

Data-Assisted Controlled Translation

Michael Carl

Institut für Angewandte Informationsforschung,
Martin-Luther-Straße 14,
66111 Saarbrücken, Germany,
carl@iai.uni-sb.de

Abstract

In this paper we try to approach controlled translation from a data-driven perspective. Two fundamental and interwoven problems occur for which we suggest different solutions: (i) the complexity of inducing translation knowledge from aligned text can be reduced when strictly discriminating between compositional and non-monotonous translations; (ii) the ‘noise’ in reference translations can be eliminated through human interaction. A interaction between the user and translation system requires a completely new set of interfaces to dynamically negotiate meanings and controlled translation equivalences.

1 Introduction

In previous papers (Carl et al., 2002; Schäler et al., 2003) it was claimed that if EBMT is to find a niche amongst the different MT paradigms, it has to offer the potential to easily adapt to new domains in a more controlled manner than other MT systems (and translation memories) currently do. In another paper presented in this conference (Carl and Langlais, 2003), it is shown that an aligned text and a general purpose dictionary can be combined and processed to quickly obtain special purpose translation knowledge as is required for sublanguage

translation. The produced quality of translations however, lags far behind the requirements of controlled language translation. In this paper we examine the potential for controlled data-based translation from a feasibility point of view.

While former generations of machine translation (MT) systems were taught exactly what to do when, recent MT systems make increasing use of learning methods. These data-driven systems (including CAT systems, such as Translation Memories) exploit reference texts for new translations and develop internal representations over which a trainer or programmer has limited access and control. The developmental capacities of current learning MT systems, however, are very restricted: once the architecture of the system is fixed and an initial process of learning has taken place, the developmental capacities are exhausted without much possibility of further enhancement.

In this paper we shall argue for a more open learning architecture where translation equivalences can be dynamically negotiated with the user. This is particularly important and useful for controlled language translation, as controlled languages are used in limited domains where sufficient training data is rarely available and high quality translations are expected.

In section 2, the paper discusses problems occurring when translating controlled languages. The acquisition and maintenance of translation grammars in particular causes major problems in conventional architectures. A claim which is

also iterated in the call-for-papers of this conference has been that these problems could be minimized or circumvented when extracting translation grammars from (controlled) reference translations.

Section 3 examines this claim more closely by considering the complexity when inducing translation grammars from aligned texts. Amongst the huge number of grammars which can be extracted from a given sentence or text, we shall argue in section 4 that for controlled translation a homomorphic translation grammar is appropriate where every transfer rule is (i) unique and (ii) re-duplicated in a compositional and in a non-compositional manner. Given the ‘noise’ even in controlled translations, we suggest that the translation grammar should be negotiated between the user and the system. In order for a data-assisted controlled language translation system to comply with these demands, we suggest a number of properties for the learning component in section 5. Similar to a classroom situation, where a teacher assists the learners in the collaborative construction of their mental model, learning MT systems are desirable which negotiate possible translations with their users. The system builds up a viable representation which enables it to function efficiently and consistently with the controlled translation task at hand.

2 Controlled Languages and Translation

Controlled languages define a writing standard for domain specific documents where linguistic expressions are restricted to a subset of natural language. 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 (Lehrndorfer and Schachtl, 1998, p.8), “the concept of controlled language is a mental offspring of machine translation”. A distinction has been

made between controlled languages to enhance readability and controlled languages to enhance translatability. However, Reuther (2003) finds that rules which enhance readability are a subset of translatability rules.

Controlled languages have been developed for restricted domains, such as technical documentation for repair, maintenance and service documents in large companies (e.g. BMW, Boeing, Scania, GM etc.). Within the Controlled Language Authoring Technology (CLAT, (IAI, 2002)), controlled languages are limited on three levels (i) general language correctness, (ii) lexical limitations and (iii) syntactic limitations (Reuther and Schmidt-Wigger, 2000) where each of these modules can be adjusted to the needs of a particular user and/or a particular (controlled) language.

Caterpillar’s CTE (Caterpillar Technical English), defines constraints on the lexicon, on the complexity of sentences and on the use of generalized markup language. However, when using this controlled language for translation in the KANT 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).

Translating controlled languages involves more than the translation of a controlled language into the target language. As the citation above suggests, adjusting a source language text with a controlled language tool is not sufficient for achieving high quality translation: the transfer into the target language also needs to be controlled and the generation should adhere to the requirements of a controlled target language. A conventional transfer-based MT system typically involves three processing steps:

1. the segmentation and parsing of the source text (i.e. analysis);
2. the transfer of the source segments into the target language (lexical and structural mapping);
3. the recombination and ordering of the tar-

get language segments according to target language syntax (generation).

