

A Bayesian Model of Transliteration and Its Human Evaluation When Integrated into a Machine Translation System

Andrew FINCH^{†a)}, *Nonmember*, Keiji YASUDA^{†b)}, *Member*, Hideo OKUMA^{†c)}, *Nonmember*, Eiichiro SUMITA^{†d)}, and Satoshi NAKAMURA^{†,††e)}, *Members*

SUMMARY The contribution of this paper is two-fold. Firstly, we conduct a large-scale real-world evaluation of the effectiveness of integrating an automatic transliteration system with a machine translation system. A human evaluation is usually preferable to an automatic evaluation, and in the case of this evaluation especially so, since the common machine translation evaluation methods are affected by the length of the translations they are evaluating, often being biased towards translations in terms of their length rather than the information they convey. We evaluate our transliteration system on data collected in field experiments conducted all over Japan. Our results conclusively show that using a transliteration system can improve machine translation quality when translating unknown words. Our second contribution is to propose a novel Bayesian model for unsupervised bilingual character sequence segmentation of corpora for transliteration. The system is based on a Dirichlet process model trained using Bayesian inference through blocked Gibbs sampling implemented using an efficient forward filtering/backward sampling dynamic programming algorithm. The Bayesian approach is able to overcome the overfitting problem inherent in maximum likelihood training. We demonstrate the effectiveness of our Bayesian segmentation by using it to build a translation model for a phrase-based statistical machine translation (SMT) system trained to perform transliteration by monotonic transduction from character sequence to character sequence. The Bayesian segmentation was used to construct a phrase-table and we compared the quality of this phrase-table to one generated in the usual manner by the state-of-the-art GIZA++ word alignment process used in combination with phrase extraction heuristics from the MOSES statistical machine translation system, by using both to perform transliteration generation within an identical framework. In our experiments on English-Japanese data from the NEWS2010 transliteration generation shared task, we used our technique to bilingually co-segment the training corpus. We then derived a phrase-table from the segmentation from the sample at the final iteration of the training procedure, and the resulting phrase-table was used to directly substitute for the phrase-table extracted by using GIZA++/MOSES. The phrase-table resulting from our Bayesian segmentation model was approximately 30% smaller than that produced by the SMT system's training procedure, and gave an increase in transliteration quality measured in terms of both word accuracy and F-score.

Key words: transliteration, human evaluation, machine translation, Dirichlet process model, Bayesian, segmentation

1. Introduction

Machine translation has to handle wide variety of proper

nouns to achieve high translation performance. However, it is impossible to discover all of these words from limited quantities of bilingual training data, or even to manually maintain a bilingual dictionary because of the huge number of distinct proper nouns in the real world including new words and expression being created day by day. Therefore, machine translation is a the promising application for the application of transliteration technology. In this paper we present a Bayesian technique for building transliteration models that avoids the overfitting issues commonly encountered with non-Bayesian techniques. We evaluate this approach automatically with respect to prior work, but more importantly we evaluate the worth of our approach in a real-world translation system.

Automatic machine translation metrics, for example the BLEU score [1], and the NIST score [2] often have a natural bias towards shorter translations. That is to say, it is often a better strategy for a machine translation system to output nothing or an very short erroneous translation than to output a long, incorrect translation. Within these metrics attempts to correct this this length bias are included, but unfortunately they are not completely effective and some degree of bias remains. As a evidence of this phenomenon, it is common practice in competitive machine translation evaluation campaigns for the systems to delete untranslated unknown words from their machine translation output, rather than keep them or attempt to transliterate them (for example, [3]). For this reason it is important to study the effect of introducing transliteration into a machine translation system through a human evaluation experiment, even though such experiments are expensive in terms of human effort. We now move on to motivate the Bayesian co-segmentation scheme using in training our transliteration model.

1.1 Motivation

It is possible to couch the problem of transliteration as a problem of machine translation at the character level. In this paradigm, decoding is usually assumed to proceed in a monotone order, but otherwise the technique remains essentially the same, except that the tokens used in the system are characters rather than words. Recently systems based on phrase-based statistical machine translation technology are being actively researched [4]–[6], and have achieved state-of-the-art performance on this task. The approach makes no linguistic assumptions about the data and no intermediate

Manuscript received December 28, 2010.

