

Chapter 1

EXAMPLE-BASED TRANSLATION IN A HYBRID INTEGRATED ENVIRONMENT

Reinhard Schaler

Localisation Research Centre

University of Limerick

Limerick, Ireland

Reinhard.Schaler@ul.ie

Michael Carl

Laboratoire de Recherche Appliquée en Linguistique Informatique

Département d'informatique et recherche opérationnelle

Université de Montréal, Canada

carl@IRO.UMontreal.CA

Andy Way

School of Computer Applications

Dublin City University

Dublin 9, Ireland

away@computing.dcu.ie

Keywords: Translation Memories, Phrasal Lexicon, Controlled Translation

Abstract One key to the success of EBMT is the removal of the boundaries limiting the potential of translation memories (TMs). We discuss a linguistically enhanced TM system, a *Phrasal Lexicon* (PL), which takes advantage of the huge, underused resources available in existing translation aids. We claim that PL and EBMT systems can only provide valuable translation solutions for restricted domains, especially where controlled language restrictions are imposed. When integrated into a hybrid and/or multi-engine MT environment, the PL will yield significant improvements in translation quality. We establish a future model of translation

usage and anticipate that EBMT and the PL will have a central place in future hybrid integrated translation platforms.

1. INTRODUCTION

It is right for the wider MT community to acknowledge that there have been a number of success stories where rule-based MT (RBMT) systems have provided a general solution to the problems of translation. Examples include the *Meteo* system (Chandioux, 1976) and PAHO's *ENGSPAN* and *SPANAM* (Vasconcellos and Leøn, 1985) systems, to name but two.

However, at the same time it is generally agreed that the overall quality of today's MT systems is somewhat short of what might be desired. Nevertheless, translators remain wary that MT software might pose a considerable threat to well established work practices, despite clear evidence to the contrary. Indeed, it is hard to know how such entrenched views may be overcome, despite the best efforts of MT educators and researchers to level the playing field in this emotive issue. It is clear that RBMT *can* be a solution: the problem is in knowing in what circumstances. We will address this issue in more depth later in the paper.

Despite the widespread view that RBMT will never be good enough to warrant serious consideration as systems capable of high quality general purpose translation, research and development in RBMT systems continues to this day. Furthermore, their deployment, especially in limited domains, cf. *PaTrans*, Ørsnes *et al.*, 1996, is growing, and real savings are being made where business professionals are prepared to keep an open mind and take the time to see where RBMT can be of use, and where it (probably) cannot.

At the same time, Translation Memory (TM) systems have rapidly come to be regarded as an extremely useful tool in the translator's armoury. Notwithstanding the widespread acceptance of such tools, the emergence of TM applications have continued to keep some translators on their guard. This is even harder to understand: TM systems do *not* translate. All they do is find close matches for the input string in their database of previously seen translations and display these matches together with their translations for the translator himself to manipulate into the final, output translation. At all stages in the translation process, the translator is the integral figure: he is free to accept or reject any suggested matches, and may or may not insert any suggested translations into the target document and the TM itself wherever he deems this to be appropriate.

In the view of some researchers, TM technology may be considered as some sort of sophisticated search-and-replace engine. For example, Macklovitch and Russell, 2000 liken TM systems to information retrieval tools, in that all the translator does in effect is search for 'documents' (aligned sentence pairs) which may help translate a given sentence. The TM system formulates the

query itself from the input string to be translated, i.e. in most cases, the query *is* the source language sentence, while the retrieval process tries to match the query as closely as possible to existing examples in the system database.

A *sine qua non* for TM (and Example-based Machine Translation (EBMT)) applications is an aligned parallel corpus. The translator may be presented with a pre-aligned TM, but may also be a central figure in its construction: using a tool such as Trados *WinAlign*, for instance, the translator may manually overwrite some of the decisions made by the aligner by linking ⟨source, target⟩ sentence pairs by using the graphical interface provided.

However, the potential of TM technology to facilitate the task of the translator still further is limited by the fact that the fragments contained in the TM are linked only at sentential level. If sub-sentential alignment could be integrated into the TM databases, more useful fragments could be put at the disposal of the translator. If a more radical view were to be taken, once sub-sententially aligned chunks exist, then automatic translation via EBMT can be made available reasonably straightforwardly.

Although EBMT as a paradigm has been described in research papers as far back as the mid 1980's (Nagao, 1984) and although it has managed to capture the interest and enthusiasm of many researchers, it has—so far—not reached the level where it could be transformed from a research topic into a technology used to build a new generation of machine translation engines—and new approaches, technologies and applications are badly needed in MT.

We believe that the time is ripe for the transformation of EBMT into demonstrators and, eventually, commercially viable products. We see the niche area for EBMT as being somewhere between RBMT and TMs. In this paper we shall claim that EBMT—like all corpus-based approaches to MT, be it human-assisted or fully automatic—may successfully be applied only for restricted domain, special purpose translations. We investigate the relation of controlled language and translation more precisely and come to the conclusion that EBMT is a paradigm especially suited to controlled translation. We develop a model for the future of translation technology, and propose an integrated hybrid environment for translation mining, management and scheduling. We believe that example-based translation technologies will play a central role in such a hybrid translation environment, as they are by their very nature dynamic, have the ability to learn, are extensible and scalable to new, large sublanguage domains.

In section 2 we argue for using phrases (and not just words) as basic translation units in any MT system and contrast our approach to mainstream MT research. We examine TMs as a possible application for a phrasal lexicon in section 3 and point to their structural limitations. As a consequence we require in section 4 that TMs are enhanced with more sophisticated processing devices. Here we investigate the impact of controlled language and MT. Our

conclusion is that EBMT is especially suited to obtaining translations in restricted domains, as they are easy to adapt and able to produce high quality translations. We then investigate multi-engine and hybrid system architectures in section 5, and develop a model for the future use of translation technology in section 6. We anticipate a hybrid integrated translation platform which, as a kernel, contains an example-based component capable of adapting quickly to different user needs and which is complemented by linguistic, statistical and rule-based technologies.

2. EBMT AND THE PHRASAL LEXICON

EBMT has been proposed as an alternative methodology capable of replacing RBMT, initially by Nagao, 1984, followed by extensions reported in Sato and Nagao, 1990 and Sadler and Vendelmans, 1990. EBMT has also been proposed as a solution to specific translation problems, as reported in Sumita and Iida, 1991, and culminating in the current volume.

There is an enormous variety of approaches to the use of examples in natural language processing (NLP), with different areas of focus and with different motivations. In the area of automatic translation, this provides a strong testimony to the high level of interest in EBMT. Taking existing parallel texts as their starting point, some of the areas that researchers have worked on include:

- Word-sense disambiguation;
- Translation ambiguity resolution;
- Lexicography;
- Extraction of bilingual collocations or translation patterns from parallel corpora;
- Translation Quality Measures;
- Extensions to and variations on the basic idea of EBMT.

One idea which precedes all of the approaches mentioned and which, surprisingly, has until recently not been taken up by researchers to any significant degree (cf. Simard and Langlais, 2001), is that of the ‘Phrasal Lexicon’, described first by Becker, 1975, and applied by Schaler, 1996.

Becker’s model was radically different from the mainstream linguistic theories prevalent in the mid-1970’s: instead of considering language production as the process of combining units the size of words or morphemes to form utterances, he identified phrases consisting of more than one word as the building blocks for the formation of utterances.