Large general purpose MT systems — such as many RBMT systems — can only be tuned to produce controlled translations with considerable difficulty. Translating controlled languages in traditional rule-based system implies controlling and adjusting all three processing steps. As each of these steps requires independent knowledge resources, adjusting a conventional RBMT system to new controlled language is non-trivial. The well-known ‘knowledge-acquisition bottleneck’, i.e. where and how to acquire the necessary resources is multiplied with the problem of how to maintain, adjust and homogenize the different resources.

Data-driven MT systems, in contrast, extract translation knowledge from translated texts, thus avoiding manual updating and balancing of independent knowledge resources.

3 Data-driven Translation

Given a set of alignments, data-driven MT systems try to figure out what pieces of knowledge might have led to these translations and then extract and/or reuse this hidden knowledge for new translations.

Figuring out these pieces of knowledge is a non-trivial task. To estimate the complexity of the task, Wu (Wu, 1995) counts the number of ways an alignment might have been generated based on an inversion transduction grammar. He counts the number of binary mappings where source tokens can map into an empty target element and vice versa.

In this paper we count the number of isomorphic translations of the alignments. Then we examine how many transfer rules can be extracted from these isomorphic translations. This gives an upper bound for the grammar induction task. We will claim that only a small subset of the possible transfer rules are required.

3.1 Isomorphic Translation

Isomorphic translations, as examined here, are not restricted to binary derivations and they do not allow the mapping of an empty element: in

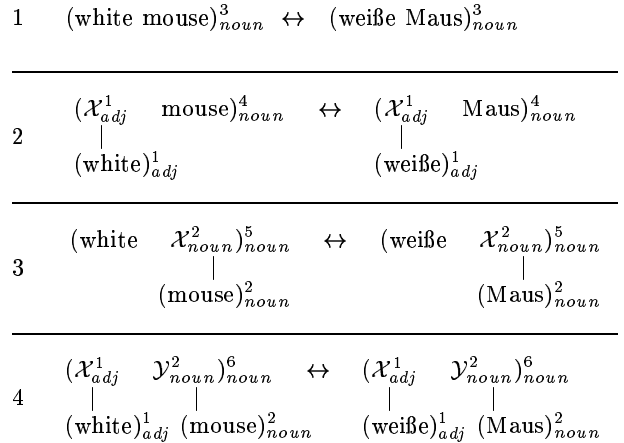


Figure 1: Four Isomorphic Translations of Alignment 1




Figure 2: Homomorphic Translation Grammar

an isomorphic translation every partial derivation of the source side has a translation in the target language with a derivation which has the same depth and arity as the source derivation tree.

We extract translation knowledge from alignments in two steps (i) generate at least one isomorphic translation for each alignment and (ii) extract homomorphic transfer rules from these isomorphic translations. In addition we want the extracted transfer rules to be as compositional and unique as possible.¹

For the sake of explanation, assume the following translation equivalence 1 has been aligned and extracted from a bilingual text:

$$1 \text{ white mouse} \leftrightarrow \text{weiße Maus}$$

This alignment can be generated though four different isomorphic translations which are shown in figure 1. Note that the translations

¹We refer the interested reader to an algorithm presented and discussed in (Carl and Langlais, 2003), in these proceedings.

from top to bottom become more compositional, as the trees contain more internal nodes. While the trees in translation 1 have a depth of 0, translations 2 to 4 have depth 2. The arity of the derivation equals the number of variables as explained below.

From the four isomorphic translations shown in figure 1 can be extracted 6 homomorphic transfer rules² which are shown in figure 2. Transfer rules 1, 2 and 3 contain only lexical tokens. These rules have the arity 0. Transfer rules 4, 5 and 6 contain variables which represent substituted isomorphic derivation trees. Rules 4 and 5 have arity 1 since they contain one variable while rule 6 has 2 variables and thus arity 2. Note that the superscribed indexes of the variables in the derivation trees of figure 1 coincide with the enumeration of the extracted transfer rules in figure 2. Note also that rules 1, 2 and 6 generate the most compositional translation in figure 1, while transfer rule 3 generates a non-compositional translation.

