

# Integrating a Phrase-based SMT Model and a Bilingual Lexicon for Human in Semi-Automatic Acquisition of Technical Term Translation Lexicon

Yohei Morishita    Takehito Utsuro    Mikio Yamamoto

Graduate School of Systems and Information Engineering

University of Tsukuba

Tsukuba, 305-8573, JAPAN

## Abstract

This paper presents an attempt at developing a technique of acquiring translation pairs of technical terms with sufficiently high precision from parallel patent documents. The approach taken in the proposed technique is based on integrating the phrase translation table of a state-of-the-art statistical phrase-based machine translation model, and compositional translation generation based on an existing bilingual lexicon for human use. Our evaluation results clearly show that the agreement between the two individual techniques definitely contribute to improving precision of translation candidates. We then apply the Support Vector Machines (SVMs) to the task of automatically validating translation candidates in the phrase translation table. Experimental evaluation results again show that the SVMs based approach to translation candidates validation can contribute to improving the precision of translation candidates in the phrase translation table.

## 1 Introduction

For both high quality machine and human translation, a large scale and high quality bilingual lexicon is the most important key resource. Since manual compilation of bilingual lexicon requires plenty of time and huge manual labor, in the research area of knowledge acquisition from natural language text, automatic bilingual lexicon compilation have been studied for more than a decade. Techniques invented so far include translation term pair acquisition based on statistical co-occurrence measure from

parallel sentences (Matsumoto and Utsuro, 2000), translation term pair acquisition from comparable corpora (Fung and Yee, 1998), compositional translation generation based on an existing bilingual lexicon for human use (Tonoike et al., 2006), and translation term pair acquisition by collecting partially bilingual texts through the search engine (Huang et al., 2005).

However, most of those techniques invented so far have not been reliable enough in any practical situation of semi-automatically developing a bilingual lexicon. This is especially true in the case of techniques which use resources other than parallel sentences, since searching comparable corpora or the search engine snippet for a translation of a term into another language is much harder compared with when searching parallel sentences for a translation pair. Even in the case of techniques on translation term pair acquisition from parallel sentences, those techniques do not seem to be reliable enough for those who are actually working on semi-automatically or manually compiling a bilingual lexicon using parallel sentences.

For example, we have been working with a Japanese organization which is responsible for translating Japanese patent applications published by the Japanese Patent Office (JPO) into English. Among various document genres where machine and/or human translation of documents is really required in industrial situation, patent document is one of the most important and have substantial impact in a number of practical applications and services, such as cross-lingual patent retrieval and filing patent applications to foreign countries. Here, in the process

Table 1: Distribution of IPC (International Patent Classification) Categories in 1.8M Parallel Patent Sentences

IPC Category	# of documents	%	# of sentences	%
A. Human Necessities	1,606	3.5	41,180	2.4
B. Performing Operations, Transporting	5,948	12.8	165,994	9.2
C. Chemistry, Metallurgy	1,606	3.5	22,933	1.3
D. Textiles, Paper	331	0.7	7,148	0.4
E. Fixed Constructions	255	0.6	5,906	0.3
F. Mechanical Engineering, Lighting, Heating, Weapons	3,941	8.5	113,604	6.3
G. Physics	16,533	35.7	786,650	43.7
H. Electricity	16,127	34.8	642,163	35.7
Total	46,347	100.0	1,798,571	100.0

of patent document translation, a bilingual lexicon of technical terms is one of the crucial resource, and furthermore, it is definitely necessary to continuously update and extend the lexicon as new patent applications including invention of novel technologies and novel technical terms are published by JPO. Therefore, the organization is continuously working on manually extending its Japanese-English lexicon of technical terms by utilizing Japanese-English parallel patent sentences as certain reference text data for searching for a translation of a Japanese technical term into English.

Through our personal communication with the organization, it is claimed that automatic techniques for translation term pair acquisition are mostly useless. This is because it is often necessary to manually validate acquired translation term pairs by referring to parallel sentences, where this validation process usually takes as much time as when without automatic translation term pair acquisition techniques. According to the organization, when employing certain statistical techniques on automatic acquisition of translation pairs of technical terms from parallel patent sentences, the primary requirement is precision rather than recall. This is because when translation candidates suggested by such statistical techniques are with more than 90% precision, it saves time for persons who work on compiling bilingual lexicon to searching for English translation of a Japanese technical term. Even with relatively low recall, the organization has sufficient number of patent documents so that, for many years, they can continue working on compiling bilingual lexicon only by accepting translation candidates highly

confidently suggested by a statistical technique, but rejecting those suggested with less confidence.