For Becker, language generation is compositional¹ in the way illustrated in Figure 1.1. The phrasal lexicon provides patterns that can provide (at least

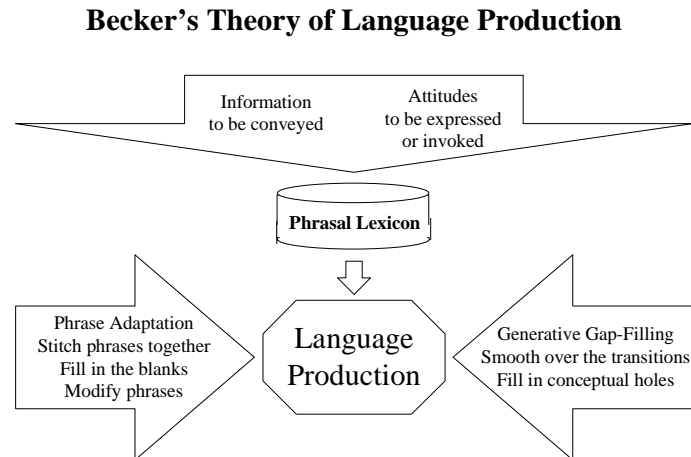


Figure 1.1. Becker's Theory of Language Production

some of) the expressions needed to convey a message in a certain 'tone'. These phrases are then stitched together, blanks filled in and phrases modified where necessary. If this is not sufficient to generate the utterance, new phrases are generated to smooth over transitions and fill in conceptual holes.

It is not difficult to see that the notion of language generation as envisaged by Becker is in exactly the same spirit as many approaches to EBMT. With respect to TM technology, translations cannot be produced automatically (except in the case of exact matches) as sub-sentential ⟨source, target⟩ chunks are not identified by the system: the translator himself has to use his knowledge of how the source language relates to the target, extract the relevant parts of any close matches suggested by the TM system and stitch them together himself to produce the final translation. If there are any 'gaps' or 'conceptual holes', the translator uses his expertise to fill these in and produce the output text in the target language.

2.1 TM Resources are underused

We believe that the time is ripe for the transformation of EBMT into demonstrators, technologies and eventually commercially viable MT engines along the lines suggested by Schaler, 1996 and Macklovitch, 2000, both of which are based on the belief that existing translations contain more solutions to more translation problems than any other available resource (Isabelle *et al.*, 1993).

The key to the success of this development, we suggest, is the removal of the boundaries limiting the potential of translation memories. To bring EBMT to fruition, researchers and developers have to go beyond the self-imposed limitations of what is now traditional—in computing terms almost old fashioned—TM technology.

In Translation Memory systems, two intellectually challenging problems have to be addressed which cannot just be solved by clever engineering:

1 **Contents of the TM (alignment).**

The first problem occurs when translators want to create *translation memories* by aligning previous translations with their source equivalents on a segment-by-segment basis in order to import these aligned segments into a TM and then use this for the translation of a new version of the same source material. Developers generally offer alignment tools which work in either interactive or fully automatic mode.

2 **Retrieving similar entries.**

The other is the decision of how to deal with cases where no exact matches can be found. Developers generally opt to search for similar matches and to calculate a ranking of identified ‘fuzzy’ matches which are then offered to the translator as a possible basis for the translation of the new segment.

Current TM systems do not have to carry out any significant amount of linguistic processing. They do not need to know practically anything about the target language as the processing (matching, calculation of fuzziness, identification of changes etc.) is performed on the source language only.

While the latter problem still remains to be solved both in theory and in practice (the work of Planas and Furuse, 1999 is a significant step in this direction, cf. Chapter 5, this volume), the former has often been claimed to be solved to a large extent. However, we will show that there is still considerable work to be done if TM technology is to go beyond its current limitations.

2.1.1 Contents of the translation memory. The availability of alignment tools linked with the now widespread use of translation memory systems has led to the creation of massive bilingual and multilingual parallel corpora aligned at sentence level, the smallest segment level currently accessible by TM systems. However, matching segments at sentence level unnecessarily restricts the potential and the usefulness of translation memories as extremely valuable linguistic resources. Thus, translation knowledge at a phrasal or sub-sentential level is lost if alignment and matching only works on a sentence level.²

A second shortcoming in current TM technology is the unavailability of tools to control and/or cross-check the extracted segments. Furthermore, there

is no easy way to control the quality of the segments and their translation. For instance, it is possible (in principle at any rate) to map any source sentence onto the same target expression. Given a real reference text, this is, of course, an extremely unlikely case. Nevertheless, any TM system is likely to store redundant, ambiguous or wrong translations in the TM with no possibility of reducing or avoiding such entries other than manually checking the entire database.

Inconsistent translation segments are likely to be produced during initial alignment. They can also be produced when multiple users share a TM in real-time, adding different translations for the same sentence. This is usually controlled using meta-data fields such as <user>, <date>, etc., but unless the process for sharing TMs is properly controlled, this can still add to the inconsistencies present in the TM. While in some cases such inconsistencies might be required, in other cases it is undesirable.

Furthermore, last minute changes are frequently made to translated documents *outside* the TM environment. If the process of using TMs is controlled properly, these changes will be added to the TM. However, this does not happen on a frequent basis and leads to gaps between the reference text and the new translations, which reduces the precision of the TM and increases the amount of fuzzy matching required when the TM is used for a new translation.

2.1.2 Retrieving similar entries. If a TM system cannot find an exact match in a TM, it can only propose fuzzy matches. Ideally, a TM system should find a segment (or a set of segments) in its database which is similar in meaning to the new sentence in the hope that their translations will be similar as well. However, given that there is no known algorithm to compute the meaning of arbitrary sentences and to quantify their similarity, this approach seems not to be feasible (cf. Reinke, 1999 for an in-depth discussion on similarity of meaning).

Fuzzy matching can be a highly complex operation: in fact, it may prove so cumbersome that translators often opt out of the fuzzy match proposal operation by setting the percentage threshold of the fuzzy match component so high that high percentage matches which could contain matching phrases are hidden away from them. Instead, they prefer to translate the new source document without the support of the TM system in order to save time. In the language of statistical NLP, they impose a high Precision threshold, but low Recall, for this matching process. The consequence, of course, is that potentially valuable matches at sub-sentence level are lost. Adding a facility capable of providing (source, target) chunks at the phrasal level will overcome this shortcoming.

2.1.3 Recombination of Target Segments. TMs are usually fed with domain-specific reference material in order to achieve maximum precision for translating specialized texts and their segments. However, TMs have little or no capacity to check that such reference texts are consistent. Thus, if the reference text contains two (or more) different translations for the same source language segment, the user is asked which of the translations, if any, he prefers in the given context, without the TM learning from this decision. If the same sentence were to be translated later, exactly the same scenario would arise. Adding more reference translations to the TM increases the chance of adding more ambiguous translations.

The way translations are produced in TM systems is, therefore, *descriptive*, since retrieved and proposed translations reflect the quality of the reference translations and their alignment functions. As the quality of translations available in the TM decreases and the new text or sentence differs from reference examples in the TM, the quality of proposed segment translations decreases in a corresponding manner.

3. THE PHRASAL LEXICON AND EBMT

While we discussed in section 2.1.1 some other possible ways in which the translation process in a TM system may be distorted, there remain two main factors which may adversely affect the way in which the TM is used:

- the way entries are retrieved from the TM (i.e. fuzzy matching);
- inconsistencies in the contents of the TM.

Both factors may cause translation quality to deteriorate, especially where these two problematic factors co-occur. The Phrasal Lexicon avoids these shortcomings by permitting exact phrasal matches together with a manner of controlling the contents of the memory.