3.2 Reducing Complexity

As shown in figure 3, the number of isomorphic translations $I(n)$ increases exponentially with the length n of the alignment.³ An alignment of length 2 (such as the one above) has 4 isomorphic translations, an alignment of length $n = 3$ has 24 isomorphic translations and an alignment with 20 tokens has more than 10^{18} isomorphic translations.

As can be seen in figure 3, the number of possible isomorphic translations grows to the order of 5^n more steeply than the number of extracted transfer rules. A translation grammar which generates all isomorphic translations for an alignment is exhaustive. Such an exhaustive translation grammar is suggested by the DOT approach (Poutsma, 2000). To generate

²In theory, many more transfer rules could be possible, e.g. “white mouse $\leftrightarrow \epsilon$ ”, “ $\epsilon \leftrightarrow$ weiße Maus”; “white \leftrightarrow Maus”, “mouse \leftrightarrow weiße”, etc. While the former transfer rules are unlikely, the latter ones are probably inconsistent with larger sets of alignments. We will, therefore, exclude such transfer rules from further complexity consideration.

³Where n is the minimum number of tokens in the source or the target side of the alignment.

n	$I(n)$	$T(n)$
1	1	1
2	4	6
3	24	18
4	176	45
5	1440	103
6	12608	224
7	115584	472
8	1095424	975
9	10646016	1989
10	105522176	4026
11	1062623232	8110
12	10840977408	16289
13	111811534848	32659
14	1163909087232	64412
15	12212421230592	130932
16	129027376349184	261987
17	1371482141884416	524113
18	14656212306231296	1048382
19	157369985643577344	2096938
20	1696975718802522112	4194069
21	18369603773021552639	8388351

$$\frac{I(n)}{T(n)} > \frac{10^{n-2}}{4 * 2^n} = \frac{5^n}{4 * 10^2} = O(5^n)$$

Figure 3: Number of Isomorphic Translations $I(n)$ and Transfer Rules $T(n)$ for an alignment of length n

translations for a new input sentence based on an exhaustive translation grammar, all possible partial target derivations are generated and the most likely translations are computed based on probabilities associated to partial derivations. The shortcoming of this approach is exponential computation time. A linguistically richer version of the DOT approach is presented by Way (Way, 2003), which does not only predict correct translations with higher probability but also reduces the number of generated derivations by making use of LFG-feature structures.

As Turcato & Popowich (Turcato and Popowich, 2003) point out, translations can be compositional and/or monotonous. A translation is compositional if each source token trans-

lates exactly into one target token. A translation is non-monotonous if it is only partially compositional. However it is hard to know for an uninformed learner up to what extent a translation is compositional and when or whether it starts to become non-monotonic (Carl and Langlais, 2003). Assume, for instance, the sub-sentential alignment 2.

2 clean record \leftrightarrow weiÙe Weste

Four isomorphic translations can be generated analogue to those in figure 1 and a translation grammar similar to the one in figure 2 can be extracted. Since—without any prior knowledge⁴—a learner cannot know which of the two alignments is compositional and which is not, we are bound to keep at least two translations for each alignment: the compositional and the non-compositional. All other transfer rules, i.e. the transfer rules 4 and 5 in figure 2, can be discarded from the translation grammar since they add no further translation knowledge. This does not imply that every rule contains only lexical token or variables. In fact, in (Carl and Langlais, 2003) we extract only the highest weighted transfer rules such that non-monotonous translation template, such as 3 below, could be extracted. In translation template 3 the variables translate compositionally while the remaining tokens do not.

3 \mathcal{X}_{noun} is afraid of \mathcal{Y}_{noun} \leftrightarrow
 \mathcal{X}_{noun} hat Angst vor \mathcal{Y}_{noun}

The translation grammar in figure 4 shows the merged most and least compositional grammars for the two alignments. Note that the number of extracted transfer rules decreases to at most $3n - 2$: $2n - 1$ lexical transfer rules and $n - 1$ translation templates.

4 Data-driven Controlled Translation

According to the description given in section 2, a controlled language implies a vocabulary with

⁴The algorithm discussed in (Carl and Langlais, 2003) can be fed with external knowledge which marks idiomatic expressions as non-compositional.

