NEW POSSIBILITIES IN MACHINE TRANSLATION

Eduard H. Hovy

Information Sciences Institute of USC

4676 Admiralty Way

Marina del Rey, CA 90292-6695

Tel: 213-822-1511

Email: HOVY@ISI.EDI

ABSTRACT

There is a growing need for language translation of documents in commerce, government, science, and international organizations. At the same time, translation by computer (MT) is reaching the stage where it can deliver significant cost savings (systems are being sold in Japan that reputedly reduce the time required for translation by up to 50%). Although fully automated high-quality translation is technically not feasible today or in the near future, a number of recent theoretical developments make possible MT systems that are more powerful and effective than existing ones. These developments include: better representation techniques, a clearer understanding of semantics for translation, more complete grammars, and better generation and parsing technology. By making optimal use of existing technology, new MT projects can reach a sophisticated level of performance within a short time. This paper provides reasons for starting a new MT program and recommends the establishment of three small MT projects that address the same domain but use different theoretical frameworks.

INTRODUCTION

The possibility of using computers to perform the translation of documents among various languages was one of the earliest goals of Natural Language Processing and, indeed, one of the earliest of Artificial Intelligence. In the typical approach taken in the sixties, a parser program was equipped with a grammar and lexicon of the source language and a generator program with a grammar and lexicon of the target language, and the remainder consisted of a set of rules of correspondences among syntactic structures or lexical items. These approaches were soon proved naive by translations such as the now-famous “the vodka is strong but the meat is rotten” from “the spirit is willing but the flesh is weak”. It was apparent that semantic information had somehow to be taken seriously (at least to the point of knowing that, for example, “spirit” may indeed be “vodka”, but not when used as an active agent who can be “willing”).

Since the early 60's, Machine Translation (MT) as a field of inquiry has largely lain dormant in the U.S., with the exception of a few large projects (such as at the University of Texas, Austin [Bennett 82]) and a few smaller projects (such as at Yale University [Lytinen 84]). In recent years, however, especially under the impetus of Japanese and European efforts at addressing the problem, U.S. interest in MT research has been on the increase.

The principal reason for the increase is the ongoing development of tools and techniques that enable the performance of certain tasks with more thoroughness and success than was possible earlier (see,

for example, [Carbonell et al. 81, Carbonell & Tomita 87, Nirenburg 87, Arnold 86, Nakamura et al. 88, Laubsch et al. 84, Amano 86]). There has been a steady growth of the capabilities of parsers and generators, the coverage of grammars, and the power and sophistication of knowledge representation techniques. In addition, two recent developments have the nature of breakthroughs that will greatly enhance future MT systems: the incorporation of disjunction in KL-ONE-like representation systems, and the development of general-purpose language-based taxonomic ontologies of representation.

Overall, the field has grown wiser since the 60's: the newer MT projects are all less ambitious in scope than the early ones. Though nobody today would promise to deliver a system that performs perfect translation in even a relatively restricted domain, researchers feel comfortable about proposing systems that perform the first pass of a translation, producing a rough copy of the text in the target language, which would then be edited for stylistic smoothness and fluent cadence by a human editor. Such systems are called Machine-Aided Translation (MAT) systems. Since such systems significantly reduce the problems and costs of translation, they are in high demand in industry and industrial research throughout the world. For example, three MAT systems currently in use in Japan reputedly reduce the time of translation of technical documents by about 50%; two of them are commercially available for under \$70,000 [Time 89]. The following passage from the invitation to an international seminar on MT organized by IBM (held in Munich, West Germany, in August 1989) summarizes the point:

There is a growing need for translation (estimated at 15-25 percent per annum) in commerce, science, governments, and international organizations. This is due to increased international cooperation and competition, an ever-growing volume of text to be communicated, often in multiple languages, world-wide electronic communication, and more emphasis in countries on the use of national language in documents and systems. The opening of the European market in 1992 will add significantly to these factors.