Manuscript revised May 5, 2011.

[†]The authors are with NICT, Kyoto-fu, 619–0289 Japan.

^{††}The author is with NIST, Ikoma-shi, 630–0192 Japan.

a) E-mail: andrew.finch@nict.go.jp

b) E-mail: keiji.yasuda@nict.go.jp

c) E-mail: hideo.okuma@nict.go.jp

d) E-mail: eiichiro.sumita@nict.go.jp

e) E-mail: s-nakamura@is.naist.jp

DOI: 10.1587/transinf.E94.D.1889

phonetic representation is required, because the transduction is directly from grapheme to grapheme. The advantages of this type of approach are that the only training corpus required is a set of bilingual word pairs, and the approach can be applied directly to a wide range of language pairs without the need to develop a set of linguistically-motivated heuristics specific to the languages involved.

At the core of all phrase-based statistical machine translation systems (SMT) is the phrase-table. This table is the basic set of building blocks that are used to construct the translation. The creation of a phrase-table during a typical training procedure for a phrase-based SMT system consists of the following steps:

1. Word alignment using GIZA++ [7]
2. Phrase-pair extraction using heuristics (for example *grow-diaf-final-and* from the MOSES [8] toolkit)

This approach works very well in practice, but a more elegant solution would be to arrive at a set of bilingual sequence-pairs (we use this term to describe analogue of the phrase-pair at the character level) in one step, from a generative model. Unfortunately, when traditional methods that use the EM algorithm to maximize likelihood are applied to the task, they produce solutions that can grossly over-fit the data. As an extreme example, the most likely segmentation of a corpus into sequence-pairs, assuming no limits on sequence-pair size would be the entire corpus as a single bilingual sequence-pair, holding all the probability mass.

GIZA++ mitigates this problem by aligning the words in a one-to-many fashion. The single word on one side of the alignment acts as a constraint on the size of the bilingual pairs. A similar approach can be taken in transliteration, where a single character in one language is permitted to align to multiple characters of the other, but not vice versa. This approach is reasonable for English-Chinese transliteration [9], [10], where one Chinese character can be assumed to map to several English characters.

In GIZA++ this one-to-many alignment is done twice: from both source-to-target and also from target-to-source. A table of word-to-word alignments is then constructed from (typically the intersection) both of these alignments. Additional word alignments that are not in the intersection are added based on evidence and heuristics, and finally all possible phrase-pairs are extracted from the table of alignments that are consistent with the table.

In [11], [12] many-to-many alignment is performed directly using maximum likelihood training, but evidence trimming heuristics that exclude part of the available training data are required to prevent the models from overfitting the data. [13] have successfully applied a similar Bayesian technique to grammar induction and [14], [15] have developed a tractable Bayesian methods for the more complex task of bilingual phrase pair extraction for SMT, which involves reordering. [16] tackle the overfitting problem in phrasal alignment by using a leave-one-out approach using a strategy that despite being a different paradigm, shares many of the characteristics of our approach.

In this paper we extend existing monolingual word segmentation models ([17], [18]) to bilingual segmentation, and provide a simple yet elegant way to directly segment a bilingual training corpus in a many-to-many fashion without overfitting, using a Bayesian model.

This paper is organized as follows. In Sect. 2 we describe the Bayesian model used in our transliteration system. Here we give an overview of the Dirichlet process model, the Chinese Restaurant process and explain how our model relates to these two representations. We also describe the blocked Gibbs sampling technique used to train the model. In Sect. 4 we describe the experiments we performed to evaluate our model: the data sets, the baseline system and the training procedure. Section 5 contains the experimental results, and in Sect. 7 we conclude and mention promising avenues for future research.

2. Methodology

Recently in the natural language processing field Bayesian models have been proposed to tackle a variety of problems, and have been found to be particularly effective in word segmentation [17], [18]. The model we use in this paper is a unigram Dirichlet process model. Using this approach to perform bilingual segmentation for the general case of machine translation with re-ordering would be a challenging undertaking, however for transliteration where the sequence lengths are short and under the assumption that there is no re-ordering, it is feasible to tackle the bilingual segmentation problem directly without the need for specialized optimization or annealing (we do use a block sampling algorithm, and a dynamic programming algorithm).

