translators, moreover, worry that advocates of machine translation believe that it will be possible to eliminate “the most expensive link in the communication chain”—the skilled bilingual.³⁴

Even those most committed to work on machine translation note the sociological problems that plague this enterprise. Viewed from this angle, no amount of R&D funding will by itself resolve the fragmentation of the research, development, and user communities in the United States. If machine translation is going to be supported, the skeptics say, it should be as part of a larger national commitment to building networks of interaction among researchers, business people, and professional translators.

THE ADVOCATES

In response to those who question whether the demand is “real” for translations of Japanese technical literature, advocates of new policy initiatives note the lack of follow-through in previous initiatives. Some believe that the

³³ Skeptics note that most machine translation systems in operation today are based on logic programming, the expert system approach. Marshall Unger argues that recent breakthroughs in artificial intelligence (those obtained in the neural-net paradigm, also called connectionism) offer the most promising direction for future research and development.

³⁴ As discussed above, it would appear that such concerns are no longer valid in view of changing perspectives on machine translation.

Office of Japanese Technical Literature, created by law in 1986, has been downgraded within the Department of Commerce and given inadequate funding.³⁵ Federal agencies could, using machine translation, do a better job disseminating existing data on Japanese science and technology to the public, research, and business communities alike. Machine translation, used as a gisting tool, does work that is simply beyond the capacity of human translation. If machine translation were developed into a systemized service, the result could be to stimulate new demand for Japanese technical information. If users do not know it exists or that it can be obtained at reasonable or no cost, demand will remain low.

It is not enough to say: let the market be the judge. Markets can fail to meet the national interest, providing a legitimate rationale for government action. In Japan, government and business are cooperating to support a long-term machine translation development effort. They are, in effect, creating a market. Because the needs of users in the United States are different from those in Japan and cannot be fully met by the Japanese to English machine translation systems now in operation, some believe that our needs cannot be met by simply buying from Japan.³⁶

There are broader motivations for advancing support for machine translation in the United States. In Japan and Europe, machine translation is seen as an important component in the new range of information technologies for the 21st century. As a technical challenge, Japanese to English machine translation represents an engineering stretch that may spin off progress in other areas of information technology, one that requires integration with other fields of information technology R&D. Viewed from this perspective, machine translation is a test bed for solving natural language processing problems in information technology. The question, advocates say, is whether the United States can afford not to pursue this technology.

The potential impacts of a focused effort on machine translation extend beyond information technologies. Bringing together researchers, developers, and users in a long-term project would provide an opportunity for experimenting with new approaches to R&D organization. There is a striking fragmentation of machine translation efforts in the United States—ineffective communication between scientists working on basic research and the engineers who apply it; weak links between the developers and the potential users. This lack of coordination makes machine translation a very fragile technology in the United

³⁵ Comments by Congressman George E. Brown Jr., on the "Federal Role in Accessing Japanese Technical Information" at the Symposium on Japanese to English Machine Translation.

³⁶ At least one participant in the symposium, however, argues the case for simply buying from Japan. As long as the technology is readily available to U.S. customers, the argument goes, it makes sense to take full advantage.

States today. There are also deeper issues at stake such as whether we should be satisfied with a situation in which the United States produces the best research papers and the Japanese produce the products.

WHAT MIGHT BE DONE?

In developing a policy agenda for machine translation in the United States, there are a number of levels at which initiatives can be taken. Initiatives could be taken at the international level. Second, a national policy effort could be defined. And third, individuals and organizations could define new approaches. The suggestions outlined below are not mutually exclusive, however, policymakers who see the need for action will also need to determine where scarce resources can best be invested.

International cooperation in machine translation is desirable in order to develop expanded and up-to-date bilingual dictionaries. It also makes sense in view of the large expenditures required and the "precompetitive" nature of some of the more basic research needed to lay a solid foundation for product development. Dr. Makoto Nagao, a leader of Japan's machine translation efforts, has initiated an International Association for Machine Translation that will bring together researchers, developers, and users in North America, the European Community, and Japan. The purpose of the organization would be to collect and compile information on machine translation so that users have a better understanding of what types of systems are available, how they can be used, what the experience of others has been, and improve evaluation of machine translation systems. For the developers, the benefits would include cooperative approaches to dictionary building, database construction, and sharing of texts as well as expected improvements in theories of machine translation and in machine translation systems. Users and developers alike could benefit from the development of standards for input, design of controlled languages, and evaluation. Conceived as a federation of regionally-based organizations, the international organization will be supported by membership dues (including both individual and corporate members) and other sources, including government funding.