At the same time, automated machine translation of natural language is reaching the stage where it can deliver significant cost savings in translation production, and vastly increase the scope of information retrieval, although fully automated high-quality translation is technically not feasible today and in the near future.

[H. Lehmann and P. Newman, IBM Scientific Centers in Heidelberg and Los Angeles, 1989.]

This paper presents a case for the establishment of a modest MAT program under Darpa support. After providing some background and describing new technological capabilities, it discusses a framework in which a few small MAT projects could be brought into existence for a modest investment and motivated toward achieving a high level of performance in five years.

MT SYSTEMS: COMPONENTS AND APPROACHES

THE COMPONENTS OF AN MT SYSTEM

In order to build an MT system, the following program modules or components are needed:

- A Parser
- A Generator

Figure 1: The Modules of an MT or MAT System.

- Grammars for each language
- Lexicons for each language
- A semantic Knowledge Base
- Interlanguage Translation Rules (in systems without an interlingua)

In all MT systems, these modules are related essentially as shown in Figure 1. We briefly discuss each module below.

Parser: Sentences of the source text are parsed into some internal form by the parser. In almost all current MT systems, the internal form represents both syntactic and semantic aspects of the input.

Interlanguage Translation Rules: Many MT systems contain a set of rules that transform certain aspects of the internal representation of the input to make them conform to the requirements of the target language. Such MT systems are known as *transfer-based*. An alternative approach is to build MT systems without transfer rules, using a single intermediate representation form called an *interlingua*; the generality and power of such systems depends on the expressiveness of the interlingua used.

Generator: The (modified) internal representation of the input is generated as sentence(s) of the target language by the generator. The output must express the semantic content of the internal form, and if possible should use syntactic forms equivalent to those present in the input.

Grammars: In some systems, the grammars (syntactic information) are intrinsic parts of the parser and generator; in others, the grammars can be separated from the procedural mechanism. In *bidirectional* systems, the parser and generator use the same grammar to analyze and produce each language. Such systems are desirable because they do not duplicate syntactic information and are therefore more

maintainable. True bidirectional grammars have proven hard to build, not least because existing knowledge representation formalisms do not provide some capabilities (such as inference over disjunction) that facilitate parsing and generation.

Semantic Knowledge Base: All sophisticated MT systems make heavy use of a knowledge base (representing underlying semantic information) containing the ontology of the application domain: the entities and their possible interrelationships. Among other uses, the parser requires these entities to perform semantic disambiguation and the generator uses them to determine acceptable paraphrases where exact 'literal' formulations are not possible.

Lexicons: All MT systems require a lexicon for the source language and one for the target language. In simple systems, corresponding entries in the two lexicons are directly linked to each other; in more sophisticated systems, lexicon entries are either accessed by entities represented in the knowledge base, or are indexed by characteristic collections of features (as built up by the parser).

APPROACHES TO MT

Using these basic modules, a number of different approaches to the problem of MT are possible.

The Lexical Approach: Many of the early MT systems, as well as some existing projects, base their approach on the lexicon to a large extent. Typically, in such systems one finds a proliferation of highly specific translation rules spread throughout the lexicon; in fact, the size and complexity of lexical entries can be used as a touchstone for the degree to which the system is lexically based or not. While this approach may work for a time for any specific domain, it lacks the power that comes from a general, well-founded theoretical underpinning. This is the reason such systems tend to become larger and seemingly less defined as they grow, while not necessarily exhibiting greatly increased performance.

The Interlingua Approach: The second approach is to use an interlingua as a language into which to parse and from which to generate. Early attempts at an interlingua (such as the Conceptual Dependency representation [Schank 75]) did not lead to much success primarily due to the difficulty of dealing with terms on a very primitive (in the sense of basic or fundamental) level: sentences, when parsed, had to be decomposed into configurations of the basic elements, and to be generated, had to be reassembled again. Given the basic level of the elements used at the time, this task was too complex to support successful MT.