3.1 Phrasal Matching

In the phrasal lexicon, exact matches are performed at a phrasal level, so the problem of fuzzy matching disappears (or at any rate, is considerably reduced). The probability of finding exact matches at a lower phrasal level (e.g. at NP, VP or PP level) is significantly higher than the probability of finding exact matches at the sentence level, the level at TM segments are currently provided.

Phrasal units are looked up in a phrasal lexicon and translated by combining already translated phrases stored in the phrasal lexicon, very much along the lines proposed originally by Becker. As an example, let us assume a TM containing the two entries in (1):

- (1) a. EN: The bullets move to the new paragraph.
DE: Die Blickfangpunkte rücken in den neuen Abschnitt.
- b. EN: The title moves to the center of the slide.
DE: Der Titel rückt in die Mitte des Dias.

Despite the fact that the TM entries in (1) are good matches for the new sentence in (2), it would not be translated automatically by the TM system:

- (2) The bullets move to the center of the slide.

At most, the system would be capable of identifying one (or both) of the two source sentences in the TM in (1) as fuzzy matches. If so, these would be displayed together with their translations, which would then have to be adapted by a translator to generate the final output string in (3):³

- (3) Die Blickfangpunkte rücken in die Mitte des Dias.

Note, however, that depending on how fuzzy matching is calculated in the TM, and at what percentage level fuzzy matching is set by the user, the useful translation pairs in (1) may not be presented at all to the translator: if we compare the source strings in (1) with the new input in (2), observe that 5 out of 7 (71%) of the words in (1a) match those in (2) exactly, while 6 out of 9 (67%) of the words in (1b) match those in (2) exactly. If we were to set our level of fuzzy matching at 75% (say), neither candidate would be accessible by the translator.⁴

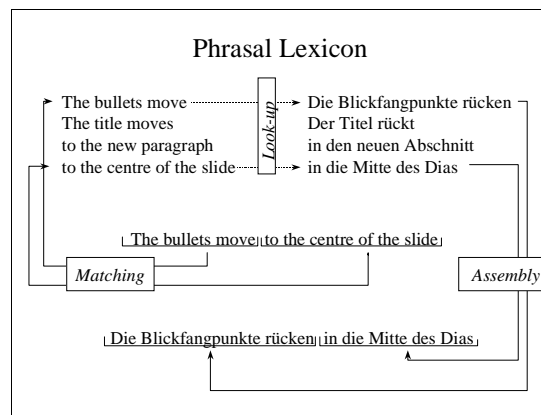


Figure 1.2. Overview of the Phrasal Lexicon

If, however, we produced a phrasal lexicon capable of providing translation units at the phrase level, and wrote a procedure to combine them so that a

correct translation of the new sentence were produced automatically, then we would have built an EBMT system. Notwithstanding this, if a phrasal lexicon (cf. Figure 1.2) were to be integrated into a TM system, then translators would be able to avail of many of the advantages of a TM, principally those of consistency and savings in terms of time and cost.

Storing, matching and proposing segments at the phrasal level has a number of advantages, including:

- Translators will be offered a higher percentage of exact matches from TMs.
- The quality of proposed translation segments will improve.
- The use of information stored in TMs will increase; matching phrases in otherwise fuzzy matching sentences will no longer fall below the match percentage threshold set by most translators.
- TM systems will be able to translate larger amounts of source text automatically without the need to manually adapt fuzzy matches.

3.2 More sophisticated processing devices

Since the PL is based on smaller translation units, it can potentially identify more exact matches than sentence-based TM systems—but at a price. In contrast to a TM, a PL would need to know—at the very least—enough about the source *and* the target language to identify phrases and describe the linguistic characteristics of their constituent parts. For instance, while the translator may be able to pick out the appropriate sub-sentential alignments from (1), namely:

- (4) a. the bullets move \implies die Blickfangpunkte rücken
 b. to the center of the slide \implies in die Mitte des Dias

in order to translate the new input string in (2), it is a non-trivial problem as to how to obtain such translationally relevant chunks automatically. Indeed, this is the topic of this book: automatic or semi-automatic phrasal alignment requires sophisticated linguistic and/or numerical processing as described in Chapters 9–14 in this volume. Typical problems when combining sub-sentential chunks in new translations involve the concept of *boundary friction* (cf. Chapters 4 and 16, this volume). For instance, the sub-sentential alignments in (4) have been chosen so as to be optimal given the new input in (2). Consider instead a second set of alignments which could have been selected from the translation pairs in (1), namely those in (5):

- (5) a. the bullets move to the \implies die Blickfangpunkte rücken in den
 b. center of the slide \implies Mitte des Dias

Combining these translation chunks would result in the mistranslation of (2), namely *Die Blickfangpunkte rücken in den Mitte des Dias*. Here we note in the target PP a masculine accusative determiner *den* together with a feminine singular noun *Mitte*. In the correct translation (3), meanwhile, we see that determiner-noun agreement has been maintained within the NP.

Dependent on the language pair, further restrictions such as the addition of case or semantic markers, or other linguistic information, have to be imposed on the entries in the PL and treated accordingly when recombining target phrases (cf. Carl *et al.*, 1999). For example, Way, 2001 (cf. Chapter 16, this volume) incorporates the syntactic functions of LFG (Kaplan and Bresnan, 1982) into an EBMT system based on aligned (source, target) tree pairs. Way notes that the ill-formed string *John swim* cannot be translated by LFG-MT (Kaplan *et al.*, 1989), as the constraints that the subject NP must be singular and the verb plural cannot be resolved by the LFG constraint solver. In his LFG-DOT system, Way *op cit.* shows that this string may be translated by relaxing these constraints, and, importantly, that the system *knows* that the input string and the translation formed are ungrammatical, unlike other EBMT systems. Perhaps more significantly, this research shows that one needs less than the full target language grammar in order to obtain translations: in RBMT, it is far from clear (a) that a translation could be obtained at all in this situation, and (b) if it could, arguably the whole target grammar would be involved in enabling some translation. In contrast, in a ‘linguistics-rich’ EBMT system such as LFG-DOT, all that is required is to relax one feature on the NP and on the VP for a translation to ensue. This demonstrates a clear advantage of EBMT systems over their rule-based counterparts.

More questions arise with respect to the quality of the translations in the PL. According to Somers and Collins (cf. Chapter 2, this volume), the EBMT literature has little to offer regarding the question as to how best to store cases in relation to each other, as well as how to decide on the organisation of the memory. It is very likely that for the PL in Figure 1.2, a phrase like *The bullets move* would be translated differently if the alignments were taken from a text on wars. While the co-occurrence of *bullet* and *paragraph* in the original alignment implies a high probability that the segment was taken from a software localization context, this information is lost in the PL. However, the translation *bullet* \implies *Blickfangpunkt* cannot be understood in a text tackling wars.

Therefore, a number of methods have been proposed for human-assisted or fully automatic alignment in order to store non-ambiguous translation pairs in the PL. To achieve an accuracy of more than 90% for automatically generated

phrase translations, Watanabe (cf. Chapter 14, this volume) uses a graphical interface where users confirm alignment proposals of the system. Menezes & Richardson (cf. Chapter 15, this volume) use a fully automatic device to filter out translation ambiguities in a post-processing step and Carl (cf. Chapter 12, this volume) investigates the induction of invertible translation grammars from reference translations.