2.1 Joint Source-Channel Model

Let us assume we are given a bilingual corpus consisting of a source sequence $\mathbf{s}_1^M = \langle s_1, s_2, \dots, s_M \rangle$ and a target sequence $\mathbf{t}_1^N = \langle t_1, t_2, \dots, t_N \rangle$. We distinguish sequences of characters from single characters by using a boldface font with an overbar.

We adopt the joint source-channel model of [9] as the underlying generative model, and we make the additional assumption that the segments are independent of each other (our approach can easily be extended to model these dependencies at the expense of some additional complexity, see [18]). Under this model, the corpus is generated through the concatenation of *bilingual sequence-pairs* (we will use this term repeated throughout this paper to refer to corresponding sequences of source and target graphemes, as defined below).

A bilingual sequence-pair is a tuple (\mathbf{s}, \mathbf{t}) consisting of a sequence of source graphemes together with a sequence of target graphemes $(\mathbf{s}, \mathbf{t}) = (\langle s_1, s_2, \dots, s_i \rangle, \langle t_1, t_2, \dots, t_j \rangle)$.

The corpus probability is simply the probability of all possible derivations of the corpus given the set of bilingual sequence-pairs and their probabilities.

$$\begin{aligned}
 p(\mathbf{s}_1^M, \mathbf{t}_1^N) &= p(s_1, s_2, \dots, s_M, t_1, t_2, \dots, t_N) \\
 &= \sum_{\gamma \in \Gamma} p(\gamma)
 \end{aligned}
 \tag{1}$$

where $\gamma = ((s_1, t_1), \dots, (s_k, t_k), \dots, (s_K, t_K))$ is a derivation of the corpus characterized by its co-segmentation, and Γ is the set of all derivations (co-segmentations) of the corpus.

The probability of a single derivation is given by the product of its component bilingual sequence-pairs:

$$p(\gamma) = \prod_{k=1}^K p((s_k, t_k))
 \tag{2}$$

The corpus for our experiments is segmented into bilingual word-pairs. We therefore constrain our model such that both source and target character sequences of each bilingual sequence-pair in the derivation of the corpus are not allowed to cross a word segmentation boundary. Equation (2) can therefore be arranged as a product of word-pair w derivations of the sequence of all word-pairs \mathcal{W} in the corpus.

$$p(\gamma) = \prod_{w \in \mathcal{W}} \prod_{(s_k, t_k) \in \gamma_w} p((s_k, t_k))
 \tag{3}$$

where γ_w is a derivation of bilingual word-pair w .

2.2 Unigram Dirichlet Process Model

A Dirichlet process is a stochastic process defined over a set S (in our case, the set of all possible bilingual sequence-pairs) whose sample path is a probability distribution on S .

The Dirichlet process model we use in our approach is a simple model that resembles the cache models used in language modeling [19]. Intuitively, the model has two basic components: a model for generating an outcome that has already been generated at least once before, and a second model that assigns a probability to an outcome that has not yet been produced. Ideally, to encourage the re-use of model parameters, the probability of generating a novel bilingual sequence-pair should be considerably lower than the probability of generating a previously observed sequence pair. This is a characteristic of the Dirichlet process model we use and furthermore, the model has a preference to generate new sequence-pairs early on in the process, but is much less likely to do so later on. In this way, as the cache becomes more and more reliable and complete, so the model prefers to use it rather than generate novel sequence-pairs. The probability distribution over these bilingual sequence-pairs (including an infinite number of unseen pairs) can be learned directly from unlabeled data by Bayesian inference of the hidden co-segmentation of the corpus. The ability of the model to assign a probability to any unseen sequence-pair gives the technique the ability to score candidate training data.