The Transfer Approach: Many later systems relied less on translation rules hidden in the lexicon and more on representation-transforming rules associated with representational features. This approach gained popularity when early experiments with interlinguas failed due to researchers' inability to develop powerful enough language-neutral representation terms. However, the approach also suffers from a proliferation of rules, especially when more than two languages are present: for n languages, $O(n^2)$ sets of translation rules are required.

At present, no single approach is the clear winner. The systems with the most practical utility at present, the commercially available Japanese systems, all use a relatively crude lexical approach and derive their power from the brute force provided by tens of thousands of rules. Most promising for newer more general systems seems to be a mixture of the interlingua and transfer approaches.

WHY A NEW ATTEMPT AT MAT?

The time is ripe for a new initiative in the investigation of MAT in the U.S.A. The principal reasons for this are both strategic and technical. In the first instance, a large amount of MT work is being done in Europe (including such multinational projects as the EEC-wide EUROTRA) and Japan, with increasing success; little MT work is done in the U.S. (most of which is funded by Japanese money). In the second instance, recent technical breakthroughs, coupled with the steady advances of the past 25 years, make possible the establishment of small MAT projects and their rapid growth to achieve a high level of sophistication.

These advances, discussed in more detail below, are the following:

- Advances in the theory of representation languages
- The maturation of a representation scheme which enables the melding of the best features of the interlingua and transfer approaches
- Steady advances in grammar development
- Steady advances in generation and parsing technology

Representation Languages:

Advances have been made in the theory of representation languages which make possible a new integrated treatment of syntax and semantics. Usually, semantic knowledge is represented in knowledge representation languages such as those of the KL-ONE family. Syntactic knowledge, on the other hand, is hardly ever (if at all) represented in these languages, and neither are the numerous intermediate structures built by parsers. This is because disjunction (the logical operator **OR**) has generally not been included in the language capabilities. The result is a serious problem, since parsers necessarily deal with multiple options due to the structural and semantic ambiguities inherent in language. The inability to represent both syntactic and semantic knowledge in the same system has precluded the development of parsers using a single inferencing technique to perform their work in a homogeneous and unified manner. Thus the lack of a general framework for computing with disjunctive knowledge structures has always been a hindrance to the development of parsing technology.

Work is currently under way to incorporate inference over disjunctions into Loom [MacGregor & Bates 87], a newly developed exemplar of the KL-ONE-like languages, at ISI. This work extends the capabilities of earlier methods for handling disjunctive descriptions in unification-based parsers (see [Kasper 87, Kasper 88]). It is expected to be completed by the end of 1989. This breakthrough will have two major effects: greatly simplified parsers and enhanced processing speed and efficiency.

In more detail, this innovation makes possible, in a single KL-ONE-like representation system, the representation of both semantic and syntactic knowledge. In this scheme, the automatic concept classifier will be used as a powerful resource to perform simultaneous syntactic and semantic-based classificatory inference under control of the parser. Until now, the flow of control between syntactic and semantic processing has always been a vexing question for parsers: for semantic processing, they have used classificatory inference of various kinds, and for syntactic processing, a variety of other methods, including unification. Since syntactic ambiguities are often resolved by semantic information, and vice versa, it is important to make the results of each type of processing available to the other as soon as possible.

Difficulties in doing so have always meant that one or the other process is made to perform more work (in some cases significantly more) than necessary, requiring the maintenance of numerous alternatives of interpretation. Under the new scheme, the representation of syntactic and semantic knowledge in the same representation system simplifies the parsing process considerably, since there is then only one inference process and its results are represented in a single formalism. Also, the speed and efficiency of the parser is increased, since each type of processing can be performed as soon as possible and no additional work need be done. This new integrated approach, enabled by the ability to handle inference over disjunction, has not been developed before.

