Termight: Identifying and Translating Technical Terminology

Ido Dagan* Ken Church

AT&T Bell Laboratories

600 Mountain Ave.

Murray Hill, NJ 07974, USA

dagan@bimacs.cs.biu.ac.il

kwc@research.att.com

Abstract

We propose a semi-automatic tool, *termight*, that helps professional translators and terminologists identify technical terms and their translations. The tool makes use of part-of-speech tagging and word-alignment programs to extract candidate terms and their translations. Although the extraction programs are far from perfect, it isn't too hard for the user to filter out the wheat from the chaff. The extraction algorithms emphasize completeness. Alternative proposals are likely to miss important but infrequent terms/translations. To reduce the burden on the user during the filtering phase, candidates are presented in a convenient order, along with some useful concordance evidence, in an interface that is designed to minimize keystrokes. *Termight* is currently being used by the translators at AT&T Business Translation Services (formerly AT&T Language Line Services).

1 Terminology: An Application for Natural Language Technology

The statistical corpus-based renaissance in computational linguistics has produced a number of interesting technologies, including part-of-speech tagging and bilingual word alignment. Unfortunately, these technologies are still not as widely deployed in practical applications as they might be. Part-of-speech taggers are used in a few applications, such as speech synthesis (Sproat et al., 1992) and question answering (Kupiec, 1993b). Word alignment is newer, found only in a few places (Gale and Church, 1991a; Brown et al., 1993; Dagan et al., 1993). It is used at IBM for estimating parameters of their statistical machine translation prototype (Brown et

al., 1993). We suggest that part of speech tagging and word alignment could have an important role in glossary construction for translation.

Glossaries are extremely important for translation. How would Microsoft, or some other software vendor, want the term "Character menu" to be translated in their manuals? Technical terms are difficult for translators because they are generally not as familiar with the subject domain as either the author of the source text or the reader of the target text. In many cases, there may be a number of acceptable translations, but it is important for the sake of consistency to standardize on a single one. It would be unacceptable for a manual to use a variety of synonyms for a particular menu or button. Customarily, translation houses make extensive job-specific glossaries to ensure consistency and correctness of technical terminology for large jobs.

A glossary is a list of terms and their translations.¹ We will subdivide the task of constructing a glossary into two subtasks: (1) generating a list of terms, and (2) finding the translation equivalents. The first task will be referred to as the monolingual task and the second as the bilingual task.

How should a glossary be constructed? Translation schools teach their students to read as much background material as possible in both the source and target languages, an extremely time-consuming process, as the introduction to Hann's (1992, p. 8) text on technical translation indicates:

Contrary to popular opinion, the job of a technical translator has little in common with other linguistic professions, such as literature translation, foreign correspondence or interpreting. Apart from an expert knowledge of both languages..., all that is required for the latter professions is a few general dictionaries, whereas a technical translator needs a whole library of specialized dictionaries, encyclopedias and

*Author's current address: Dept. of Mathematics and Computer Science, Bar Ilan University, Ramat Gan 52900, Israel.

¹The source and target fields are standard, though many other fields can also be found, e.g., usage notes, part of speech constraints, comments, etc.

technical literature in both languages; he is more concerned with the exact meanings of terms than with stylistic considerations and his profession requires certain 'detective' skills as well as linguistic and literary ones. Beginners in this profession have an especially hard time... This book attempts to meet this requirement.

Unfortunately, the academic prescriptions are often too expensive for commercial practice. Translators need just-in-time glossaries. They cannot afford to do a lot of background reading and "detective" work when they are being paid by the word. They need something more practical.

We propose a tool, *termight*, that automates some of the more tedious and laborious aspects of terminology research. The tool relies on part-of-speech tagging and word-alignment technologies to extract candidate terms and translations. It then sorts the extracted candidates and presents them to the user along with reference concordance lines, supporting efficient construction of glossaries. The tool is currently being used by the translators at AT&T Business Translation Services (formerly AT&T Language Line Services).