Based on such requirement from the organization working on compiling bilingual lexicon of technical terms from parallel patent documents, this paper presents an attempt at developing a technique of acquiring translation pairs of technical terms with sufficiently high precision from parallel patent documents. The approach taken in the proposed technique is based on integrating the phrase translation table of a state-of-the-art statistical phrase-based machine translation model (Koehn et al., 2007), and compositional translation generation based on an existing bilingual lexicon for human use (Tonoike et al., 2006).

In this approach, we first simply evaluate translation candidates in the phrase translation table as well as those generated by compositional translation generation based on an existing bilingual lexicon for human use. We also evaluate agreement between translation candidates from those two individual techniques that are different from each other with respect to their approaches as well as resource used in their approaches. Our evaluation results clearly show that the agreement between the two individual techniques definitely contribute to improving precision of translation candidates. We then apply the Support Vector Machines (SVMs) to the task of automatically validating translation candidates in the phrase translation table, where features from various sources such as translation candidates for each constituent word found in the existing bilingual lexicon for human use, as well as statistics from the whole parallel sentences used for learning the phrase trans-

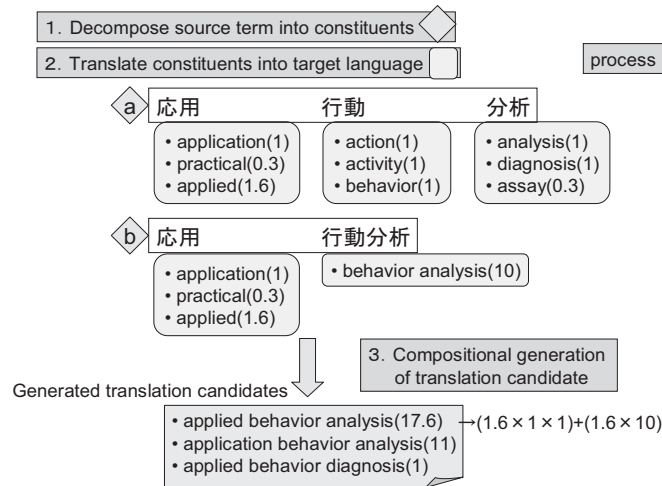


Figure 1: Compositional Translation Generation for the Japanese Technical Term “応用行動分析”

lation table, are incorporated. Experimental evaluation results again show that the SVMs based approach to translation candidates validation can contribute to improving the precision of translation candidates in the phrase translation table.

## 2 Japanese-English Parallel Patent Documents

In the NTCIR-7 workshop, the Japanese-English patent translation task is organized (Fujii et al., 2008), where parallel patent documents and sentences are provided by the organizer. Those parallel patent documents are collected from the 10 years of unexamined Japanese patent applications published by the Japanese Patent Office (JPO) and the 10 years patent grant data published by the U.S. Patent & Trademark Office (USPTO) in 1993-2000. The numbers of documents are approximately 3,500,000 for Japanese and 1,300,000 for English. Because the USPTO documents consist of only patent that have been granted, the number of these documents is smaller than that of the JPO documents.

From these document sets, patent families are automatically extracted and the fields of “Background of the Invention” and “Detailed Description of the Preferred Embodiments” are selected. This is because the text of those fields is usually translated on a sentence-by-sentence basis. Then, the method of (Utiyama and Isahara, 2007) is applied to the text of those fields, and Japanese and English sentences are aligned. Table 1 shows the distribution

of the IPC (International Patent Classification) Categories in the whole parallel documents and sentences (about 1.8M sentences in total).

## 3 Techniques of Generating Translation Candidates

### 3.1 Techniques based on a Bilingual Lexicon for Human Use

#### 3.1.1 A Bilingual Lexicon: Eijiro

As an existing Japanese-English translation lexicon for human use, we use Eijiro (<http://www.eijiro.jp/>, Ver.79, with 1.6M translation pairs).

#### 3.1.2 Compositional Translation Generation

In compositional translation generation (Tonoike et al., 2006), translation candidates of a term are compositionally generated by concatenating the translation of the constituents of the term. Here, as an existing bilingual lexicon for translating constituents, we use Eijiro and bilingual constituents lexicons (0.14M translation pairs) compiled from the translation pairs of Eijiro.