The motivation of the investigation in this chapter is to extract and store high quality translation units according to a number of criteria in order to enable the database to be scaled up in a controlled manner. Consequently, it is unrealistic to think that a corpus-based MT system can succeed as a general solution to the problems of translation. In contemplating such issues, it is clear that the notions of controlled translation and controlled language will play a role. In the next section, we investigate these matters further, and recommend that EBMT systems incorporating a phrasal lexicon might be optimally suited to controlled translation.

4. CONTROLLED LANGUAGE AND MT

Controlled languages define a writing standard for domain-specific documents. Linguistic expressions in texts are restricted to a subset of natural languages. They are characterized by simplified grammars and style rules, a simplified and controlled vocabulary with well defined meanings, and a thesaurus of frequently occurring terms. Controlled languages are used to enhance the clarity, usability, transferability, retrievability, extractability, and translatability of documents.

According to Lehrndorfer and Schachtl, 1998:8, “the concept of controlled language is a mental offspring of machine translation”. That is, one of the main *raison d’être* of controlled language was to facilitate automatic translation of documents. A number of companies (e.g. *Boeing*, *British Airways*, *Air Canada*, and *Caterpillar*) use controlled language in their writing environment. Nor is this trend restricted to English: *Siemens* use controlled German (*Dokumentationsdeutsch*: Schachtl, 1996; Lehrndorfer and Schachtl, 1998, *Aérospatiale* use controlled French (*GIFAS Rationalised French*: Lux and Dauphin, 1996; Barthe, 1998), while *Scania* use controlled Swedish (*ScaniaSwedish*: Almqvist and Sagvall Hein, 1996), for example.

4.1 Controlled Language and RBMT

By definition, therefore, controlled languages are characterized by simplified grammars and vocabulary, a consistent syntax and terminology to which the MT system has to be adjusted. Controlled languages have been developed for restricted domains, such as technical documentation for repair, main-

tenance and service documents in large companies (e.g. *Boeing*, *Siemens*, *Scania*, *GM* etc.).⁵

Caterpillar's 'Caterpillar Technical English', for instance, defines monolingual constraints on the lexicon, constraints on the complexity of sentences, and the use of generalized markup language. However, when using this controlled language for translation in the *KANT* RBMT system (Mitamura and Nyberg, 1995), it was found that "[terms] that don't appear to be ambiguous during superficial review turned out to have several context-specific translations in different target languages" (Kamprath *et al.*, 1998).

Van der Eijk *et al.*, 1996:64 state that "an approach based on fine-tuning a general system for unrestricted texts to derive specific applications would be unnecessarily complex and expensive to develop". Later work in *METAL* applications refers to there being "limits to fine-tuning big grammars to handle semi-grammatical or otherwise badly written sentences. The degree of complexity added to an already complex NLP grammar tends to lead to a deterioration of overall translation quality and (where relevant) speed" (Adriens and Schreurs, 1992:595). Despite this, from 1994, *General Motors* used *METAL* for controlled authoring and MT (Means and Godden, 1996). Furthermore, attempts at redesigning the *Meteo* system (Chandioux, 1976), probably the biggest success story in the history of MT, to make it suitable for another domain (aviation) proved unsuccessful.

Controlled translation, therefore, involves more than just the translation of a controlled language. Passing a source language text through a controlled language tool is not sufficient for achieving high quality translation. Large general purpose (rule-based) MT systems can be converted only with considerable difficulties to produce controlled translations. Controlling the translation process in a conventional transfer-based MT system involves controlling three processing steps:

- the segmentation and parsing of the source text (i.e. analysis);
- the transfer of the source segments into the target language (lexical and structural mapping);
- the recombination and ordering of the target language segments according to the target language grammar (generation).

As the resources of each of these steps require independent knowledge resources, adjusting a conventional RBMT system to a new controlled language is non-trivial.

4.2 Controlled Language and Corpus-based MT

Using traditional rule-based systems for the translation of controlled languages leads to the well known ‘knowledge acquisition bottleneck’. This may take several forms, including:

- Lack of *coverage*: knowledge is restricted to (theoretically interesting) interactions of linguistic phenomena;
- Lack of *robustness*: ‘toy’ grammars and lexica do not scale up;
- Translation data is often invented, so that *consistency* is hard to maintain.

In addition, several knowledge resources have to be adjusted and homogenized. An oft heard claim is that corpus-based MT systems can overcome this bottleneck given that available translations can be exploited. In contrast to traditional approaches, corpus-based MT systems induce the knowledge required for transfer from a reference text. To date, corpus-based MT technologies have yet to tackle controlled languages: they have not supported the acquisition of controlled translation knowledge, nor have they provided an appropriate environment for controlled translation.

This is extremely surprising: the quality of corpus-based translation systems depends on the quality of the reference translations from which the translation knowledge is learned. The more a reference text is consistent, the better the expected quality of the translations produced by the system. By contrast, therefore, translation knowledge extracted from noisy corpora has an obvious adverse impact on the overall translation quality. The only research we are aware of which attempts to detect omissions (but not *other* errors) in translations is Chen, 1996 and Melamed, 2001. However, in the context of corpus-based MT, such methods have not been used to date to eliminate noisy or mistranslated parts of the reference text, nor to enhance the quality and consistency of the extracted translation units.

If EBMT is to make the leap from research topic to industrial applications, we contend that the niche for EBMT may well be found somewhere between RBMT and TM systems. That is, there is a need for a controlled, domain-specific, easily adaptable MT engine, one which proposes translation candidates *and* recombines them in a controlled manner.

Conventional TM systems are not suitable for this task. Within the *Tetris-IAI* project (TETRIS, 1999), an experiment was undertaken whereby controlled language was fed into a TM. It was found that controlling the source language without controlling the reference material does not increase the hit-rate of the TM and thus does not increase the likelihood of high quality translations— from a company’s point of view, the bottom line is that the translation cost is not lowered. Methods for preparing and modifying reference texts to achieve

better consistency on both terminological and syntactic levels have, therefore, been proposed (e.g. Somers, 1993) and could also be a feasible way forward for TMs.

Furthermore, statistics-based MT (e.g. Brown *et al.*, 1990) is not an appropriate candidate for controlled translation. One cannot usually expect consistent reference translations in statistics-based MT, owing to the huge size of the reference texts required. In many cases, texts from different domains are necessarily merged together to compute word translation probabilities for a language pair in various contexts owing to sparse data: there is not enough textual material in one sublanguage domain to compute accurate, robust language and translation models. Of course, how words and phrases are used in different domains can differ greatly. In order to achieve high translation accuracy, therefore, it is unwise to calculate one's statistical language and translation models on corpora merged from different sublanguage areas, but this is often a prerequisite in statistical MT. It is clear from the methodology adopted so far that the implicit goal in statistical MT has been general purpose translation.

In our view, the main potential of EBMT lies in the possibility to easily generate special purpose MT systems. The more restricted the purpose, the better defined such restrictions and the greater the availability of high quality reference translations, so the potential of EBMT to produce high quality translations increases. Brown (Chapter 9, this volume) shows that coverage can be increased by a factor of 10 or thereabouts if templates are used, but it would be fanciful to think that this would scale up to domain-independent translation. Even if EBMT systems were augmented with large amounts of syntactic information (cf. Chapters 15 and 16, this volume), they would in all probability stop short of becoming solutions to the problems of translating general language. Even though a number of techniques are described to select and control segmentation, transfer and recombination, EBMT has not been seen as suited for controlled translation until now. It is our contention that EBMT systems may be able to generate controlled, domain-specific translations given a certain amount of built-in linguistic knowledge together with some preparation of the aligned corpus. This translation knowledge may, we contend, be reused to produce controlled, example-based translations of high quality.