1a	(white) _{adj}	\leftrightarrow	(weiÙe) _{adj}
1b	(clean) _{adj}	\leftrightarrow	(weiÙe) _{adj}
2a	(mouse) _{noun}	\leftrightarrow	(Maus) _{noun}
2b	(record) _{noun}	\leftrightarrow	(Weste) _{noun}
3a	(white mouse) _{noun}	\leftrightarrow	(weiÙe Maus) _{noun}
3b	(clean record) _{noun}	\leftrightarrow	(weiÙe Weste) _{noun}
6	$(\mathcal{X}_{adj} \mathcal{Y}_{noun})_{noun}$	\leftrightarrow	$(\mathcal{X}_{adj} \mathcal{Y}_{noun})_{noun}$

Figure 4: Merged Translation Grammar

well-defined meanings and a simplified grammar. That is, each expression in a controlled language can be disambiguated such that it has only one meaning. If we assume that we are translating from a controlled source language into a controlled target language, this holds true for the source and the target language. Controlled translation thus performs a mapping from well-defined source language meanings to well-defined target language meanings.

To ensure well-defined mappings from source to target we need to avoid ambiguous transfer rules. In the grammar in figure 4, German “weiÙe” has two different English translations “clean” and “white”. If we delete one of the rules 1a or 1b from the translation grammar, “weiÙe” has an unambiguous default translation while the contexts in which “weiÙe” is differently translated are explicitly listed in the grammar.

A number of alternative methods have been proposed to disambiguate transfer rules. Menezes (Menezes, 2002), for instance, proposes a machine-learning approach which disambiguates conflicting transfer mappings, using linguistic features contained in the parent and child nodes of the source language. In any case, particularly for controlled translation a more restricted grammar is required than the one shown in the grammar in figure 4.

There are thus two assumptions for controlled translation: (i) given an appropriate context, each expression has exactly one meaning which is represented in one derivation tree and (ii) every source language derivation has exactly one translation into the target language.

In practice, however, it is not always possible to derive a set of invertible mappings without external help from two (or more) alignments. The aim of data-assisted controlled translation

will thus be to approximate these ideal conditions and to provide support in the elaboration of the necessary unambiguous resources. For instance, assume the following examples 4 and 5:

- 4 Locate the outer cable. ↔
Außenseil befestigen.
- 5 Secure the outer cable. ↔
Außenseil befestigen.

Provided there is no further context, a decision has to be taken how to translate the German sentence “Außenseil befestigen.” into English. This decision can be based on the likelihood of the possible translations, for instance by counting their occurrences in a reference text. The decision could also be taken in collaboration with a translator or terminologist where the system is taught the conditions which clarify when to use “Locate” and when to use “Secure”. Given the choice and time, a translator (or terminologist) might eventually come up with completely new solutions such as 6 and 7. The solution is likely to have different translations for different expressions and same translations for the same expressions and the translation grammar will become more compositional.

- 6 Locate the outer cable. ↔
Außenseil fixieren.
- 7 Secure the outer cable. ↔
Außenseil sichern.

Note that this process also leads to a revision and refinement of concepts and to the emergence of new words and formulations in the grammar. There is thus a mutual refinement of concepts and representation in the controlled languages, the translation system and the user. This refinement cycle is however fundamentally different from the situation in rule-based controlled translation as described in section 2, as the interaction between the user and the system plays a fundamentally different role in adaptive data-driven MT system and a rule-based MT-system. This changing role of user and system is reflected in the distinction between translation knowledge and translator knowledge as described by Kiraly

(Kiraly, 2000).

5 Translator Knowledge

In his “social constructivist approach to translator education”, Don Kiraly (Kiraly, 2000) distinguishes between **translation** knowledge and **translator** knowledge. While his book addresses teachers of human translators, one can equally apply this distinction to a data-driven MT scenario. According to Kiraly, translation knowledge implies the ability to produce an adequate target text for speakers of another language on the basis of the original source text. We will assume that this is—to a certain extent—the case for current MT technology.

Translator knowledge, in addition, entails a “professional self-concept”, a profound awareness of the responsibility as an active participant in a complex communicative process. Translator competence involves the capacity to join a number of new communities conversant in specialized technical fields. As a consequence, a successful translator is able to act successfully within parallel expert communities in different linguistic-cultural communities.

It is this capacity which one would also expect from a data-driven controlled machine translation system: the ability to adjust quickly to different domains and controlled languages, sensitive to detecting inconsistencies in reference translations thus expressing “responsibility” and “professional self-concept” while interacting with a user.

With evolving expectations and a changing architecture of computer systems the definitions of knowledge and learning are also changing. Kiraly contrasts the traditional “transmissionist” versus and a modern “transformist” perspective for which he enumerates some characteristics shown in figure 5.