The underlying stochastic process for the generation of a corpus composed of bilingual phrase pairs γ is usually written in the following from:

$$\begin{aligned}
 G_{|\alpha, G_0} &\sim DP(\alpha, G_0) \\
 (s_k, t_k) | G &\sim G
 \end{aligned}
 \tag{4}$$

G is a discrete probability distribution over the all bilingual sequence-pairs according to a *Dirichlet process prior* with *base measure* G_0 and concentration parameter α . The concentration parameter $\alpha > 0$ controls the variance of G ; intuitively, the larger α is, the more similar G_0 will be to G .

2.2.1 The Chinese Restaurant Process

Unfortunately it is not possible to estimate G directly, since there are an infinite number of possible bilingual sequence-pairs, so instead we integrate over its possible values. To do this we cast the bilingual sequence-pair generation process as an instance of the Chinese Restaurant Process (CRP) [20]. According to this representation, every bilingual sequence-pair corresponds to the dish served at its table in a potentially infinite set of tables in a Chinese restaurant. The number of customers seated at each table represents the cumulative count of the bilingual sequence-pair. A new customer to the restaurant can take a seat at an occupied table with a probability proportional to the number of customers at that table, and must eat that table’s dish, or can take a seat at an unoccupied table with a probability proportional to a constant, in which case they must eat a dish (a bilingual sequence-pair) chosen by the chef (in this analogy the chef’s choice is in accordance with the base distribution G_0).

2.2.2 The Base Measure

For the *base measure* that controls the generation of novel sequence-pairs, we use a joint spelling model that assigns probability to new sequence-pairs according to the following joint distribution:

$$\begin{aligned}
 G_0((s, t)) &= p(|s|)p(s||s) \times p(|t|)p(t||t) \\
 &\simeq \frac{\lambda_s^{|s|}}{|s|!} e^{-\lambda_s} v_s^{-|s|} \times \frac{\lambda_t^{|t|}}{|t|!} e^{-\lambda_t} v_t^{-|t|}
 \end{aligned}
 \tag{5}$$

where $|s|$ and $|t|$ are the length in characters of the source and target sides of the bilingual sequence-pair; v_s and v_t are that vocabulary (alphabet) sizes of the source and target languages respectively; and λ_s and λ_t are the expected lengths of source and target.

According to this model, source and target sequences are generated independently: in each case the sequence length is chosen from a Poisson distribution, and then the sequence itself is generated given the length. Note that this model is able to assign a probability to arbitrary bilingual sequence-pairs of any length in source and target sequence, but favors shorter sequences in both.

More sophisticated methods of defining the base measure are possible, for example [14], [15] use the IBM model 1 likelihood of one phrase conditioned on the other in the base model to encourage the formation of bilingual pairs that follow the word alignments in the corpus. This idea

Input: Random initial corpus segmentation
Output: Unsupervised co-segmentation of the corpus according to the model

```

foreach iter=1 to NumIterations do
  foreach bilingual word-pair  $w \in \text{randperm}(W)$  do
    foreach co-segmentation  $\gamma_i$  of  $w$  do
      Compute probability  $p(\gamma_i|h)$ 
      where  $h$  is the set of data (excluding  $w$ ) and its hidden co-segmentation
    end
    Sample a co-segmentation  $\gamma_i$  from the distribution  $p(\gamma_i|h)$ 
    Update counts
  end
end

```

Algorithm 1: The blocked Gibbs sampling algorithm.

can be transferred to the character level and used in our approach, but remains as future work at this point.

Following [17] we assign the parameters λ_s , λ_t and α , the values 2, 2 and 0.3 respectively. Ideally these parameters should be learned from the data, however in our experiments the settings were sufficient to give a useful co-segmentation of the training corpus. Moreover, the system proved to be insensitive to changes in these parameters in a set of pilot experiments, converging to very similar final iteration samples for a range of parameter settings.