Melding interlingua and transfer approaches:

A second breakthrough is the maturation of a representation scheme which enables the melding of the best features of the interlingua and transfer approaches. Problems arise with the interlingua approach either when the interlingua is too 'shallow' to capture more than the surface form of the source text (and hence requires nuance-specific translation rules) or when it is too 'deep' to admit easy parsing and generation, as is the case with Conceptual Dependency [Schank 75].

Knowledge representation experience over the past 15 years has resulted in a much better understanding of the different types of representation schemes and of the ways to define representation terms that support the tasks at hand (the literature contains much work in this regard; see for example [Hobbs 85, Hobbs et al. 86]). However, the organization of such terms to facilitate optimal language translation has been a problem until the recent recognition that a taxonomy of abstract generalizations of linguistically motivated classes can be used as a type of generalized transfer rule. It has become clear that, using the abstract conceptual categories necessary to support the generation of the source and target languages (such as the Upper Model for English in the Penman system; see [Bateman et al. 89c]), it is possible to exploit the commonalities across languages to bypass the need for numerous transfer rules. To the extent that English shares with the other languages a linguistically motivated underlying ontology of the world (especially at the more abstract levels, taxonomizing the world into objects, qualities, and processes such as actions, events, and relations), such a conceptual model can act as a type of interlingua in an MAT system, where differences are taken care of by transfer rules of the normal type. For example, the fact that actions have actors is general enough to be part of the generalized 'interlingua', while particularities of tenses in various languages is not.

By building a suitable taxonomic organization of these terms, both the abovementioned problems can be avoided: by defining enough specific terms in the taxonomy, nuances present in the domain can be represented; and by basing the terms of the taxonomy on linguistically derived generalizations (instead of, say, on notions about the underlying reality of the physical universe as in the case of CD), the ease of parsing and generation can be guaranteed. The use of such a taxonomy for MAT has been investigated; a pilot study is reported in [Bateman et al. 89a, Bateman et al. 89b]. The central ideas are described in some considerable detail in [Bateman 89]. This semi-interlingua approach is preferable to the lexically based and pure transfer approaches, since it minimizes the number of special-purpose rules required to customize the system to a new domain, and hence increases the power and portability of the MAT system.

Grammar development:

One of the steady advances in the field of Natural Language Processing is the development of more complete grammars. There exist today computational grammars that cover English (and other languages such as German, Chinese, Japanese, and French) far more extensively than the most comprehensive grammars of 20 years ago did. Modern MAT system developers thus need spend much less effort on grammar development and can concentrate on less understood issues.

Generation and parsing technology:

Another advance is in generation and parsing technology. The issues in single-sentence parsing and

generation have been studied to the point where a number of well-established paradigms and algorithms exist, each with known strengths and weaknesses (in fact, in the last 5 years a number of general-purpose generators have been distributed, including Penman [Penman 88], MUMBLE [Meteer et al. 87], and SEMSYN [Rösner 88]). Obviously, this situation greatly facilitates the construction of new MAT systems.

Knowledge about Machine Translation:

The amount of knowledge about MT available today is much larger than it was 20 years ago. More than one journal is devoted to the topic (for example, *Computers and Translation*). Books on the subject include [Nirenburg 87, Slocum 88, Hutchins 86]. Some larger MT systems developed over the past decade are the EEC-sponsored EUROTRA project [Arnold & des Tombe 87, Arnold 86], the METAL project [Bennett 82], the Japanese-German SEMTEX-SEMSYN project [Rösner 88]. Two current MT projects in the U.S. are KBMT [Nirenburg et al. 89] and a project at the CRL (New Mexico State University).

WHAT WOULD AN MAT PROGRAM INVOLVE?

The three cornerstones of an MT system are the parser, the generator, and the knowledge representation system. Computational Linguistics research has developed far enough today that there are available in the world at least four general-purpose language generators, two of them English-based, and a number of limited-purpose parsers. A number of knowledge representation systems are also available, some of which commercially (such as KEE, manufactured by Intellicorp), and others in the public domain (such as NIKL and Loom [Kaczmarek et al. 86, MacGregor & Bates 87]).