In addition, he enumerates a number of “key principles of social-constructivist education” which are in line with new requirements and responsibilities of professional translators and translator education. These key principles are equally desirable for new generations of data-assisted controlled translation systems and

Transmissionist Perspective (Rule-Based MT)	Transformist Perspective (Data-driven MT)
knowledge is transferred from teacher to student	knowledge is constructed individually in the learner
knowledge is content acquisition of fixed facts	knowledge is a process, creating personal meanings
knowledge is public and equal for all students	knowledge is private personal and idiosyncratic
learning is molecular in discrete chunks	learning is holistic in complex slices of reality
learning characteristics are shared	every learner is unique
learning is individual	learning is social

Figure 5: Transmissionist versus Transformist Perspective

discussed in the following subsections.

5.1 Scaffolding

Scaffolding refers to the support of a teacher to assist learners in a collaborative construction of their mental models.

Scaffolding is not just any assistance which helps a learner to accomplish a task. It is help which will enable a learner to accomplish a task which they would not have been quite able to manage on their own, and it is help which is intended to bring the learner closer to a state of competence which will enable them eventually to complete such a task on their own. (Király, 2000, p.97)

Applied to the translation task, this principle requires the system (i.e. the learner) to integrate different pieces of knowledge to produce translation for new texts. In addition the system must have a minimum amount of learning capacity to generalize the information provided and to create an internal representation of the task. It must be able to grade the information, check its consistency and interact with the teacher (i.e. the user) to communicate gaps or major uncertainties.

5.2 Viability

Rather than searching for truth, the term viability implies that the constructions of reality are maintained as long as they work for us. This implies that a learner needs to acquire tools which enable him to function efficiently in a physical and socio-cultural environment. It also implies that learning is dynamic and understanding is gradually and continuously refined.

At first glance, this principle seems to contradict the idea of a controlled language. However, controlled languages, including the entire conception of the controlled language checking tools is flexible and susceptible to change and evolution (as an example see, for instance, Godden (2000)). To reiterate the example from the previous section: while the verb “befestigen” might be completely sufficient in a German controlled language, a finer grained distinction between the two translations of “locate” and “secure” becomes necessary for controlled language translation.

5.3 Appropriation

Appropriation is based on the idea that learning entails the internalization of socio-cultural knowledge. Learning is an interactive constructive process where knowledge is not transferred but instead it is constructed by the individual learner through dialogue with other people in a linguistic community. This suggests that learning is a function of the situation in which it occurs. It implies for a teacher to expand and recast existing structures rather than solely working with concepts that are new and foreign.

Appropriation for data-assisted controlled language translation suggests that a system can be built, extended and tuned to a certain sub-language, domain or task. Once a certain degree of translation competence has been acquired by the system, completely new knowledge resources, translation equivalences, structures or dictionaries from a different domain should not be added to the system. Rather, knowledge resources should be extended smoothly and new data would have to overlap with the already available data. It will also be inappropriate to make very fine-grained distinctions for some

phenomena where the domain is still not well elaborated.

5.4 Zone of proximal development

This concept, originally developed in (Vygotsky, 1994), refers to a “virtual space of potential growth, a window of opportunity that is created within a specific learning situation”. Although learning is a necessary aspect of the process of developing psychological (and other) functions, the developmental processes lag behind the learning processes. A constructivist teacher thus aims “to hold the learners in their zones of proximal development by providing just enough help and guidance, but not too much.” Eventually, the learner develops the appropriate functions and is capable of self-directed assistance without intervention. When a skill has been mastered, it is executed in a automatic manner without control from outside.

This principle implies the capacity of the system to incrementally construct representations for controlled translations from data. The system accumulates a number of basic structures or examples until it can draw viable inferences and develop a deeper understanding and mastering of the task. To reach a higher level of development and performance, the system needs to possess the ability to exploit data in an appropriate way. On the other hand the principle indicates that a user cannot expect the system to produce translations which are beyond its learned possibilities. A user has to provide appropriate and sufficient material to guide the system, such that it can make maximum use of its learning capacities.