An example of compositional translation generation for the Japanese technical term “応用行動分析” is illustrated in Figure 1. First, the Japanese technical term “応用行動分析” is decomposed into its constituents by consulting an existing bilingual lexicon and retrieving Japanese headwords. In this case, the result of this decomposition can be given as in

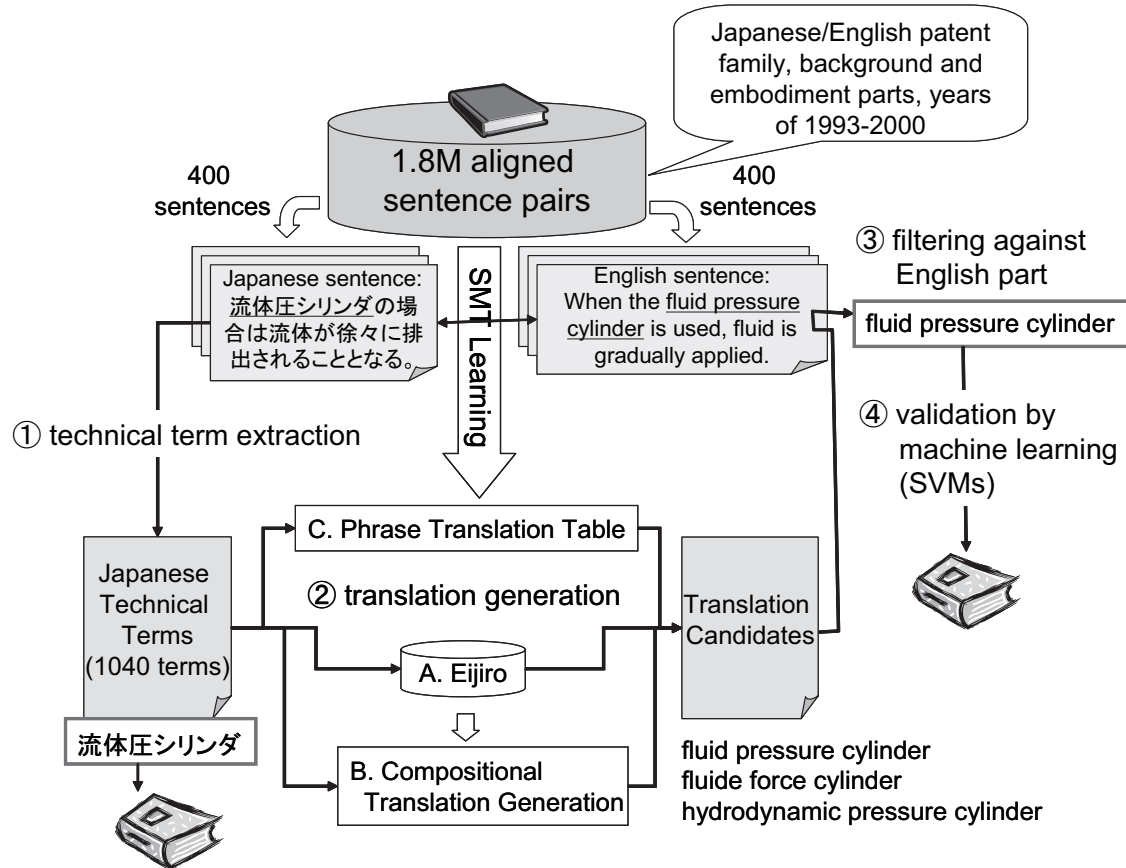


Figure 2: Generating/Filtering/Validating Translation Candidates of Technical Terms in Parallel Patent Documents

the cases “a” and “b” (in Figure 1). Then, each constituent is translated into the target language. A confidence score is assigned to the translation of each constituent. Finally, translation candidates are generated by concatenating the translation of those constituents according to word ordering rules considering prepositional phrase construction.

Each constituent is assigned a score based on the number of morphemes and the frequencies of translation pairs in the bilingual constituent lexicons. Then, the score of the concatenated translation candidates is calculated as the product of the scores of their constituents. When more than one translation candidates are generated as in the case of Figure 1, they are ranked in descending order of their scores.

### 3.2 Phrase Translation Table of an SMT Model

As a toolkit of a phrase-based statistical machine translation model, we use Moses (Koehn et al., 2007) and apply it to the whole 1.8M parallel patent sentences. In Moses, first, word alignment of par-

allel sentences are obtained by GIZA++ (Och and Ney, 2003) in both translation directions and then the two alignments are symmetrised. Next, any phrase pair that is consistent with word alignment is collected into the phrase translation table and a phrase translation probability is assigned to each pair (Koehn et al., 2003). We finally obtain 76M translation pairs with 33M unique Japanese phrases, i.e., 2.29 English translations per Japanese phrase on average, with Japanese to English phrase translation probabilities  $P(p_E | p_J)$  of translating a Japanese phrase  $p_J$  into an English phrase  $p_E$ . For each Japanese phrase, those multiple translation candidates in the phrase translation table are ranked in descending order of Japanese to English phrase translation probabilities.