In essence, an MAT project under the new effort would perform six steps:

- Step 1: The *selection of an approach*, and the resultant development (in transfer and lexicon-based systems) of the transfer rules. The success of a translation system can depend greatly on the theoretical approach taken, which can hinder the principal task of identifying the major bottlenecks for MT.
- Step 2: The selection or development of *parsing and generation mechanisms*, together with auxiliary information sources such as lexicons and knowledge representation systems. Given the existing work in Natural Language Processing at various centers in the U.S., and the general availability of parsing and generation technology, multiple options are available to the projects.
- Step 3: The development of the *grammars* of the various languages involved, and their incorporation with the parser and generator. This task can be more difficult, depending on the availability of experts in the languages. However, international collaborations, or the use of grammars built overseas, can greatly facilitate the task.
- Step 4: The selection of an *application domain* and the representation of its elements. Given that the primary goal of the program is to produce prototype machine-aided translation systems that identify MT bottlenecks, this task should be addressed with care.
- Step 5: The actual *parsing and generation of texts* to constitute the translation.
- Step 6: The *evaluation* of the results. This task is very important in order to compare the strengths of the various approaches and to identify the major problems facing MT. A standard set of MT tests should be applied at various stages of the program.

A SCENARIO FOR AN MAT PROGRAM

This section outlines how a solid MAT capability can be achieved in the next five years for a relatively small investment.

WHAT SHOULD THE PROGRAM AIM FOR?

The program should aim at establishing a small MAT program in the U.S. to conduct good research on the basic issues, to stay abreast of the developments happening elsewhere in the world, to develop and exploit current breakthroughs in the technology in the form of prototypes that perform machine-aided translation, and to foster collaborations among the various Darpa-supported NLP projects. Its goals should be:

1. to stimulate development and incorporation of the newest techniques in order to identify and push the limits of MT possible today,
2. to focus on technologies that provide general, extensible capabilities, thereby surpassing less general foreign efforts,
3. to develop prototype systems that exemplify this work in limited domains,
4. to use the tests developed by various other MT projects (such as the EEC project EUROTRA) to measure the progress and success of the current technology, and to identify its most serious bottlenecks and limitations,
5. to stimulate collaborations and software sharing among various groups developing appropriate NLP-related software and theories in this country.

The program should *not* aim at the development of single-sentence translation systems with wide coverage of narrow domains that possess little generality (as proven by the commercially available Japanese systems, this can be achieved by brute force). It should instead aim at the development of prototype systems that illustrate the translation of multipage texts, and that are general, easily portable, and accommodate new domains and languages with a minimum of effort. That is, generality and feasibility are the properties that will propel this effort beyond current technology.

OVERVIEW OF THE PROPOSED PROGRAM

This subsection describes some important facets of the proposed MAT program.

Given the amount of existing NLP technology, a relatively small investment can result in a significant effort in MAT over a period of 5 years. By making use of existing parsers and generators and grammars, individual projects can be kept reasonably small in manpower (on the order of four to five people per project).

Limiting project size enables the support of a greater number of projects. This is important because, due to differences in their theoretical approach, systems can be variously successful simply by virtue of

the limitations of the theory they embody, and can thus hinder the principal task, which is to identify the major bottlenecks of MT. Therefore the program should encourage two or three different theoretical approaches in order to help find the best one and to promote the development of technology which will deliver near-term machine-aided translation and lay the foundation for full machine translation in the long run.

The program should specify a domain of application for the MAT systems which is easily modeled and represented, and for which the language typically used is clear and relatively unambiguous. A popular domain for existing MAT systems is that of technical documents such as computer manuals, descriptions of computer architectures or operating systems, etc. Another alternative domain is intelligence reports. Beyond the obvious advantages of such domains is the fact that evaluation techniques and tests have already been developed for translated technical documents by such projects as EUROTRA.