Termight may prove useful in contexts other than human-based translation. Primarily, it can support customization of machine translation (MT) lexicons to a new domain. In fact, the arguments for constructing a job-specific glossary for human-based translation may hold equally well for an MT-based process, emphasizing the need for a productivity tool. The monolingual component of *termight* can be used to construct terminology lists in other applications, such as technical writing, book indexing, hypertext linking, natural language interfaces, text categorization and indexing in digital libraries and information retrieval (Salton, 1988; Cherry, 1990; Harding, 1982; Bourigault, 1992; Damerau, 1993), while the bilingual component can be useful for information retrieval in multilingual text collections (Landauer and Littman, 1990).

2 Monolingual Task: An Application for Part-of-Speech Tagging

Although part-of-speech taggers have been around for a while, there are relatively few practical applications of this technology. The monolingual task appears to be an excellent candidate. As has been noticed elsewhere (Bourigault, 1992; Justeson and Katz, 1993), most technical terms can be found by looking for multiword noun phrases that satisfy a rather restricted set of syntactic patterns. We follow Justeson and Katz (1993) who emphasize the importance of term frequency in selecting good candidate terms. An expert terminologist can then skim the list of candidates to weed out spurious candidates and cliches.

Very simple procedures of this kind have been remarkably successful. They can save an enormous amount of time over the current practice of reading the document to be translated, focusing on tables, figures, index, table of contents and so on, and writing down terms that happen to catch the translator's eye. This current practice is very laborious and runs the risk of missing many important terms.

Termight uses a part of speech tagger (Church, 1988) to identify a list of candidate terms which is then filtered by a manual pass. We have found, however, that the manual pass dominates the cost of the monolingual task, and consequently, we have tried to design an interactive user interface (see Figure 1) that minimizes the burden on the expert terminologist. The terminologist is presented with a list of candidate terms, and corrects the list with a minimum number of key strokes. The interface is designed to make it easy for the expert to pull up evidence from relevant concordance lines to help identify incorrect candidates as well as terms that are missing from the list. A single key-press copies the current candidate term, or the content of any marked emacs region, into the upper-left screen. The candidates are sorted so that the better ones are found near the top of the list, and so that related candidates appear near one another.

2.1 Candidate terms and associated concordance lines

Candidate terms. The list of candidate terms contains both multi-word noun phrases and single words. The multi-word terms match a small set of syntactic patterns defined by regular expressions and are found by searching a version of the document tagged with parts of speech (Church, 1988). The set of syntactic patterns is considered as a parameter and can be adopted to a specific domain by the user. Currently our patterns match only sequences of nouns, which seem to yield the best hit rate in our environment. Single-word candidates are defined by taking the list of all words that occur in the document and do not appear in a standard stop-list of "noise" words.

Grouping and sorting of terms. The list of candidate terms is sorted to group together all noun phrase terms that have the same head word (as in Figure 1), which is simply the last word of the term for our current set of noun phrase patterns. The order of the groups in the list is determined by decreasing frequency of the head word in the document, which usually correlates with the likelihood that this head word is used in technical terms.

Sorting within groups. Under each head word the terms are sorted alphabetically according to reversed order of the words. Sorting in this order reflects the order of modification in simple English noun phrases and groups together terms that denote different modifications of a more general term (see

Figure 1: The monolingual user interface consists of three screens: (1) the input list of candidate terms (upper right), (2) the output list of terms, as constructed by the user (upper left), and (3) the concordance lines associated with the current term, as indicated by the cursor position in screen 1. Typos are due to OCR errors. Underscores denote line breaks.

for example the terms *default paper size*, *paper size* and *size* in Figure 1).