Table 2: Number of Translation Candidates generated by Individual Techniques

Individual Techniques	# of Japanese Noun Phrases for which English Translation Candidates are Generated	# of Generated English Translation Candidates (rate per a Japanese Noun Phrase)
Eijiro	175	177 (1.01)
Compositional Translation Generation	450	465 (1.03)
Phrase Translation Table	950	2851 (3.00)

## 4 Evaluating Individual Techniques and their Agreements

### 4.1 The Procedure

Out of the whole 1.8M parallel sentences, we randomly select 400 for evaluating translation generation techniques<sup>1</sup>, restricting that they have uniform distribution of IPC categories. Figure 2 illustrates the procedure of generating translation of technical terms in parallel patent sentences. First, we automatically extract noun phrases from Japanese sentences by applying a simple regular expression for noun phrase extraction. Next, we manually extract 1,040 technical terms from those Japanese noun phrases<sup>2</sup>. To those 1,040 Japanese technical terms, the three techniques (i.e., A, B, and C in Figure 2) for generating English translation candidates are applied. Here, suppose that we are given a Japanese noun phrase  $t_J$  extracted from the Japanese sentence  $S_J$  of a parallel sentence pair  $\langle S_J, S_E \rangle$ , and that for  $t_J$ , the techniques for generating English translation candidates are applied. Then, those translation candidates are matched against the English sentence  $S_E$  of the parallel sentence pair, and those which are not found in the English part are filtered out. Finally, Support Vector Machines (SVMs) are applied to the task of

<sup>1</sup>Since our primary application is semi-automatic acquisition of technical term bilingual lexicon from parallel sentences, it is quite usual that a large scale parallel sentences are provided and are used both for learning a phrase translation table and for generating technical term translation pairs. If one wants to consider another task such as acquiring technical term translation pairs that do not appear in the parallel sentences used for learning the phrase translation table, it is necessary to invent a framework slightly different from the one we proposed in this paper.

<sup>2</sup>In a situation of practically applying the technique proposed in this paper, we are planning to use a large scale lexicon of Japanese technical terms when extracting Japanese technical terms for which English translation candidates are to be generated.

validating translation candidates based on features from various sources such as the existing bilingual lexicon for human use and statistics from the whole 1.8M aligned parallel sentences.

For each of the three techniques, Table 2 lists the number of Japanese noun phrases for which the technique can generate English translation candidates, as well as the number of generated English translation candidates. In Figure 3, out of the set (a) of the whole 1,040 Japanese noun phrases, we denote the set of Japanese noun phrases for which Eijiro can generate English translation candidates as  $E$ . We also denote the set of those for which compositional translation generation can generate English translation candidates as  $C$ , and the set of those for which the phrase translation table can generate English translation candidates as  $P$ . We further focus on the set  $(E \cap P)$  of Japanese noun phrases for each of which all of the three techniques can generate the same English translation candidate, and on the set  $(C \cap P) - E$  of Japanese noun phrases for each of which both compositional translation generation and the phrase translation table can generate the same English translation candidate, but the Eijiro can not. We also focus on the set  $P - (C \cap P)$  of Japanese noun phrases for which only the phrase translation table can generate English translation candidates.

For a given Japanese noun phrase, both compositional translation generation and the phrase translation table generate English translation candidates that are ranked in descending order of certain scores or probabilities. As we show in Table 3, in the following sections, we evaluate the 1st ranked translation candidate. On the other hand, only with an exception of a few technical terms, in each entry of Japanese technical terms, Eijiro lists only one English translation<sup>3</sup>. In the case of such exception