2.2.3 The Generative Model

The generative model is given in Eq. (6) below. The equation assigns a probability to the k^{th} bilingual sequence-pair $(\mathbf{s}_k, \mathbf{t}_k)$ in a derivation of the corpus, given all of the other sequence-pairs in the history so far $(\mathbf{s}_{-k}, \mathbf{t}_{-k})$. Here $-k$ is read as: “up to but not including k ”.

$$p((\mathbf{s}_k, \mathbf{t}_k) | (\mathbf{s}_{-k}, \mathbf{t}_{-k})) = \frac{N((\mathbf{s}_k, \mathbf{t}_k)) + \alpha G_0((\mathbf{s}_k, \mathbf{t}_k))}{N + \alpha} \quad (6)$$

In this equation, N is the total number of bilingual sequence-pairs generated so far (the number of customers so far), $N((\mathbf{s}_k, \mathbf{t}_k))$ is the number of times the sequence-pair $(\mathbf{s}_k, \mathbf{t}_k)$ has occurred in the history (the number of people seated at its table). G_0 and α are the base measure and concentration parameter as before.

3. Bayesian Inference

3.1 Gibbs Sampling

We used a blocked version of a Gibbs sampler for training. In [19] they report issues with mixing in the sampler that were overcome using annealing. In [18] this issue was overcome by using a blocked sampler together with a dynamic programming approach. Our algorithm is similar to that of [18], and we found our sampler converged rapidly without annealing (see Fig. 4). The number of iterations was set by hand after observing the convergence behavior of the algorithm in pilot experiments. We used a value of 30 iterations through the corpus in all our experiments.

The sampling algorithm is shown in Algorithm 1 and the iterative component proceeds as follows. Firstly the training set of bilingual word-pairs is permuted randomly, and a bilingual word-pair is sampled from this permutation without replacement. Secondly, a probability distribution over all possible co-segmentations of the chosen bilingual word-pair is calculated by obtaining probabilities with respect to a model that does not include the bilingual word-pair, its previous segmentation information and respective counts. Due to the short sequence lengths involved in transliteration, it is possible to use a brute force approach to calculate this distribution, however for efficiency we extended the forward filtering/backward sampling (FFBS) dynamic programming algorithm of [18] to deal with bilingual segmentation. We implemented this algorithm graphically as explained below.

We use a segmentation graph (shown in Fig. 1) to guide the process. This directed graph is a compact representation of all possible ways in which to co-segment a bi-lingual pair. Each node represents a set of partial co-segmentation hypotheses of the whole sequence that share the same sequences of source and target tokens, and each arc represents the bilingual phrase pair used to transition from the tail of the arc to the head. In the figure the arcs are labelled with the log-probability of this sequence-pair (given by the model in Eq. (6)), therefore the log-probability of a full segmentation hypothesis is given by the sum of the arc labels on the respective path from the source node ‘<s>’ to the sink node ‘abba’. The most probable co-segmentation is indicated with bold arcs in the figure and corresponds to the segmentation ‘a b ba’, this is reasonable since both ‘a/A’ and ‘ba/BA’ are associated with their phonetic equivalents in Japanese, and the Japanese ‘TSU’ indicates that the consonant immediately to the right is to be repeated. The least probable segmentation in the graph is given by ‘abb/A a/TSU-BA’. The log-probabilities in the graph are real values taken from the third iteration of the training, and here the most probable segmentation is already by far the most likely.

Nodes in the graph can have multiple in- and out-degree. Two nodes are combined when the unsegmented part of the bi-lingual sequence pair is the same for both, giving rise to a compact, efficient representation.