Controlling the translation process in EBMT implies the careful selection of a set of translation examples which are 'similar' to the input. That is, segmentation, transfer and recombination of target segments is largely determined by the set of retrieved examples. Given that only retrieved examples are recombined, the task of controlling EBMT systems is reduced to controlling the retrieval of appropriate analogous examples from the database. This process, in turn, is triggered by the segmentation of the input sentence, which is why this procedure is perhaps the most crucial in any EBMT system, and why, therefore, system descriptions examine this point most carefully and thoroughly.

We will now take a broader view of machine translation to see in which context EBMT might be a suitable component. We examine the role of EBMT in variations of a multi-engine approach, as well as as an integral component of a hybrid architecture. We compare the two basic approaches, and give our view as to how an EBMT system in such an environment might best be located and used.

5. EBMT IN A MULTI-ENGINE ENVIRONMENT

It is widely accepted that different MT paradigms have different advantages and shortcomings. TMs are fed with domain-specific reference translations and are widely used as tools for translators in the area of Computer Assisted Translation (CAT). TMs, however, do not provide sufficient control mechanisms to enable a more sophisticated translation environment. In contrast, rule-based MT systems may automatically provide high quality translations in limited domains (cf. *Meteo*, Chandioux, 1976), but are more usually designed for general purpose translations. As a consequence, RBMT systems are difficult to adjust to specialized types of texts and consequently suffer from limited portability. EBMT systems have more sophisticated processing devices at their disposal than TMs. In addition, they make better use of available resources than TMs do. This, we have argued, gives EBMT systems the unique possibility to generate controlled translations.

Given the different advantages and shortcomings of each approach, multi-engine MT systems have been designed as an attempt to integrate the advantages of different systems without accumulating their shortcomings. We shall now describe and compare a number of different architectures which have been proposed whereby a number of different MT systems and techniques may be integrated.

5.1 Parallel vs. Sequential Linkage

In order to describe these systems and attempt some classification of them, one distinction can be made along the lines of whether entire translation engines are triggered in parallel or sequentially. In a parallel multi-engine scenario, each system is fed with the source text and generates an independent translation. The candidate translations output by each system are then collected and (manually or automatically) recombined to give the ‘best’ translation.

There are a number of projects which incorporate different MT components in parallel in a multi-engine system. The most notable of these are *Verbmobil* (Wahlster, 1993; Wahlster, 2000) and *PanGloss* (Frederking and Nirenburg, 1994).

Verbmobil integrates the complementary strengths of various MT approaches in one framework, i.e. deep analysis, shallow dialogue act-based approach and

simple TM technology. Nübel, 1997 shows that the performance of the integrated system outperforms each individual system. *PanGloss* uses EBMT (Brown 1996, cf. also Chapter 9, this volume) in conjunction with KBMT—the mainline *PanGloss* engine—and a transfer-based engine.

While there is an element of redundancy in such approaches given that more than one engine may produce the correct translation (cf. Way, 2001:23) one might also treat the various output translations as comparative evidence in favour of the best, overall translation. Somers, 1999 observes that “what is most interesting is the extent to which the different approaches mutually confirm each other’s proposed translations”.

In a sequential multi-engine approach, two or more MT components are triggered on different sections of the same source text. The output of the different systems is then concatenated without the need for further processing. This dynamic interaction is monitored by one system—usually the most reliable amongst the available systems. The reasoning behind this approach is that if one knows the properties of the involved translation components, reliable translations can be produced by using fewer resources than in a parallel multi-engine approach.

Integration of a TM with a rule-based component is a common strategy in commercial translation. A dynamic sequential interaction between a TM (*Trados*) and an MT system (*Logos*) is described in Heyn, 1996. In the case where only poorly matching reference translations are available in the *Trados* TM, the input sentence is passed to the *Logos* MT system for regular translation. The user is then notified which of the systems has processed the translation, since the *Logos* system is less likely to produce reliable results.⁶

A similar scenario is described in Carl and Hansen, 1999, where a TM is linked with an EBMT system. This research shows that the quality of translations is likely to be higher for EBMT translation than for TM translation, in cases where the fuzzy match score of the TM falls below 80%.

The OTELO⁷ project is another example of how both local and remote translation resources such as MT, TM, and other NLP applications may be integrated. In OTELO, two MT systems, *Logos* and *IBM LMT*, were combined with the *IBM TranslationManager* and OTELO’s HyperLink Translator.

5.2 Hybrid MT Systems

In each example multi-engine MT system discussed in the previous section, each module in the integrated environment has its own resources and data structures. In a strong integration of two or more MT techniques, however, the same data structures are shared by the different components. Some of the components of such hybrid systems may, therefore, modify or adjust

certain processing resources of another component in order to bring about an improvement in coverage or translation quality.

Coupling statistical data and RBMT often leads to a strong integration to improve translation quality. In some hybrid systems, statistical data is added to the lexical resources of the RBMT system, which adjudge different translation candidates as more or less felicitous for a given textual or thematic context. In particular, it has been shown that statistically enriched RBMT systems can handle collocational phenomena (for example, establishing the most likely translation match in the context of a particular word). The statistical data are drawn from monolingual corpora (Nomiyama, 1991) or independent corpora of source and target language (Doi and Maraki, 1992).

Rayner and Bouillon, 1995 describe an application of statistical data during the rule-based transfer phase. Statistical data are derived by manually scoring translation variants produced by the system. Since the training is based on texts belonging to one specific subject field, typical mistakes made by the system can be corrected. The probability of a transfer candidate is calculated via the transfer probability and the probability of the resulting target structure. As such a multiplication of probabilities requires large amounts of data in order to be effective, such an approach is applicable only to very restricted subject fields where only a few examples may suffice in order to produce reliable data. In such cases, translation quality is traded for improved coverage.

In a hybrid stratificational integration of example-based and rule-based techniques some processing steps are carried out by the rule-based component while for others examples are used.

Menezes & Richardson (Chapter 15, this volume) combine rule-based analysis and generation components with example-based transfer. Bond *et al.* (Chapter 7, this volume) generate translation templates for new sentences on the fly from a set of alignments. The differing sections in the source template and the input sentence are identified and translated by a rule-based noun-phrase translation system. Carl *et al.*, 1999 use EBMT as a front and backend to a RBMT system for translation of simple phrases and multi-word units.

All these approaches attempt to make the most of the strengths of the hybrid approach while compensating for the weaknesses of each as much as is possible. The main idea is to integrate a significant share of human translation experience (as accumulated in Translation Memories and Term Banks) with the rule-based paradigm.

Even a very large TM or EBMT system is unlikely to be able to translate a completely new sentence correctly, let alone an entire new text. However, such systems are able to 'learn' in that new examples can be added to the system database, so that subsequent encounters with previously unknown source strings will be translated successfully. In RBMT systems there is no such analogous process. That is, they do not store translation results for later reuse, so

that all post-editing effort is wasted: RBMT systems will deal with the same input in exactly the same way in perpetuity.

A hybrid model




Figure 1.3. The Phrasal Lexicon as part of a Hybrid MT Environment

A hybrid system (such as that shown in Figure 1.3), in contrast, will be able to learn and adapt itself easily to new types of text. In addition, the rule-based component provides sophisticated language models to a hybrid set up. Consequently, one can envisage that even if none of the individual engines can translate a given sentence correctly, the overall system may be able to do so if the engines are allowed to interact. Even if the individual components improve, the integrated system should always outperform the individual systems with respect to either the quality of the translation, the performance, or the tunability of the system.