<sup>3</sup>Out of the whole 1,040 Japanese noun phrases for evalua-

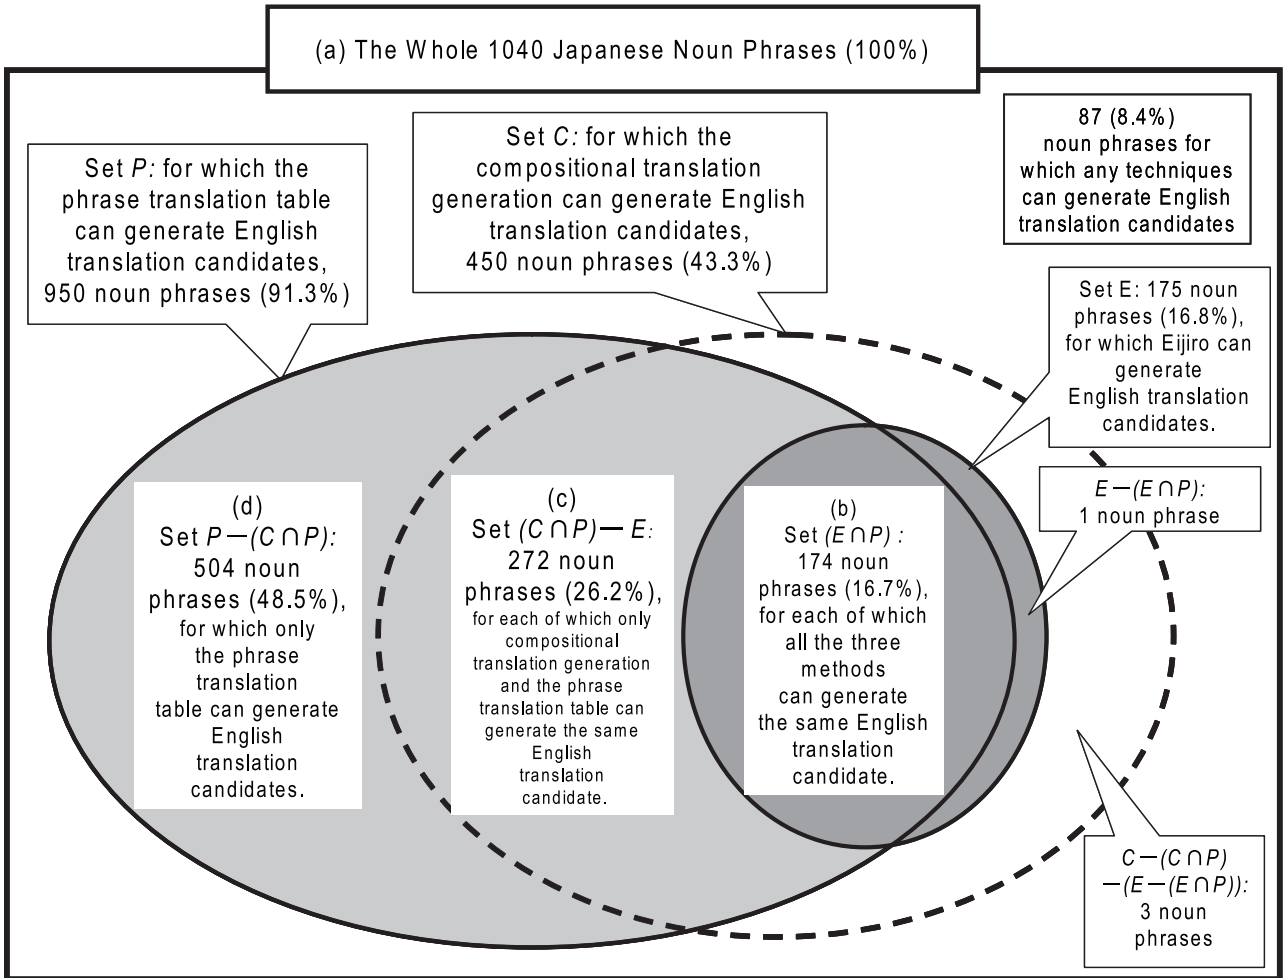


Figure 3: Number of Japanese Noun Phrases for which Each Individual Technique can generate English Translation Candidates

where Eijiro lists more than one English translations for one Japanese technical term, we regard those multiple translations as equally correct in the evaluation of the subsequent sections.

Among the whole procedure in Figure 2, the next section presents results of evaluating the recall/precision/F-measure of the three techniques individually, as well as that of agreement among two or three of the individual techniques. Furthermore, Section 5 presents results of applying SVMs to the task of validating translation candidates.

tion, Eijiro includes 321 of them. Each of those 321 Japanese noun phrases has about 2.31 English translation on the average. However, among those 321 Japanese noun phrases, only 175 have their English translation in the English part of parallel sentences for learning the phrase translation table. Furthermore, almost all of them have only one translation in the English part.

## 4.2 Evaluation Results

In the left half of Table 3, we show results of evaluating each of individual techniques against

- (a) the whole 1,040 Japanese noun phrases,
- (b) the set  $(E \cap P)$ ,
- (c) the set  $(C \cap P) - E$ ,
- (d) the set  $P - (C \cap P)$ .