Concordance lines. To decide whether a candidate term is indeed a term, and to identify multi-word terms that are missing from the candidate list, one must view relevant lines of the document. For this purpose we present a concordance line for each occurrence of a term (a text line centered around the term). If, however, a term, t_1 , (like 'point') is contained in a longer term, t_2 , (like 'insertion point' or 'decimal point') then occurrences of t_2 are not displayed for t_1 . This way, the occurrences of a general term (or a head word) are classified into disjoint sets corresponding to more specific terms, leaving only unclassified occurrences under the general term. In the case of 'point', for example, five specific terms are identified that account for 61 occurrences of 'point', and accordingly, for 61 concordance lines. Only 20 concordance lines are displayed for the word 'point' itself, and it is easy to identify in them 5 occurrences of the term 'starting point', which is missing from the candidate list (because 'starting' is tagged as a verb). To facilitate scanning, concordance lines are sorted so that all occurrences of identical preceding contexts of the head word, like 'starting', are grouped together. Since all the words of the document, except for stop list words, appear in the candidate list as single-word terms it is guaranteed that every term that was missed by the au-

tomatic procedure will appear in the concordance lines.

In summary, our algorithm performs the following steps:

- Extract multi-word and single-word candidate terms.
- Group terms by head word and sort groups by head-word frequency.
- Sort terms within each group alphabetically in reverse word order.
- Associate concordance lines with each term. An occurrence of a multi-word term t is not associated with any other term whose words are found in t .
- Sort concordance lines of a term alphabetically according to preceding context.

2.2 Evaluation

Using the monolingual component, a terminologist at AT&T Business Translation Services constructs terminology lists at the impressive rate of 150-200 terms per hour. For example, it took about 10 hours to construct a list of 1700 terms extracted from a 300,000 word document. The tool has at least doubled the rate of constructing terminology lists, which was previously performed by simpler lexicographic tools.

2.3 Comparison with related work

Alternative proposals are likely to miss important but infrequent terms/translations such as 'Format Disk dialog box' and 'Label Disk dialog box' which occur just once. In particular, mutual information (Church and Hanks, 1990; Wu and Su, 1993) and other statistical methods such as (Smadja, 1993) and frequency-based methods such as (Justeson and Katz, 1993) exclude infrequent phrases because they tend to introduce too much noise. We have found that frequent head words are likely to generate a number of terms, and are therefore more important for the glossary (a "productivity" criterion). Consider the frequent head word *box*. In the Microsoft Windows manual, for example, almost any type of box is a technical term. By sorting on the frequency of the headword, we have been able to find many infrequent terms, and have not had too much of a problem with noise (at least for common headwords).

Another characteristic of previous work is that each candidate term is scored independently of other terms. We score a group of related terms rather than each term at a time. Future work may enhance our simple head-word frequency score and may take into account additional relationships between terms, including common words in modifying positions.

Termight uses a part-of-speech tagger to identify candidate noun phrases. Justeson and Katz (1993) only consult a lexicon and consider all the possible parts of speech of a word. In particular, every word that can be a noun according to the lexicon is considered as a noun in each of its occurrences. Their method thus yields some incorrect noun phrases that will not be proposed by a tagger, but on the other hand does not miss noun phrases that may be missed due to tagging errors.

3 Bilingual Task: An Application for Word Alignment

3.1 Sentence and word alignment

Bilingual alignment methods (Warwick et al., 1990; Brown et al., 1991a; Brown et al., 1993; Gale and Church, 1991b; Gale and Church, 1991a; Kay and Roscheisen, 1993; Simard et al., 1992; Church, 1993; Kupiec, 1993a; Matsumoto et al., 1993; Dagan et al., 1993). have been used in statistical machine translation (Brown et al., 1990), terminology research and translation aids (Isabelle, 1992; Ogden and Gonzales, 1993; van der Eijk, 1993), bilingual lexicography (Klavans and Tzoukermann, 1990; Smadja, 1992), word-sense disambiguation (Brown et al., 1991b; Gale et al., 1992) and information retrieval in a multilingual environment (Landauer and Littman, 1990).