In the context of U.S.-Japan relations, cooperation in machine translation has special significance. Machine translation has been identified by both countries as an area for potential cooperation under the United States-Japan Agreement for Cooperation in Science and Technology. Machine translation appears to be well suited for U.S.-Japan cooperation for a number of reasons. First, the precompetitive aspects of the R&D challenge are significant. Second, a long-term effort is needed. Third, the costs are high and the commercial payoffs over the horizon. Finally, the United States could potentially learn from the efforts already under way in Japan, provided that the participants include senior

researchers from Japan's premier corporate laboratories. The Japanese would, of course, benefit from opportunities to interact with the American user community and from collaboration in development of translation tools. Assuming that a U.S.-Japan cooperative project would be thoughtfully focused to ensure that the potential benefits to the United States are maximized, the significance could be broader in demonstrating the value of collaborative R&D efforts more generally.

A compelling case can also be made for a defensive United States strategy where the U.S. government takes the initiative. If consensus is built on the need for a policy initiative, the U.S. government would have to be involved. The small, struggling companies that specialize in machine translation in the United States lack the resources and the motivation to mount a serious effort in Japanese to English machine translation. For the large companies with big R&D budgets, on the other hand, the high risks and uncertain returns of expenditures on machine translation research present significant disincentives. Universities active in this field do not see it as their mission to build commercially viable, general-purpose machine translation systems. (See Figure 8.)

This leaves the U.S. government as the only potential catalyst for a serious, product-oriented effort on Japanese to English machine translation. An important requirement for such an initiative would be the participation of government, industry, and university researchers and strategists. A three-stage approach could be considered. In stage one, a steering committee would be assembled to take an inventory of technology sources and user needs. The major tasks of the steering committee would be to: determine what are the bottlenecks to the development of high-quality machine translation systems; explore the issues of evaluation criteria; develop plans for design of a new type of R&D organization. Stage one could be completed within an 18-month time frame. Stage two of a U.S. initiative, lasting 2-5 years, would focus R&D funding on the development of two prototype systems (Japanese to English; English to Japanese). Stage two efforts would be highly focused on problem solving and the evaluation of technical results. The R&D would be carried out by two competing consortia involving participation by industry, government, and university researchers. If two competing consortia were established, however, the costs would double. Another possibility would be to have two competing consortia, both working on a Japanese to English prototype system. Product-oriented implementation would be the goal of stage three, during the 5-10 year time horizon. During this stage, interim products would be developed by a number of subcontractors who would work on machine translation systems for use in a wide range of domains.

Critical to the success of such an initiative is a recognition that some of the most significant barriers to previous machine translation projects have been organizational. Involving the right mix of researchers who combine strengths in theory as well as engineering problem-solving would be essential. Building mechanisms that integrate developers and potential users at an early stage would also be important. The infrastructure already established by the National

	US Univer- sities	US Machine Translation Firms	US Industry	US Gov't	Japanese Industry *
• motivation				√	√
• resources					
- financial			√	√	√
- organizational			√	√	√
• expertise					
- core machine translation technology	√	√			√
- machine translation production development		√			√
• user-friendly interfaces					
• user feedback					
• develop pitfalls					
• large-scale databases					
• format handling					
- integration	√		√		√
• platforms					
• networks					
• database management systems					
• graphics					
• scanners					
- high-quality MT technology		?			?

* Note contributions are also made by Japanese universities and government but Japanese industry has all of these resources.

FIGURE 8 Sources for J-E development. SOURCE: Bernard Scott.

Technical Information Service for accessing and translating Japanese language documents could serve as a foundation for this effort. Machine translation could be a test bed for a focused, U.S. national initiative in information technology development. Developing procedures for rapid processing of requests to use copyrighted materials is also essential, as well as agreements that encourage timely dissemination and low cost or free usage of jointly developed dictionaries and other tools.

At another level, policy initiatives can also be developed at the organizational or corporate level. Those with long experience in providing Japanese language technical information argue that a commitment is necessary to create and maintain demand at the organizational level. MCC's International Liaison Office,

the information specialists at Bell Laboratories—these are the people and places where demand can be developed internally. Building better ways to organize, disseminate, and utilize Japanese technical literature provides, therefore, the context for machine translation. Building a constituency of users for translated documents means working directly with the users to tailor the search, selection, and output to their needs. Unfortunately, only the largest organizations today have invested the resources needed and only a few have made expanding access to foreign technical literature a high priority.

CONCLUSION

A range of options is available to policymakers (public and private) who determine that initiatives in machine translation should be pursued. These options—international collaboration with the European Community and Japan, a U.S. national project, government support for stronger linkages between university research centers of excellence and potential developers, a focused government procurement strategy designed to build machine translation expertise in the United States—are not mutually exclusive. It is, however, not clear that pursuing an international cooperation strategy would by itself build a fully integrated research and production base in the United States or a new generation of machine translation systems that fit the special needs of the U.S. user community.

Machine translation, broadly defined, offers interesting research and development challenges in building new information technologies. It also offers significant potential rewards to businesses operating in global markets, to researchers and engineers who need to know what is going on in Japan, as well as to translators whose work can be augmented by the use of machine translation tools. Machine translation is not a panacea or a fully developed technology—but it is a dynamic and challenging area worthy of serious policy consideration.