Against the whole set (a), both Eijiro and compositional translation generation based on Eijiro have very low recall, while their precisions are over 90%. On the other hand, as can be easily expected, the phrase translation table has nearly 80% recall, but its precision is around 87%. When considering our

Table 3: Recall/Precision/F-measure of 1st Ranked Translation Candidates (%)

(a) Individual techniques against the whole 1,040 Japanese noun phrases

Eijiro	Compositional Tran. Generation	Phrase Tran. Table
16.3 (170/1040)	40.3 (419/1040)	79.3 (825/1040)
97.1 (170/175)	93.1 (419/450)	86.8 (825/950)
28.0	56.2	82.9

(b) Against the set  $(E \cap P)$  (for each of which all of the three techniques can generate the same English translation candidate, 174 Japanese noun phrases)

Eijiro	Compositional Tran. Generation	Phrase Tran. Table	Agreements of the Three Tech- niques
97.7 (170/174)	97.1 (169/174)	96.0 (167/174)	96.0 (167/174)
97.7 (170/174)	97.1 (169/174)	96.0 (167/174)	<b>98.8</b> (167/169)
97.7	97.1	96.0	<b>97.4</b>

(c) Against the set  $(C \cap P) - E$  (for each of which both compositional translation generation and phrase translation table can generate the same English translation candidate, but the Eijiro can not, 272 Japanese noun phrases)

Compositional Tran. Generation	Phrase Tran. Table	Agreements of Compositional Tran. Generation and Phrase Tran. Table	Validation by SVMs
91.9 (250/272)	90.8 (247/272)	89.3 (243/272)	93.0 (253/272)
91.9 (250/272)	90.8 (247/272)	<b>93.1</b> (243/261)	93.0 (253/272)
91.9	90.8	<b>91.2</b>	<b>93.0</b>

(d) Against the set  $P - (C \cap P)$  (for which only the phrase translation table can generate English translation candidates, 504 Japanese noun phrases)

Phrase Tran. Table	Validation by SVMs
81.5 (411/504)	$\left( \begin{array}{c} 57.5(290/504) \\ \underline{\mathbf{90.1}}(290/322) \\ 70.2 \end{array} \right)$
81.5 (411/504)	$\left( \begin{array}{c} 72.8(366/504) \\ \underline{\mathbf{87.1}}(366/420) \\ 79.2 \end{array} \right)$
81.5	

primary application of semi-automatic acquisition of technical term bilingual lexicon, we prefer precision to recall, and regard this precision (around 87%) of the phrase translation table against the whole set (a) as a baseline of the evaluation of this paper.

Compared with this baseline, against the set (b) (i.e., Japanese noun phrases for which all the three techniques can generate English translation candidates) and (c) (i.e., Japanese noun phrases for which both compositional translation generation and the phrase translation table can generate English translation candidates, but the Eijiro can not), agreements of the three or the two techniques have precisions and F-measures over 90%. For both sets (b) and (c), agreements of the three or the two techniques essentially represent agreement of the two resources that have quite different nature, i.e. a bilingual lexicon for human use and a statistical technique. Because of this difference in nature of the resource,

we can achieve high precision in their agreement. Union of those sets (b) and (c) cover 43% of the whole set of 1,040 Japanese noun phrases, and we can have around 95% precision for the union in total. We can claim that such a high precision is definitely an advantage in terms of our application of semi-automatic acquisition of technical term bilingual lexicon.

## 5 Validating Translation Candidates by SVMs

### 5.1 The Procedure

This section describes the procedure and the results of applying Support Vector Machines (SVMs) (Vapnik, 1998) to the task of validating translation candidates generated by the three techniques.

As a tool for learning SVMs, we use TinySVM (<http://chasen.org/~taku/software/>

Table 4: Features of SVMs Learning

Feature Type	Features
Monolingual (for the set (d))	number of morphemes in the Japanese noun phrase
	number of words in the English translation candidate
Bilingual — based on Eijiro	score and rank of compositional translation generation given to the English translation candidate (for the set (c))
	whether at least one translation pair of constituents of the Japanese noun phrase and the English translation candidate is included in Eijiro (for the set (d))
Bilingual — based on statistics in the parallel sentences	probability and rank of the phrase translation table given to the English translation candidate
	frequencies $freq(t_E, t_J)$ , $freq(t_E, \neg t_J)$ , and $freq(\neg t_E, t_J)$ in the contingency table

TinySVM/). Each training/test instance of SVMs learning is represented as a tuple  $\langle t_J, t_E, c \rangle$ , where  $t_J$  and  $t_E$  denote a Japanese noun phrase and an English translation candidate generated by at least one of the three techniques, and the class  $c$  denotes whether  $t_E$  is a correct translation of  $t_J$  found in the English part of the parallel sentence (i.e., “ $c = +$ ”), or not (i.e., “ $c = -$ ”). Out of the whole 1,040 Japanese noun phrases, at least one English translation candidate is generated for 954 of them and the total number of generated English translation candidates is 2,851. Thus, we have 2,851 instances for training/testing SVMs in total. As the kernel function, we compare the linear and the polynomial (2nd order) kernels, where the latter performs better.