The FFBS algorithm operates directly on the segmen-

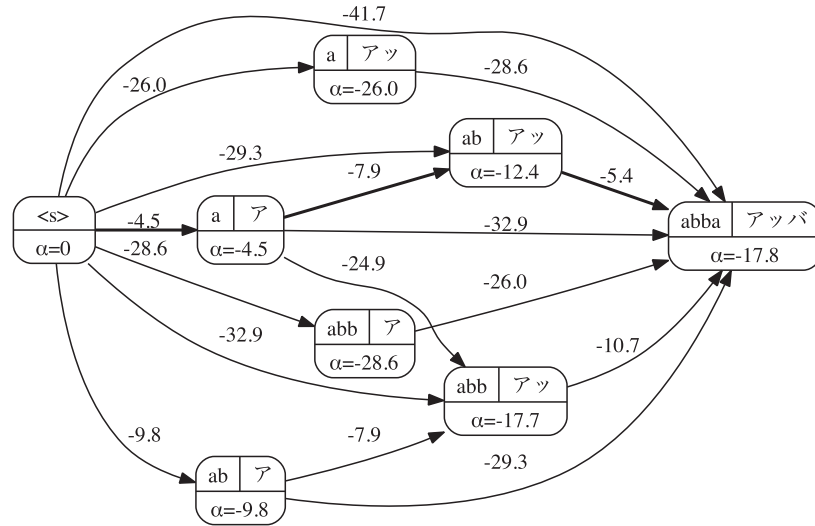


Fig. 1 A graph representing all possible co-segmentations of the character sequences “abba” in English and “A-TSU-BA” in Japanese. The α labels on the nodes represent the log-probability of subgraph (including the node itself) to the left of the node. The labels on the arcs are the log-probabilities of bi-lingual phrase pairs used to transition from tail-to-head, and are given by the model of Eq. (6).

```

INPUT:      a sequence of bilingual sequence pairs: ARRAY-of-sequence-pairs derivation
OUTPUT:   a set of all pairs formed by agglomerating the pairs in derivation

SET-of-sequence-pairs    agglomerations    = empty
ARRAY-of-sequence-pairs  chunk              = empty

FOR chunk in all subsequences of derivation
{
    sequence-pair pair = concatenation of sequence pairs in chunk

    IF (pair has source sequence length <= MAX_SOURCE_SEQUENCE_LENGTH AND
        pair has target sequence length <= MAX_TARGET_SEQUENCE_LENGTH)
    {
        INSERT pair into agglomerations
    }
}

PRINT all_agglomerations
    
```

Fig. 2 The sequence-pair agglomeration algorithm.

tation graph, and has two steps. The *forward filtering* step, calculates for each node in the graph, the probability of the subgraph (including the node itself) to the left of the node, back to the source node. This probability α , is stored in the node itself (these α 's are shown in Fig. 1). This process proceeds recursively in a depth-first post-order traversal of the graph, starting at the sink node. Nodes for which the probability has been calculated are marked as done, ensuring α gets calculated only once for the node.

The *backward sampling* step samples a derivation of the bi-lingual word pair according to the probability distribution over all possible segmentations. This is done easily using the α values stored in the graph by the forward filtering process. The backward sampling also proceeds recursively from the sink node. For each incoming arc, the probability of including that arc in the sample is given by the

product of the arc probability and the α value at the tail of the arc. This value is calculated for each incoming arc, and one arc from the set is sampled according to the probability distribution over the arcs. The sampling procedure is called recursively on the tail of the sampled arc until the source node of the graph is reached. The sequence of arcs traversed defines the sampled derivation of the bi-lingual pair for the current iteration of the training process, and this sample is in accordance with the probability distribution over all derivations with respect to the model.

3.2 Sequence-Pair Extraction

During the phrase-table generation process of a typical phrase-based SMT system, GIZA++ is run twice to generate alignments at the word level, from source-to-target and from