Most alignment work was concerned with alignment at the sentence level. Algorithms for the more

difficult task of word alignment were proposed in (Gale and Church, 1991a; Brown et al., 1993; Dagan et al., 1993) and were applied for parameter estimation in the IBM statistical machine translation system (Brown et al., 1993).

Previously translated texts provide a major source of information about technical terms. As Isabelle (1992) argues, "Existing translations contain more solutions to more translation problems than any other existing resource." Even if other resources, such as general technical dictionaries, are available it is important to verify the translation of terms in previously translated documents of the same customer (or domain) to ensure consistency across documents.

Several translation workstations provide sentence alignment and allow the user to search interactively for term translations in aligned archives (e.g. (Ogden and Gonzales, 1993)). Some methods use sentence alignment and additional statistics to find candidate translations of terms (Smadja, 1992; van der Eijk, 1993).

We suggest that word level alignment is better suitable for term translation. The bilingual component of *termight* gets as input a list of source terms and a bilingual corpus aligned at the word level. We have been using the output of *word_align*, a robust alignment program that proved useful for bilingual concordancing of noisy texts (Dagan et al., 1993). *Word_align* produces a partial mapping between the words of the two texts, skipping words that cannot be aligned at a given confidence level (see Figure 2).

3.2 Candidate translations and associated concordance lines

For each occurrence of a source term, *termight* identifies a *candidate translation* based on the alignment of its words. The candidate translation is defined as the sequence of words between the first and last target positions that are aligned with any of the words of the source term. In the example of Figure 2 the candidate translation of *Optional Parameters box* is *zone Parametres optionnels*, since *zone* and *optionnels* are the first and last French words that are aligned with the words of the English term. Notice that in this case the candidate translation is correct even though the word *Parameters* is aligned incorrectly. In other cases alignment errors may lead to an incorrect candidate translation for a specific occurrence of the term. It is quite likely, however, that the correct translation, or at least a string that overlaps with it, will be identified in some occurrences of the term.

Termight collects the candidate translations from all occurrences of a source term and sorts them in decreasing frequency order. The sorted list is presented to the user, followed by bilingual concordances for all occurrences of each candidate translation (see Figure 3). The user views the concordances to verify correct candidates or to find translations that are

Figure 2: An example of *word_align*'s output for the English and French versions of the Microsoft Windows manual. The alignment of *Parameters* to *optionnels* is an error.

missing from the candidate list. The latter task becomes especially easy when a candidate overlaps with the correct translation, directing the attention of the user to the concordance lines of this particular candidate, which are likely to be aligned correctly. A single key-stroke copies a verified candidate translation, or a translation identified as a marked emacs region in a concordance line, into the appropriate place in the glossary.

3.3 Evaluation

We evaluated the bilingual component of *termight* in translating a glossary of 192 terms found in the English and German versions of a technical manual. The correct answer was often the first choice (40%) or the second choice (7%) in the candidate list. For the remaining 53% of the terms, the correct answer was always somewhere in the concordances. Using the interface, the glossary was translated at a rate of about 100 terms per hour.

3.4 Related work and issues for future research

Smadja (1992) and van der Eijk (1993) describe term translation methods that use bilingual texts that were aligned at the sentence level. Their methods find likely translations by computing statistics on term cooccurrence within aligned sentences and selecting source-target pairs with statistically significant associations. We found that explicit word alignments enabled us to identify translations of infrequent terms that would not otherwise meet statistical significance criteria. If the words of a term occur at least several times in the document (regardless of the *term* frequency) then *word_align* is likely to align them correctly and *termight* will identify the correct translation. If only some of the words of a term are frequent then *termight* is likely to identify a translation that overlaps with the correct one, directing the user quickly to correctly aligned concordance lines. Even if all the words of the term were not aligned by *word_align* it is still likely that most concordance lines are aligned correctly based on other words in the near context.