In the testing of a SVMs classifier, given a Japanese noun phrase  $x_J$ , we collect all the tuples  $\langle x_J, t_E, c \rangle$  which have  $x_J$  in the Japanese part, and classify each tuple by the SVMs classifier. Here, we regard the distance from the separating hyperplane to each test instance as a confidence measure, and choose a tuple which satisfies the following: i.e., one for which the classifier outputs the class as “+”, and furthermore, one with the greatest distance from the separating hyperplane. In the actual evaluation of the “Validation by SVMs” column in Table 3, we train/test an SVMs classifier separately for each of the sets (c) of 272 Japanese noun phrases and (d) of 504 Japanese noun phrases. For both sets, the “Validation by SVMs” column in Table 3 shows the evaluation results by 10-fold cross-validation.

## 5.2 Features

Table 4 lists the features used in the SVMs learning. As monolingual features, we use the number of morphemes constituting the Japanese noun phrase as well as that of words constituting the English translation candidate. We evaluated these features for both of the sets (c) and (d) in Table 3, where for the set (c), we had better performance without these features. Thus, we use these features only for the set (d).

Bilingual features can be classified into two types: one is based on translation knowledge in the bilingual lexicon Eijiro for human use, while the other is based on statistics obtained from the parallel sentences used for learning the phrase translation table. As bilingual features based on Eijiro, first we use the score and the rank of compositional translation generation given to the English translation candidate, which are used only for the set (c). Second, for the set (d), although compositional translation generation can not generate any translation candidate, we lookup the bilingual lexicon Eijiro and examine whether any translation pair for a constituent of the Japanese noun phrase and that of the English translation candidate can be found. Then, we use whether at least one translation pair is included in Eijiro as a bilingual feature. For example, in the case of a Japanese technical term “応用行動分析” and its English translation candidate “*application behavior analysis*”, the value of this feature is true if a translation pair such as “分析” and “*analysis*” is included



in Eijiro.

As bilingual features based on statistics obtained from the parallel sentences, first we use the probability and the rank of the phrase translation table given to the English translation candidate. Second, as another type of bilingual features based on statistics from the parallel sentences, we use statistics previously used for measuring statistical co-occurrence of translation pairs such as the mutual information, the  $\phi^2$  statistic, the dice coefficient, and the log-likelihood ratio (Matsumoto and Utsuro, 2000). Given an English term  $t_E$  and a Japanese term  $t_J$ , as bilingual features, we use co-occurrence frequencies of  $t_E$  and  $t_J$  in the contingency table below:

	$t_J$	$\neg t_J$
$t_E$	$freq(t_E, t_J)$	$freq(t_E, \neg t_J)$
$\neg t_E$	$freq(\neg t_E, t_J)$	$freq(\neg t_E, \neg t_J)$

We also evaluated the  $\phi^2$  statistic as a feature, but we had better performance without the feature, and thus we use those co-occurrence frequencies directly as features.

### 5.3 Evaluation Results

First, for the set (c), we regard the result of the agreement between compositional translation generation and the phrase translation table as a baseline. Here, in the column “Validation by SVMs” of Table 3, we have the F-measure as 93.0, which is slightly higher than the baseline 91.2, although their difference is not statistically significant.

Next, for the set (d), we regard the precision and the F-measure of the phrase translation table as a baseline. For the set (d), only the phrase translation table can generate English translation candidates, where the precision and the F-measure are 81.5% which is lower than those for other sets (b) and (c). In this case, it is required for the SVMs classifier to validate the English translation candidates in the phrase translation table and to reject incorrect candidates. In order to realize this, we introduce a lower bound against the distance from the separating hyperplane to each test instance, where English translation candidates with this distance smaller than the lower bound are rejected. By examining various values of this lower bound with other held out data, we can achieve the highest precision 90.1%, or slightly less precision 87.1% with higher F-measure. Differ-

ences between those precisions and the baseline are statistically significant at a level of 0.05. With these improvement in precisions, again we can claim that the approach of applying the SVMs learning technique to the task of validating translation candidates definitely contributes to semi-automatic acquisition of technical term bilingual lexicon.