6. A MODEL FOR THE USE OF MT

So far in this chapter, we have discussed the Phrasal Lexicon as a link between TM and MT systems. We have focussed specifically on the role of the PL, TM and EBMT in multi-engine and hybrid MT environments. We recommended that the best role for EBMT is in providing high quality translations in restricted domains written with controlled language rules very much in mind.

We will now bring this chapter to a close by presenting a future model for the use of MT and CAT which has been developed in a recent study following discussions with localisation and translation technology experts (Schaler, 1999, Carl *et al.*, 2002). This, we propose, will provide an effective basis for an MT environment in a number of different future situations. The model is based on the assumption that information can be categorised into three types.

The intended translation model visualises translation in the form of a pyramid, as in Figure 1.4.




Figure 1.4. A Model for the Future Use of Translation Technology

At the bottom of this pyramid, we see ‘non-mission-critical’ information, known more widely as the ‘gisting’ market. An example of this type of information might be an on-line article about Napoleon written in French and published on a web site in France, of which a Spanish speaker with no knowledge of French but interested in Napoleon wants a rough and ready translation at minimal cost. We view this as the ideal application scenario to facilitate the widest usage of MT. Indeed, this is currently the biggest growth area of MT: people are translating web pages or communicating with one another in their own languages via email, using on-line MT systems as the translation engine.

In the middle of the pyramid, we see large amounts of material that must be translated accurately; here gisting is simply unacceptable. Examples of this type of information are product manuals or other technical documentation. Most examples of this type of translation are domain-specific. While MT is currently being used at this level, it is more the exception than the rule. However, we have argued that it is in this area of translation that EBMT has the capability to play a central role.

At the top of the pyramid, we see small amounts of mission-critical or creative material to be read where accuracy and presentation are of paramount importance. Examples of this are marketing material, laws, user interfaces and creative literature. These are all areas which are not especially suited to MT, unless expert post-editing is available.

The model presumes (i) that the pyramid is expanding in two directions and (ii) that improvements in translation technology (MT in its widest sense) will open up new markets for developers of MT systems. We began this Chapter by discussing how unnerved some translators continue to be by the availability of MT systems. We noted in the general case that the MT systems of today are simply not capable of generating translations of sufficiently high quality to justify this wariness of MT technology. However, one area where MT has met a demand for translation that up to now was not being met is in the continuous updating of webpages and their translations. Thus MT has provided a solution to this new area of demand, and is not replacing translators at all, as human translation in this area and on this scale simply cannot be envisaged.

The expansion of the pyramid will be driven by two factors:

- a growing demand for translated material given the further trend towards globalisation of the economy (horizontal expansion);

- the increasing availability and accessibility of information in a variety of languages to end-users on the web (vertical expansion).

At the same time, research and development in the area of translation technology will allow MT to push its way up the pyramid and be used for higher quality translation. Translation service providers will offer a variety of on-line translation facilities, from high quality human translation to low-end, cheaper MT, with a range of mixed options in between. These options will include human-edited MT using specialised, fine-tuned lexical and semantic databases, TM-based translation combined with MT, and alignment and maintenance of previously translated legacy material.

We anticipate a hybrid MT platform which integrates together in one environment a wide range of applications, techniques and resources, including (but not limited to):

- multilingual alignment,
- terminology mining,
- automatic and computer-assisted terminology structuring, management and control,
- automatic or semi-automatic induction of grammars and translation templates,
- automatic consistency checks etc.

We foresee that such a hybrid MT platform will also integrate together a number of different techniques and resources such as example-based, statistics-based and rule-based approaches to translation, as well as a variety of linguistic resources and corpora.

Some researchers have given some thought to the suitability of texts for MT. We are all aware of MT systems being confronted with texts that pose unfair demands on them as some texts should *never* be translated automatically, but as Kenny and Way, 2001:13 observe:

“Those of us who have developed MT systems in the past and demonstrated them in various fora can only hope that the days of someone typing in a 50-word sentence consisting of strings of auxiliaries, prepositional phrases and containing ellipses, and the system either keeling over or else coming up with a hopeless ‘translation’ after some minutes, followed by our system tester uttering ‘MT is not for me!’, are long gone”.

Despite the importance of this topic, the work that we are aware of regarding translatability and MT focuses only on what texts should be sent to *rule*-based MT systems. One possible translatability indicator for the use of MT in general is the identification of (sets of) phenomena which are likely to cause problems for MT systems (e.g. Gdaniec, 1994, with respect to the *Logos* MT system;

Bernth and Gdaniec, 2001). Based on their work with the *PaTrans* system (Ørsnes *et al.*, 1996), a descendant of the *Eurotra* system, Underwood and Jongejan, 2001 provide a definition of translatability:

“the notion of translatability is based on so-called ‘translatability indicators’ where the occurrence of such an indicator in the text is considered to have a negative effect on the quality of machine translation. The fewer translatability indicators, the better suited the text is to translation using MT” (Underwood and Jongejan, 2001:363).

In an integrated translation environment, these definitions would have to be widened considerably. Future translatability indicators will have to be more fine-grained. MT systems will have to have the capacity to adapt to such indicators and learn from them if progress in this area is to be made. Translatability indicators will have to give specific reasons as to why any text might not (yet) be suitable for automatic translation. If this can be achieved, a tool would be triggered to incrementally modify such texts and/or the system’s resources in order to render the text suitable for automatic translation. That is, a hybrid integrated translation environment has to provide a means of separating the translatable parts from the non-translatable parts of a source text in a much more sophisticated manner than TMs currently do. In order to improve translation quality, one would have to estimate:

- the expected quality of the translation;
- the cost (both in terms of human effort as well as price in real terms) of upgrading the system’s resources;
- the cost of amending the source text to make it suitable for processing by an MT system.

Finally, and perhaps more ambitiously, the integrated system would have to be aware of gaps in the source text which it cannot tackle and provide intelligent inference mechanisms to generate solutions for bridging these gaps.

Translations will be routed through the available translation options according to criteria such as the type of text at hand, the value of the information to be translated, the quality requirements of the customer, and the resources in terms of time and money available to them. Finally, it is important that customers receive accurate information on the quality, pricing and time implications of their choice prior to selecting their preferred translation option.

7. SUMMARY

On various occasions in recent decades, MT companies have claimed that the linguistic technology developed by them has made human translation redundant. These claims have so far not had a significant impact on the reality of

translation as a profession and as a business. The only impact such hype has really had is in unrealistically raising the expectation of users as to the quality of such systems, only to disappoint them when they actually interact with MT and turn them away from the undoubted benefits to be made.

The one technology that has had a considerable impact on translation has been TM—it has changed the way translators work, as can be seen when examining the impact it had in the localisation industry, one of the largest employers of technical translators. Ironically, TM technology works without any of the sophisticated linguistic technologies developed over decades by MT developers—it is little more than a sophisticated search and replace engine.

Because of the enormous success of TM systems, large amounts of aligned, bilingual material are now available—exactly how many can only be estimated: individual products, which are frequently translated into 30 languages and more, can easily contain up to one million words. However, the highly successful approach taken by TM developers is also the cause of the inherent restrictions and limitations in TMs.