In order to ensure that the systems developed are reasonably flexible and general, they should be encouraged to be more than bilingual. This can be achieved by developing the systems first to handle English and one other language and then to incorporate a third afterward. This suggests a 5-year plan broken into three phases: a startup phase of one year for English-to-English paraphrasing, a second phase of two years to include a second language, and a final phase of two years to refine the second language and include a third language.

This scenario involves a 5-year plan, at an investment of between \$1 million and \$2.5 million per year, as follows:

- **Year 1:** \$1 million — startup (phase 1) (*paraphrase*)
- **Year 2:** \$1.5 million — construction of phase 2 of system (*bilingual translation*)
- **Year 3:** \$2 million — completion and demonstration of phase 2
- **Year 4:** \$2.2 million — refinement of phase 2, construction of phase 3 (*trilingual translation*)
- **Year 5:** \$2.4 million — further refinement, demonstration, and evaluation of final system

This money should support three groups of between 3 and 5 people per group. At various times, each group would use the services of a parser specialist, a generator specialist, a knowledge representation specialist, and a text specialist, as well as of programmers. Since it is unlikely that any single group will have available such depth of experience, this requirement would foster collaborations among NLP research projects in this country.

PROGRAM TIMETABLE

In order to minimize the amount of wasted effort, projects under this program should be encouraged to use as much existing NLP technology as possible. This is to some degree enforced by the requirement of a demonstration after 3 years, which is quite reasonable given the availability of general-purpose generators, parsing techniques, and knowledge representation systems. An additional saving of effort can be achieved by using, as second and third languages, grammars that have been developed by grammarians and computational linguists in other countries. It is suggested that German be used as the second language, since a number of computational grammars of German exist in the public domain, and since German is structurally very close to English. The third language could be the choice of individual projects so as to allow them to capitalize on their strengths, but should be a language structurally quite different from English (such as Japanese or Chinese), so as to test the generality of the underlying theoretical approach. Thus the program can be structured as follows:

Year 1:

- Selection and adaptation of a parser.
- Selection and adaptation of a generator.
- Selection and incorporation of the English grammar(s).
- Representation of the domain, construction of a domain model.
- Selection and establishment of the English lexicon.
- Demonstration of the first stage of the system by a limited paraphrase task: parsing English texts and then generating English paraphrases of them.

Year 2:

- Selection and initial incorporation of the German grammar.
- Selection and incorporation of the German lexicon.
- Integration of the initial German grammar with the parser and generator.
- Demonstration of the second stage of the system by parsing some German texts and generating English equivalents and vice versa.

Year 3:

- Refinement of the German grammar.
- Refinement of the German lexicon.
- Completion of additional data sources such as domain models, transfer rules, etc.
- Integration of the completed German grammar with the parser and generator.
- Demonstration of the third stage of the system by parsing German texts and generating English equivalents and vice versa.
- Establishment of a prototype English-German MAT system.

Year 4:

- Selection and incorporation of the third language grammar (e.g. Japanese)
- Selection and incorporation of the third language lexicon.
- Refinement of the English-German translations by development of additional techniques and transfer rules.
- Demonstration of the refined English-German MAT system.

Year 5:

- Completion of the third language grammar.
- Completion of the third language lexicon.

- Completion of additional data sources such as domain models, transfer rules, etc.
- Integration of the completed third language grammar with the parser and generator.
- Demonstration of the final stage of the system, comprising translation in all six directions between English, German, and the third language.
- Evaluation of the coverage and sophistication of the translations using the test and measures developed by EUROTRA, as applicable to the domain.
- Reports of the major shortcomings and bottlenecks that stand in the way of more complete MT.