Termight motivates future improvements in word alignment quality that will increase recall and precision of the candidate list. In particular, taking into

account local syntactic structures and phrase boundaries will impose more restrictions on alignments of complete terms.

Finally, *termight* can be extended for verifying translation consistency at the proofreading (editing) step of a translation job, after the document has been translated. For example, in an English-German document pair the tool identified the translation of the term *Controls menu* as *Menu Steuerung* in 4 out of 5 occurrences. In the fifth occurrence *word_align* failed to align the term correctly because another translation, *Steuermenu*, was uniquely used, violating the consistency requirement. *Termight*, or a similar tool, can thus be helpful in identifying inconsistent translations.

4 Conclusions

We have shown that terminology research provides a good application for robust natural language technology, in particular for part-of-speech tagging and word-alignment algorithms. Although the output of these algorithms is far from perfect, it is possible to extract from it useful information that is later corrected and augmented by a user. Our extraction algorithms emphasize completeness, and identify also infrequent candidates that may not meet some of the statistical significance criteria proposed in the literature. To make the entire process efficient, however, it is necessary to analyze the user's work process and provide interfaces that support it. In many cases, improving the way information is presented to the user may have a larger effect on productivity than improvements in the underlying natural language technology. In particular, we have found the following to be very effective:

- Grouping linguistically related terms, making it easier to judge their validity.
- Sorting candidates such that the better ones are found near the top of the list. With this sorting one's time is efficiently spent by simply going down the list as far as time limitations permit.
- Providing quick access to relevant concordance lines to help identify incorrect candidates as well as terms or translations that are missing from the candidate list.

Figure 3: The bilingual user interface consists of two screens. The lower screen contains the constructed glossary. The upper screen presents the current term, candidate translations with their frequencies and a bilingual concordance for each candidate. Typos are due to OCR errors.

- Minimizing the number of required key-strokes.

As the need for efficient knowledge acquisition tools becomes widely recognized, we hope that this experience with *termight* will be found useful for other text-related systems as well.

Acknowledgements

We would like to thank Pat Callow from AT&T Buiseness Translation Services (formerly AT&T Language Line Services) for her indispensable role in designing and testing *termight*. We would also like to thank Bala Satish and Jon Helfman for their part in the project.

References

- Didier Bourigault. 1992. Surface grammatical analysis for the extraction of terminological noun phrases. In *Proc. of COLING*, pages 977–981.
- P. Brown, J. Cocke, S. Della Pietra, V. Della Pietra, F. Jelinek, R.L. Mercer, and P.S. Roossin. 1990. A statistical approach to language translation. *Computational Linguistics*, 16(2):79–85.
- P. Brown, J. Lai, and R. Mercer. 1991a. Aligning sentences in parallel corpora. In *Proc. of the Annual Meeting of the ACL*.
- P. Brown, S. Della Pietra, V. Della Pietra, and R. Mercer. 1991b. Word sense disambiguation using statistical methods. In *Proc. of the Annual Meeting of the ACL*, pages 264–270.
- Peter Brown, Stephen Della Pietra, Vincent Della Pietra, and Robert Mercer. 1993. The mathematics of statistical machine translation: parameter estimation. *Computational Linguistics*, 19(2):263–311.
- L. L. Cherry. 1990. Index. In *Unix Research System Papers*, volume 2, pages 609–610. AT&T, 10 edition.
- Kenneth W. Church and Patrick Hanks. 1990. Word association norms, mutual information, and lexicography. *Computational Linguistics*, 16(1):22–29.
- Kenneth W. Church. 1988. A stochastic parts program and noun phrase parser for unrestricted text. In *Proc. of ACL Conference on Applied Natural Language Processing*.
- Kenneth W. Church. 1993. Char.align: A program for aligning parallel texts at character level. In *Proc. of the Annual Meeting of the ACL*.
- Ido Dagan, Kenneth Church, and William Gale. 1993. Robust bilingual word alignment for machine aided translation. In *Proceedings of the Workshop on Very Large Corpora: Academic and Industrial Perspectives*, pages 1–8.
- Fred J. Damerau. 1993. Generating and evaluating domain-oriented multi-word terms from texts. *Information Processing & Management*, 29(4):433–447.
- William Gale and Kenneth Church. 1991a. Identifying word correspondence in parallel text. In *Proc.*