In order to try and overcome these, we have proposed an implementation of EBMT based on the idea of a phrasal lexicon, a linguistically enhanced version of a TM system working at phrasal level. We propose that EBMT should be viewed as a solution where the requirement is special purpose, high quality translation. Accordingly, we recommend that such systems be restricted to using texts which are suitable for translation in the domain in question. Such texts should be written according to controlled language guidelines in identifiable sublanguage domains. If phrasal lexicons and EBMT systems are used according to such guidelines, we are confident that as more is known about these relatively new fields, the PL will come to play a central role and provide substantial support for special purpose MT. This will bring about a paradigm shift from TM to EBMT via the PL: given their attitude towards such technology, adding an automated translation facility to today's TM environment is likely to raise the hackles of many translators. Rather, our envisaged scenario will demonstrate clearly to translators the benefit of having phrasal alignments made available to them, and once they are happy with this additional mode of operation, it is a reasonably short step to enabling an automated solution via the recombination element of EBMT systems such as those described in this volume.

If EBMT is to find a niche amongst the different MT paradigms, we believe it has to offer the potential to easily adapt to new domains in a more controlled manner than TMs do. The adaptation process differs from TM technology with respect to how translation knowledge is stored, what kind of such knowledge is stored, how it is retrieved and how it is recomposed to build a new translation. This requires sophisticated processing based on linguistic resources and/or advanced numerical processing. We have developed a model for the future use

of translation technology which is intended to be a valuable aid to translators, capable of generating descriptive, controlled or general translations according to the needs of the user and the effort he is willing to invest. Finally, the setting we propose is an hybrid translation environment which integrates together a number of MT technologies, linguistic and processing resources in harmony with the actual translator, the latter remaining *the* most valuable resource in the translation chain.

Acknowledgments

Schaler's initial research on the Phrasal Lexicon was conducted at University College Dublin under Allan Ramsay. It was part-funded by the Irish Industrial Development Agency.

Thanks also to Sharon O'Brien (*DCU*) for bringing one of the authors up to speed on recent developments in the area of translatability, as well as providing many useful comments on an early draft of this paper.

Notes

1. Becker's notion of 'compositionality' is different from that used in the context of semantics or, in fact, that used in rule-based MT (RBMT), i.e. that a target structure can be *generated* in a compositional manner following a detailed analysis of the source structure and the establishment of correspondences between grammatical descriptions of the source and the target structures.

2. Segments currently used by TM systems can be defined to a certain degree by users of TM systems and can also include text strings defined in documents such as headers or members of lists. Segments, however, can never be defined using linguistic criteria. The only counterexample to this that we are aware of is *MultiCorpora*, cf. <http://www.multicorpora.ca>, which permits segmentation at levels other than the sentence.

3. One point of interest is that while many evaluations of EBMT systems exist, we are unaware of any research into the accuracy of translations where multiple phrases are combined to form one translated segment in the TM environment.

4. Fuzzy matching may also be calculated in terms of characters, of course, or by means of some more sophisticated algorithm.

5. Note, however, that *GM* have since abandoned the use of their Controlled Automotive Service Language (CASL) English. From an objective point of view, this seems hard to understand given that Godden (personal communication, cited in Bernth and Gdaniec, 2001) reports on a contrastive experiment designed to test the translations obtained from a set of texts rewritten to conform to 30 CASL rules against those obtained from the unwritten texts. These were rated by both a translator and an expert bilingual automotive technician. Godden reports "a *very* significant increase in percentage of correct translations for the pre-edited version over the original version, as well as a *very* significant decrease in percentage of incorrect translations" (Bernth and Gdaniec, 2001:207).

6. Note that compared to translations produced by humans, any translations derived via MT which are inserted into the TM are automatically penalized by *Trados* and other TM systems. Thus an exact match of the input string against a source string whose translation was obtained automatically would not be deemed an 100% match.

In this context, note that *Trados* offers a special filter to permit translations input from *Systran*, which is complemented by *Systran's* special import format for *Trados*. *SDLX* offers MT functionality (via *Transcend*) within the TM environment itself. Of course, in principle any TM tool may be combined with any MT engine if the user is sufficiently patient to manipulate input and output files using the Translation Memory Exchange (TMX) format (Shuttleworth, 2002).

7. <http://www.hltcentral.org/projects/detail.php?acronym=otelo>

References

- Adriens, G. and D. Schreurs (1992). 'From Cogram to Alcogram: toward a controlled English grammar checker'. In *COLING: Proceedings of 14th International Conference on Computational Linguistics*, Nantes, France, pp.595–601.
- Almqvist, I. and A. Sagvall Hein (1996). 'Defining ScaniaSwedish - a Controlled Language for Truck Maintenance'. In *Proceedings of the First International Workshop on Controlled Language Applications (CLAW 96)*, Leuven, Belgium, pp.159–164.
- Barthe, K. (1998). 'GIFAS Rationalised French: Designing one Controlled Language to Match Another'. In *Proceedings of the Second International Workshop on Controlled Language Applications (CLAW98)*, Pittsburgh, PA, pp.87–102.
- Becker, J.D. (1975). 'The Phrasal Lexicon'. In *Proceedings of Theoretical Issues in Natural Language Processing*, Cambridge, MA., pp.70–73.
- Bernth, A. and C. Gdaniec (2001). 'MTranslatability'. *Machine Translation* **16**(3):175–218.
- Brown, P., J. Cocke, S. Della Pietra, F. Jelinek, V. Della Pietra, J. Lafferty, R. Mercer and P. Rossin (1990). 'A Statistical Approach to Machine Translation'. *Computational Linguistics* **16**:79–85.
- Brown, R. (1996). 'Example-Based Machine Translation in the PanGloss System'. In *COLING: Proceedings of 16th International Conference on Computational Linguistics*, Copenhagen, Denmark, pp.169–174.
- Carl, M. and S. Hansen (1999). 'Linking Translation Memories with Example-Based Machine Translation'. In *Proceedings of Machine Translation Summit VII*, Singapore, pp.617–624.
- Carl, M., L. Iomdin, C. Pease and O. Streiter (1999). 'Towards a Dynamic Linkage of Example-Based and Rule-Based Machine Translation'. In *Proceedings of ISMT and CLIP*, Beijing, China.
- Carl, M., A. Way and R. Schaler (2002). 'Toward a Hybrid, Integrated Translation Environment'. In S. Richardson (ed.) *Machine Translation: From Research to Real Users*, LNAI 2499, Berlin/New York: Springer, pp.11–20.
- Chandioux, J. (1976). 'MÉTÉO: un système opérationnel pour la traduction automatique des bulletins météorologiques destinés au grand public'. *META* **21**:127–133.
- Chen, S. (1996). *Building Probabilistic Models for Natural Language*. PhD Thesis, Harvard University, Cambridge, MA.
- Doi, S. and K. Maraki (1992). 'Translation ambiguity resolution based on text corpora of source and target language'. In *COLING: Proceedings of the 14th International Conference on Computational Linguistics*, Nantes, France.