The program should encourage evaluation of the prototype systems at every stage, using a well-conceived set of measures such as those of the EUROTRA project. One measure of evaluation is to count the number of sentences translated correctly (i.e., without requiring other than stylistic changes by the editor). This measure can be subdivided according to the type(s) of error made: syntactic, semantic, lexical, unknown word (lexical), unknown concept (semantic), etc. The projects should aim at a 50% correct sentence rate by the end of phase 2 and for a 75% rate for the German translation by the end of the program. Another measure is to compare the time required to translate a piece of text by a human alone with the time taken by a human in conjunction with the system. Existing commercial systems, using brute-force techniques, claim a speedup rate of 50%, establishing a bottom line which the projects' prototypes can strive to improve.

CONCLUSION

The time is ripe for a new program in machine-aided translation of natural language. New technology from the fields of generation, parsing, and knowledge representation can be brought together into prototype MAT systems that can lead the way for working commercial systems (just as MT technology developed in the early 60's is currently being embodied and sold in Japan).

A number of technical reasons make this an opportune time to start such a program. They are developments of the following kind:

- Better representation techniques, such as the ability to handle disjunction in KL-ONE-like representation languages.
- Clearer understanding of semantics, including the development of very general conceptual taxonomies to capture generalized transfer rules.
- More complete grammars.
- Better existing generation and parsing technology.
- Greatly enhanced MT experience and developed evaluation techniques.

As outlined in this document, a very moderate investment over 5 years can result in the creation of three distinct MAT prototype systems, each supporting translations between English, German, and one other language. This is an opportunity which should be seized before the breakthrough technology currently being developed in the U.S. is copied and taken further elsewhere.

The benefit to Darpa and the Natural Language Computational community is clear. For relatively little expense, a major new MAT effort will come into being in the next few years. Much leverage will be gained from the collaborations among projects in the research community, utilizing the existing generation and parsing capabilities to optimal effect.

Acknowledgments

For ideas and help thanks to John Bateman, Bob Kasper, Ron Ohlander, and Richard Whitney.

References

- [Amano 86] Amano, S. The Toshiba Machine Translation System. In *Japan Computer Quarterly*, Vol. 64, 'Machine Translation — Threat or Tool', pp. 32-35, 1986.
- [Arnold 86] Arnold, D. Eurotra: A European Perspective on MT. In *Proceedings of the IEEE*, Vol. 74, pp. 979-992, 1986.
- [Arnold & des Tombe 87] Arnold, D.J. and des Tombe, L. Basic theory and methodology in EUROTRA. In *Machine Translation: Theoretical and Methodological Issues*, Nirenburg, S. (ed), Cambridge University Press, Cambridge, 1987.
- [Bateman 89] Bateman, J.A. Upper Modeling for Machine Translation: A level of abstraction for preserving meaning. Unpublished Penman Project document, ISI/USC, Marina del Rey, 1989.
- [Bateman et al. 89a] Bateman, J.A., Kasper, R.T., Schütz, J. and Steiner, E. A New View on the Process of Translation. In *Proceedings of the European ACL Conference*, Manchester, 1989.
- [Bateman et al. 89b] Bateman, J.A., Kasper, R.T., Schütz, J. and Steiner, E. Interfacing an English Text Generator with a German MT Analysis. To be published as *Proceedings of the Gesellschaft für linguistische Datenverarbeitung*, Springer, 1989.
- [Bateman et al. 89c] Bateman, J.A., Kasper, R.T., Moore, J.D. and Whitney, R.A. The Penman Upper Model. Unpublished Penman Project document, ISI/USC, Marina del Rey, 1989.
- [Bennett 82] Bennett, W.S. The linguistic component of METAL. Working paper, Linguistic Research Center, University of Texas at Austin, 1982.
- [Carbonell et al. 81] Carbonell, J.G., Cullingford, R.E. and Gershman, A.V. Steps towards Knowledge-Based Machine Translation. In *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 3, pp. 376-392, 1981.
- [Carbonell & Tomita 87] Carbonell, J.G. and Tomita, M. Knowledge-Based Machine Translation, the CMU Approach. In *Machine Translation: Theoretical and Methodological Issues*, Nirenburg, S. (ed), Cambridge University Press, Cambridge, 1987.
- [Hobbs 85] Hobbs, J.R. Ontological Promiscuity. In *Proceedings of the Conference of the Association for Computational Linguistics (ACL)*, Chicago, 1985.
- [Hobbs et al. 86] Hobbs, J.R., Croft, W., Davies, T., Edwards, D. and Laws, K. Commonsense Metaphysics and Lexical Semantics. In *Proceedings of the Conference of the Association for Computational Linguistics (ACL)*, New York, 1986.
- [Hutchins 86] Hutchins, W. (ed). *Machine Translation: Past, Present, Future*. Ellis Horwood Ltd, Chichester, 1986.
- [Kaczmarek et al. 86] Kaczmarek, T.S., Bates, R. and Robins, G. Recent Developments in NIKL. In *Proceedings of the 5th AAAI Conference*, Philadelphia, 1986.
- [Kasper 87] Kasper, R.T. A Unification Method for Disjunctive Feature Descriptions. In *Proceedings of the 25th Annual Conference of the Association for Computational Linguistics*, Stanford, 1987.