- of the DARPA Workshop on Speech and Natural Language.
- William Gale and Kenneth Church. 1991b. A program for aligning sentences in bilingual corpora. In *Proc. of the Annual Meeting of the ACL*.
- William Gale, Kenneth Church, and David Yarowsky. 1992. Using bilingual materials to develop word sense disambiguation methods. In *Proc. of the International Conference on Theoretical and Methodological Issues in Machine Translation*, pages 101–112.
- Michael Hann. 1992. *The Key to Technical Translation*, volume 1. John Benjamins Publishing Company.
- P. Harding. 1982. Automatic indexing and classification for mechanised information retrieval. BLRDD Report No. 5723, British Library R & D Department, London, February.
- P. Isabelle. 1992. Bi-textual aids for translators. In *Proc. of the Annual Conference of the UW Center for the New OED and Text Research*.
- John Justeson and Slava Katz. 1993. Technical terminology: some linguistic properties and an algorithm for identification in text. Technical Report RC 18906, IBM Research Division.
- M. Kay and M. Roscheisen. 1993. Text-translation alignment. *Computational Linguistics*, 19(1):121–142.
- J. Klavans and E. Tzoukermann. 1990. The bicord system. In *Proc. of COLING*.
- Julian Kupiec. 1993a. An algorithm for finding noun phrase correspondences in bilingual corpora. In *Proc. of the Annual Meeting of the ACL*.
- Julian Kupiec. 1993b. MURAX: A robust linguistic approach for question answering using an on-line encyclopedia. In *Proc. of International Conference on Research and Development in Information Retrieval, SIGIR*, pages 181–190.
- Thomas K. Landauer and Michael L. Littman. 1990. Fully automatic cross-language document retrieval using latent semantic indexing. In *Proc. of the Annual Conference of the UW Center for the New OED and Text Research*.
- Yuji Matsumoto, Hiroyuki Ishimoto, Takehito Utsuro, and Makoto Nagao. 1993. Structural matching of parallel texts. In *Proc. of the Annual Meeting of the ACL*.
- William Ogden and Margarita Gonzales. 1993. Norm – a system for translators. Demonstration at ARPA Workshop on Human Language Technology.
- Gerard Salton. 1988. Syntactic approaches to automatic book indexing. In *Proc. of the Annual Meeting of the ACL*, pages 204–210.
- M. Simard, G. Foster, and P. Isabelle. 1992. Using cognates to align sentences in bilingual corpora. In *Proc. of the International Conference on Theoretical and Methodological Issues in Machine Translation*.
- Frank Smadja. 1992. How to compile a bilingual collocational lexicon automatically. In *AAAI Workshop on Statistically-based Natural Language Processing Techniques*, July.
- Frank Smadja. 1993. Retrieving collocations from text: Xtract. *Computational Linguistics*, 19(1):143–177.
- R. Sproat, J. Hirschberg, and D. Yarowsky. 1992. A corpus-based synthesizer. In *Proceedings of the International Conference on Spoken Language Processing*, pages 563–566, Banff, October. IC-SLP.
- Pim van der Eijk. 1993. Automating the acquisition of bilingual terminology. In *EACL*, pages 113–119.
- S. Warwick, J. Hajic, and G. Russell. 1990. Searching on tagged corpora: linguistically motivated concordance analysis. In *Proc. of the Annual Conference of the UW Center for the New OED and Text Research*.
- Ming-Wen Wu and Keh-Yih Su. 1993. Corpus-based compound extraction with mutual information and relative frequency count. In *Proceedings of ROCLING VI*, pages 207–216.