- Frederking, R. and S. Nirenburg (1994). 'Three Heads are better than one'. In *Proceedings of Applied Natural Language Processing (ANLP)*, Stuttgart, pp.95–100.
- Gale, W.A. and K.W. Church (1991). 'A program for aligning sentences in bilingual corpora'. In *Proceedings of the 29th Annual Meeting of the Association for Computational Linguistics*, Berkeley, CA.
- Gdaniec, C. (1994). 'The Logos Translatability Index'. In *Proceedings of the First Conference for Machine Translation in the Americas*, Columbia, MD, pp.97–105.
- Heyn, M. (1996). 'Integrating machine translation into translation memory systems' In *Proceedings of TKE 96, EAMT Machine Translation Workshop*, pp.111–123.
- Isabelle, P., Dymetman M., Foster G., Jutras J-M. , Macklovitch E., Perrault F., Ren X., and Simard M. (1993). 'Translation Analysis and Translation Automation'. In *Proceedings of the Fifth International Conference on Theoretical and Methodological Issues in Machine Translation*, Kyoto, Japan.
- Kamprath, C., E. Adolphson, T. Mitamura and E. Nyberg (1998). 'Controlled Language Multilingual Document Production: Experience with Caterpillar Technical English.' In *Proceedings of the Second International Workshop on Controlled Language Applications (CLAW98)*, Pittsburgh, PA, pp.51–61.
- Kaplan, R. and J. Bresnan (1982). 'Lexical Functional Grammar: A Formal System for Grammatical for Grammatical Representation'. In J. Bresnan (ed.) *The Mental Representation of Grammatical Relations*, MIT Press, Cambridge, Mass., pp.173–281.
- Kaplan, R., K. Netter, J. Wedekind and A. Zaenen (1989). 'Translation by Structural Correspondences'. In *Proceedings of the Fourth Conference of the EACL*, Manchester, pp.272–281.
- Kenny, D. and A. Way (2001). 'Teaching Machine Translation & Translation Technology: A Contrastive Study'. In M. Forcada, J-A. Pérez-Ortiz & D. Lewis (eds) *Proceedings of the MT Summit Workshop on Teaching Machine Translation*, Santiago de Compostela, Spain, pp.13–17.
- Lehrndorfer, A. and S. Schachtl (1998). 'Controlled Siemens Documentary German and TopTrans'. *TC-FORUM*, Stuttgart, Germany.
- Lux, V. and E. Dauphin (1996). 'Corpus Studies: a Contribution to the Definition of a Controlled Language'. In *Proceedings of the First International Workshop on Controlled Language Applications (CLAW 96)*, Leuven, Belgium, pp.193–204.
- Macklovitch, E. (2000). 'Two types of Translation Memory'. In *Translating and the Computer 22*, Proceedings of the ASLIB Conference, London.
- Macklovitch, E. and G. Russell (2000). 'What's Been Forgotten in Translation Memory'. In *Proceedings of AMTA-2000*, Cuernavaca, Mexico.

- Means, L. and K. Godden (1996). 'The Controlled Automotive Service Language (CASL) Project'. In *Proceedings of the First International Workshop on Controlled Language Applications (CLAW 96)*, Leuven, Belgium, pp.106–114.
- Melamed, D.I. (2001). *Empirical Methods for Exploiting Parallel Texts*. MIT Press, Cambridge, MA.
- Mitamura, T. and E. Nyberg (1995). 'Controlled English for Knowledge Based MT: Experience with the KANT System'. In *Proceedings of Sixth International Conference on Theoretical and Methodological Issues in Machine Translation*, Leuven, Belgium.
- Nagao, M. (1984). 'A framework of a mechanical translation between Japanese and English by analogy principle'. In A. Elithorn and R. Banerji (eds), *Artificial and Human Intelligence*, North Holland, Amsterdam, pp.173–180.
- Nomiyama, H. (1991). 'Lexical selection mechanism using target language knowledge and its learning ability'. In *Proceedings of the Semi-Annual Meeting of the Information Processing Society of Japan*, NL86-8 (in Japanese, cited in [Doi and Maraki, 1992]), Tokyo, Japan.
- Nübel R. (1997). 'End-to-end evaluation in Verbmobil I'. In *Proceedings of MT-Summit VI*, San Diego, CA.
- Ørnsnes, B., B. Music and B. Maegaard (1996). 'PaTrans - A Patent Translation System'. In *COLING: Proceedings of the 16th International Conference on Computational Linguistics*, Copenhagen, Denmark.
- Planas, E. and Furuse O. (1999). Formalizing Translation Memories. In *Proceedings of MT Summit VII*, Singapore, pp.331–339.
- Rayner, M. and P. Bouillon (1995). 'Hybrid Transfer in an English-French Spoken Language Translator'. In *Proceedings of IA'95*, Montpellier, France, pp.153–162.
- Reinke, U. (1999). „Evaluierung der linguistischen Leistungsfähigkeit von Translation Memory-Systemen - Ein Erfahrungsbericht“. In *LDV-Forum. Forum der Gesellschaft für Linguistische Datenverarbeitung (GLDV)*, 1-2.
- Sadler, V. and R. Vendelmans (1990). 'Pilot Implementation of a Bilingual Knowledge Bank'. In *COLING: Proceedings of the 13th International Conference on Computational Linguistics*, Helsinki, Finland, pp.449–451.
- Sato, S. and M. Nagao (1990). 'Toward Memory-based Translation'. In *COLING: Proceedings of the 13th International Conference on Computational Linguistics*, Helsinki, Finland, 1990, pp.247–252.
- Schachtl, S. (1996). 'Requirements for Controlled German in Industrial Applications'. In *Proceedings of the First International Workshop on Controlled Language Applications (CLAW 96)*, Leuven, Belgium, pp.143–149.
- Schaler, R. (1996). 'Machine translation, translation memories and the phrasal lexicon: the localisation perspective'. In *Proceedings of TKE 96, EAMT Machine Translation Workshop*, Vienna, Austria, pp.21–33.

- Schaler, R. (1999). New media localisation - a linglink report for the European Commission DGXIII. Technical report, Luxembourg (1999).
- Simard, M. and P. Langlais (2001). 'Sub-sentential Exploitation of Translation Memories'. In *Proceedings of MT-Summit VIII*, Santiago de Compostela, Spain, pp.335–339.
- Shuttleworth, M. (2002). 'Combining MT and TM on a Technology-oriented Translation Masters: Aims and Perspectives'. In *Proceedings of 6th EAMT Workshop on Teaching Machine Translation*, Manchester, UK, pp.123–129.
- Somers, H.L. (1993). The Current State of Machine Translation. In *Proceedings of MT-Summit IV*, Kobe, Japan, pp.115–124.
- Somers, H.L. (1999). 'Review Article: Example-based Machine Translation'. *Machine Translation* **14**(2):113–157 (also Chapter 4, this volume).
- Sumita, E. and H. Iida (1991). 'Experiments and Prospects of Example-based Machine Translation'. In *Proceedings of the 29th Annual Meeting of the Association for Computational Linguistics*, Berkeley, CA., pp.185–192.
- TETRIS (1999). 'Technologie-Transfer intelligenter Sprachtechnologie'. IAI, Saarbrücken, Germany.
- Underwood, N. and B. Jongejan (2001). 'Translatability Checker: A Tool to Help Decide Whether to Use MT'. In *Proceedings of MT-Summit VIII*, Santiago de Compostela, Spain, pp.363–368.
- Van der Eijk, P., M. de Koning and G. van der Steen (1996). 'Controlled language correction and translation'. In *Proceedings of the First International Workshop on Controlled Language Applications (CLAW 96)*, Leuven, Belgium, pp.64–73.
- Vasconcellos, M. and M. León (1985). 'SPANAM and ENGSPAN: Machine Translation at the Pan American Health Organization'. *Machine Translation* **11**(2-3):122–136.
- Wahlster, W. (1993). 'Verbmobil–Translation of Face-to-Face Dialogs'. In *Proceedings of Machine Translation Summit IV*, Kobe, Japan.
- Wahlster, W. (2000). *Verbmobil : foundations of speech-to-speech translation*. Berlin/New York: Springer.
- Way, A. (2001). *LFG-DOT: A Hybrid Architecture for Robust MT*. PhD Thesis, University of Essex, Colchester, UK.