- [Kasper 88] Kasper, R.T. Conditional Descriptions in Functional Unification Grammar. In *Proceedings of the 26th Annual Conference of the Association for Computational Linguistics*, Buffalo, 1988.
- [Laubsch et al. 84] Laubsch, J., Rösner, D., Hanakata, K. and Lesniewski, A. Language Generation from Conceptual Structure: Synthesis of German in a Japanese/German MT Project. In *Proceedings of COLING 84*, Stanford, 1984.
- [Lytinen 84] Lytinen, S.L. *The Organization of Knowledge in a Multi-Lingual, Integrated Parser*. Ph.D. dissertation, Yale University Research Report # 340, 1984.
- [MacGregor & Bates 87] MacGregor, R. and Bates, R. The Loom Knowledge Representation Language. In *Proceedings of the Knowledge-Based Systems Workshop*, St. Louis, 1987. Also available as USC/Information Sciences Institute Research Report RS-87-188, 1987.
- [Meter et al. 87] Meter, M., McDonald, D.D., Anderson, S., Foster, D., Gay, L., Huettner, A. and Sibun, P. Mumble-86: Design and implementation. University of Massachusetts Technical Report COINS-87-87, 1987.
- [Nakamura et al. 88] Nakamura, J., Tsujii, J. and Nagao, M. GRADE: A Software Environment for Machine Translation. In *Computers and Translation*, Vol. 3:1, pp. 69-82, 1988.
- [Nirenburg 87] Nirenburg, S. (ed). *Machine Translation: Theoretical and Methodological Issues*. Cambridge University Press, Cambridge, 1987.
- [Nirenburg et al. 89] Nirenburg, S., Tomita, T., Carbonell, J., Nyberg, E. and others. KBMT-89 Project Report. Center for Machine Translation, CMU, Pittsburgh, 1989.
- [Penman 88] *The Penman Primer, User Guide, and Reference Manual*. Unpublished USC/ISI documentation, 1988.
- [Rösner 88] Rösner, D. The generation system of the SEMSYN project: Toward a task-independent generator for German. In Zock, M. and Sabah, G. (eds), *Advances in Natural Language Generation: An Interdisciplinary Perspective*, Frances Pinter, London, 1988.
- [Schank 75] Schank, R.C. *Conceptual Information Processing*. North-Holland Press, Amsterdam, 1975.
- [Slocum 88] Slocum, J. (ed). *Machine Translation Systems*. Cambridge University Press, Cambridge, 1988.
- [Time 89] Hillenbrand, B. Trying to decipher Babel. In *Time Magazine*, July 24, 1989.