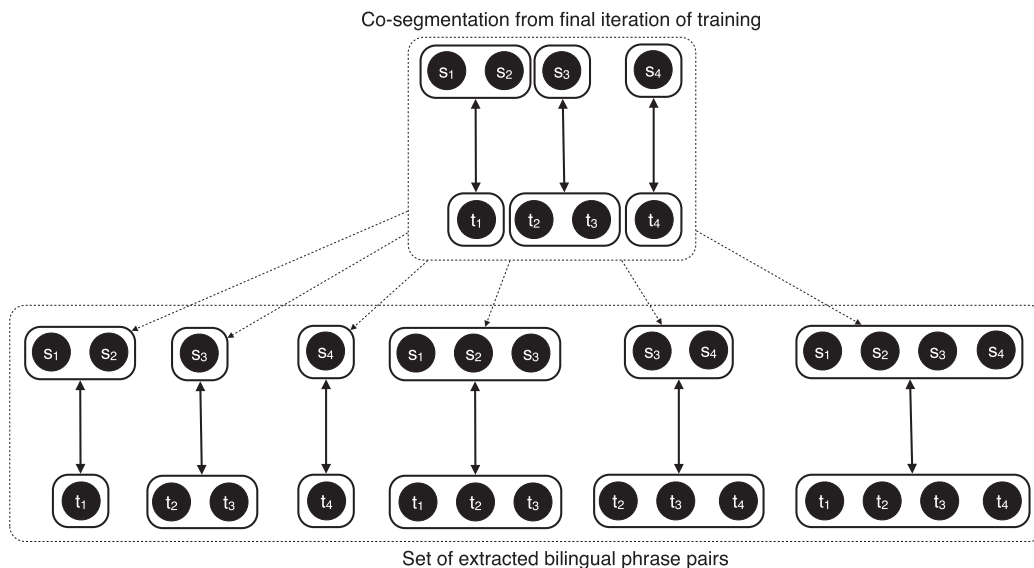


Fig. 3 The sequence-pair extraction process for a single bilingual word pair, using sequence-pair agglomeration.

target-to-source. Following this step, the *grow-diag-final-and* procedure is used to extract *all phrases* consistent with the word alignments arising from the two GIZA++ runs. When building a phrase-table from the alignment achieved at final iteration of our Gibbs sampling procedure, we use a much simpler heuristic that is in the same spirit to derive a larger set of phrases consistent with the initial co-segmentation. Our experiments show that this is a necessary step that considerably improves system performance.

The algorithm we use for phrasal extraction from the co-segmented corpus is as follows: within a single bilingual word-pair, agglomerate all contiguous bilingual sequence-pairs in all possible ways, but limit the size of the agglomerated source and target phrases to match the *maximum phrase length* parameter used to train the SMT system (this was set to 7 in our experiments). This is not strictly necessary, but we performed this step to keep the phrase-table generated from our Bayesian segmentation comparable to that generated by the baseline system. The algorithm used to perform the sequence-pair agglomeration is outlined in Fig. 2. An example word pair together with its segmentation and the set of agglomerated phrases that can be extracted from it is illustrated in Fig. 3.

4. Transliteration Experiments

4.1 Baseline System

For our experiments we use the phrase-based machine translation techniques introduced by [8], integrating our models within a log-linear framework [21]. Word alignment was performed using GIZA++ [7] and sequence-pair extraction using the MOSES [8] tools. The decoder used was an in-house phrase-based machine translation decoder that operates according to the same principles as the publicly avail-

Table 1 Statistics of the English-Japanese bilingual corpora.

Corpus	word-pairs	Characters		Avg. Word Len.	
		En	Ja	En	Ja
Training	27993	188941	131275	6.75	4.69
Development	3606	24066	16651	6.67	4.62
Evaluation	1935	11863	8199	6.13	4.24

able MOSES [8] SMT decoder.

In these experiments 5-gram language models built with Witten-Bell smoothing were used. The system was trained in a standard manner, using a minimum error-rate training (MERT) procedure [22] with respect to the BLEU score on the held-out development data to optimize the log-linear model weights.

Rama and Gali [23], evaluated several techniques for sequence-pair extraction for transliteration and found the *grow-diag-final-and* heuristic to be the most effective, we therefore adopt this method in the baseline system our experiments.

4.1.1 Decoding Constraints

The experiments reported in this paper were conducted using a beam width of 100, with no stack thresholding, and a strictly monotone decoding process.

4.2 Experimental Data

Our training data consisted of 27993 bilingual single word-pairs that were used in the NEWS2010 workshop transliteration shared task. The development data consisted of 3606 bilingual word-pairs drawn from the same sample. The evaluation data consisted of a further 1935 bilingual word-pairs not contained in the other two data sets. The corpus statistics for the three corpora are given in Table 1.

We used the data to train a phrase-based SMT system to perform transliteration from English to Japanese. We trained our Dirichlet process model on the same parallel data set, and extracted transliteration phrase-tables from the co-segmentation of the corpus at the final iteration (iteration 30).

4.3 Training Procedure

For the Gibbs sampling, we chose to start the sampling from a random co-segmentation of the corpus. That is, for each bilingual word-pair in the corpus, a single co-segmentation was sampled from a uniform distribution over all possible co-segmentations of the pair. We believe that it might be advantageous, and certainly more efficient to start the sampling from a more intelligent starting point, for example one derived from a pre-processing pass of GIZA++. However, the training was able to arrive at a good segmentation (by visual inspection) of the training corpus, its usefulness being borne out by the experimental results in the next section.