6 Conclusion

In the call-for-papers of this conference it was asked why until now only very few attempts have been made where Example-based MT systems have been designed specifically for controlled language application. This paper tries to give an answer and points out future directions for data-driven controlled translation. The paper investigates the complexity for inducing translation grammars from aligned texts and proposes a number of formal properties suitable

to indicate the state of the grammar. We seek to find answers of what Collins and Somers (2003) find to be a lack in current EBMT systems:

EBMT case-bases on the whole tend not to be structured according to any concept classification or discriminatory indexing schemes ...and it is therefore difficult to see how a newly solved problem can be re-entered into the system ...in order to allow the system to learn. General domain knowledge (e.g. knowledge of linguistic rules; the translation domain) is also rarely used to assist in the categorization/generalization of cases. (Collins and Somers, 2003, p. 126)

To make better use of different resources and to integrate reference translations in the controlled translation task we suggest an interactive module, where translation correspondences are dynamically negotiated between the user and the system.

If it is true — as Martin Kay states — that the best way to find out what translation is, is to see what (human) translators do then it is important to investigate how translation is taught in order to design machine translation systems. A particularly insightful approach has recently been put forward by Don Kiraly (Kiraly, 2000) who distinguishes between translation knowledge and translator knowledge. While translation knowledge refers to the ability of generating an adequate target text, translator knowledge requires, in addition, a “professional self concept” and the ability to quickly accommodate and adjust to different sublanguages.

The paper suggests that translator knowledge is exactly what is required for controlled machine translation: a learning MT system which expresses professional self concept by careful induction of translation knowledge from reference alignments.

Acknowledgement

My warmest thanks to Elliott Macklovitch, Jörg Schütz, Ursula Reuther and an anonymous reviewer for critical comments and useful hints.

References

- Michael Carl and Philippe Langlais. 2003. Tuning general purpose translation knowledge to a sub-language. In *Proceedings of EAMT/CLAW*.
- Michael Carl and Andy Way, editors. 2003. *Recent Advances in Example-Based Machine Translation*. Kluwer, Academic Publisher, Boston/Dordrecht/London. in press.
- Michael Carl, Andy Way, and Reinhard Schäler. 2002. Toward a hybrid integrated translation environment. In *Proceedings of AMTA*, pages 11–21.
- Bróna Collins and Harold Somers. 2003. EBMT Seen as Case-based Reasoning. In *(Carl and Way, 2003)*.
- Kurt Godden. 2000. The Evolution of CASL Controlled Authoring at General Motors. In *Proceedings of CLAW*, pages 14–20.
- IAI, Saarbrücken, Germany, 2002. *Controlled Language Authoring Technology*. <http://iai.uni-sb.de/iaide/de/clat.htm>.
- Christine Kamprath, Eric Adolphson, Teruko Mitamura, and Eric Nyberg. 1998. Controlled Language for Multilingual Document Production: Experience with Caterpillar Technical English. In *CLAW '98*, pages 51–62.
- Donald C. Kiraly. 2000. *A social constructivist approach to translator education: empowerment from theory to practice*. St. Jerome Publishing, Manchester, UK.
- Anne Lehrndorfer and Stefanie Schachtl. 1998. Controlled Siemens Documentary German and TopTrans. In *TC-FORUM*. <http://www.tcforum.org/topictr/tr9contr.htm>.
- Arul Menezes. 2002. Better contextual translation using machine learning. In *Machine Translation from Research to Real Users*, pages 124–134, Berlin Heidelberg. Springer.
- Teruko Mitamura and Eric Nyberg. 1995. Controlled english for knowledge-based mt: Experience with the kant system. In *TMI'95*.
- A. Poutsma. 2000. Data-Oriented Translation. In *Proceedings COLING-2000*, Saarbrücken.
- Ursula Reuther and Antje Schmidt-Wigger. 2000. Designing a multi-purpose cl application. In *Proceedings of CLAW*, pages 72–83.
- Ursula Reuther. 2003. Two in one - can it work? readability and translatability by means of controlled language. In *Proceedings of CLAW*.
- Reinhard Schäler, Andy Way, and Michael Carl. 2003. Ebmt in a controlled environment. In *(Carl and Way, 2003)*.
- Davide Turcato and Fred Popowich. 2003. What is Example-Based Machine Translation. In *(Carl and Way, 2003)*.
- Lev S. Vygotsky. 1994. Extracts from *thought and language and mind in society*. In Barry Stierer and Janet Maybin, editors, *Language, Literacy and Learning in Educational Practice*, pages 45–58, Clevedon. Multilingual Matters.
- Andy Way. 2003. Translating with Examples. In *(Carl and Way, 2003)*.
- Dekai Wu. 1995. Grammarless extraction of phrasal translation examples from parallel texts. In *TMI-95*.