## 6 Related Works

Among the techniques studied so far in the research area of automatic bilingual lexicon compilation as well as empirical approaches to machine translation such as statistical machine translation models, (Itagaki et al., 2007) is most closely related to the approach taken in this paper. (Itagaki et al., 2007) focused on automatic validation of translation pairs available in the phrase translation table learned by a statistical machine translation model. One of the major differences between (Itagaki et al., 2007) and the approach taken in this paper is that we focus on integrating the phrase translation table with compositional translation generation based on an existing bilingual lexicon for human use (Tonoike et al., 2006). As we showed in the experimental evaluation, translation knowledge resource of an existing bilingual lexicon for human use definitely contributes to improving the precision of translation candidates both in the agreement of two or three techniques and in validation by SVMs learning.

The system combination approaches to machine translation (Rosti et al., 2007; Matusov et al., 2006) are another related research in a broader perspective. One of the major differences between such system combination approaches to the whole sentence MT and the task focused in this paper is apparently in that we concentrate on application of semi-automatic acquisition of technical term bilingual lexicon, where the primary requirement is precision rather than recall of the acquired translation pairs.

## 7 Conclusion

This paper presented an attempt at developing a technique of acquiring translation pairs of technical terms with sufficiently high precision from parallel patent documents. The approach taken in the proposed technique is based on integrating the

phrase translation table of a state-of-the-art statistical phrase-based machine translation model, and compositional translation generation based on an existing bilingual lexicon for human use. Our evaluation results clearly showed that the agreement between the two individual techniques definitely contribute to improving precision of translation candidates. We then applied the Support Vector Machines (SVMs) to the task of automatically validating translation candidates in the phrase translation table. Experimental evaluation results again showed that the SVMs based approach to translation candidates validation can contribute to improving the precision of translation candidates in the phrase translation table.

## References

- A. Fujii, M. Utiyama, M. Yamamoto, and T. Utsuro. 2008. Toward the evaluation of machine translation using patent information. In *Proc. 8th AMTA*.
- P. Fung and L. Y. Yee. 1998. An IR approach for translating new words from nonparallel, comparable texts. In *Proc. 17th COLING and 36th ACL*, pages 414–420.
- F. Huang, Y. Zhang, and S. Vogel. 2005. Mining key phrase translations from Web corpora. In *Proc. HLT/EMNLP*, pages 483–490.
- M. Itagaki, T. Aikawa, and X. He. 2007. Automatic validation of terminology translation consistency with statistical method. In *Proc. MT summit XI*, pages 269–274.
- P. Koehn, F. J. Och, and D. Marcu. 2003. Statistical phrase-based translation. In *Proc. HLT-NAACL*, pages 127–133.
- P. Koehn, H. Hoang, A. Birch, C. Callison-Burch, M. Federico, N. Bertoldi, B. Cowan, W. Shen, C. Moran, R. Zens, C. Dyer, O. Bojar, A. Constantin, and E. Herbst. 2007. Moses: Open source toolkit for statistical machine translation. In *Proc. 45th ACL, Companion Volume*, pages 177–180.
- Y. Matsumoto and T. Utsuro. 2000. Lexical knowledge acquisition. In R. Dale, H. Moisl, and H. Somers, editors, *Handbook of Natural Language Processing*, chapter 24, pages 563–610. Marcel Dekker Inc.
- E. Matusov, N. Ueffing, and H. Ney. 2006. Computing consensus translation for multiple machine translation systems using enhanced hypothesis alignment. In *Proc. 11th EACL*, pages 33–40.
- F. J. Och and H. Ney. 2003. A systematic comparison of various statistical alignment models. *Computational Linguistics*, 29(1):19–51.
- A.-V. Rosti, S. Matsoukas, and R. Schwartz. 2007. Improved word-level system combination for machine translation. In *Proc. 45th ACL*, pages 312–319.
- M. Tonoike, M. Kida, T. Takagi, Y. Sasaki, T. Utsuro, and S. Sato. 2006. A comparative study on compositional translation estimation using a domain/topic-specific corpus collected from the web. In *Proc. 2nd Intl. Workshop on Web as Corpus*, pages 11–18.
- M. Utiyama and H. Isahara. 2007. A Japanese-English patent parallel corpus. In *Proc. MT summit XI*, pages 475–482.
- V. N. Vapnik. 1998. *Statistical Learning Theory*. Wiley-Interscience.