4.4 Evaluation Procedure

The results presented in this paper are given in terms of official evaluation metrics used in the NEWS2010 transliteration generation shared task [24]. In our results, ACC refers to the top-1 accuracy score, that measures the percentage of the time the top hypothesis from the system exactly matches the reference. F-score measures the distance of the best hypothesis from the reference transliteration; the reader is referred to the workshop white-paper [24] for more details. For brevity, we only report our results in terms of ACC and F-score in this paper, but the results in terms of the other NEWS2010 metrics have the same character.

5. Results

5.1 Training

The convergence of the algorithm during the training procedure is shown in Fig. 4 which plots the log-probability of the sampled derivation at the end of each pass through the training corpus against iteration. It can be seen from the graph that the system rapidly improves from the poor initial segmentation, and thereafter continues to gradually improve. The log-probability of the initial random co-segmentation was $-1.5e06$ and is omitted.

5.2 Evaluation Using Automatic Metrics

Our results on the English-to-Japanese transliteration task are summarized in Table 2. It is clear from the table that using sequence-pairs from only the sample at the final iteration of the training produces gave lower performance than the baseline system. The phrase-table derived in this way contained only 3372 sequence-pairs as opposed to over 140,000 in the phrase-table extracted from the GIZA++ alignments.




Fig. 4 The evolution of the log-probability of the sampled derivation with respect to the training iteration.

Moreover these sequence-pairs were very short compared to those in from the baseline system’s phrase-table: approximately 3 characters in both source and target on average, compared to around 5 characters for the baseline system.

When a phrase-table built from agglomerations of the same set of sequence-pairs was used, a much larger phrase-table of around 100,000 phrases resulted, with sequence-pairs that are comparable in size to those of the baseline, around 5 characters. On the transliteration task, this phrase-table gave an improvement of approximately 1% in ACC over the baseline system, from a phrase-table that was about 30% smaller in size. Moreover, since the sequence-pairs are concatenations of 3372 component sequence-pairs, this model could be stored very compactly if necessary. Further gains were obtained by interpolating the agglomerated model together with the baseline model. We believe this gain may be due to the effect of smoothing.

Our experiments were designed to favor the baseline model since the system was tuned using the MERT procedure with its own phrase-table. It is possible that our proposed system would have obtained a higher score if tuned with its own phrase-table, however we chose not to as this would have introduced additional variance from the differences in the two MERT search processes into the results. However, to verify that the agglomeration step was truly necessary, we also ran an experiment that was tuned with respect the small unagglomerated phrase-table: ‘tuned on Bayesian phrase-table’ in Fig. 2, rather than the baseline. As expected, performance improved possibly due to more optimal weights for the length models, since the phrases are shorter, but did not improve to the same level as the system that used the agglomerated phrase table.

It is interesting to note that the system’s performance was improved dramatically simply by grouping the phrases into larger units. This highlights one of the advantages of the phrase-based translation approach. The agglomerated model, because of the way it was constructed, is not able to generate anything the simpler model cannot, but when larger sequence-pairs are used to build the target sequence the characters in the phrase carry with them the *implicit con-*

Table 2 The experimental results for the three systems together with some statistics of their phrase-tables. Here *+agglomerated* means the sequence-pairs were extracted by agglomeration from a single sample at the end of the training. In *+integrated* the phrase-tables from the baseline system and the agglomerated system were linearly interpolated with equal weights. Differences between systems were all found to be significant by paired t-testing at a level of 0.05, except for the ACC scores for the agglomerated and integrated systems.

Phrase Extraction Model	ACC	F-score	Phrase-table Entries	Phrase-table Overlap (%)	Avg. Phrase Length	
					En	Ja
GIZA++ and <i>grow-diag-final-and</i>	0.313	0.745	143382	100	5.41	4.80
Bayesian Segmentation (tuned on baseline phrase-table)	0.278	0.726	3372	2	2.60	2.75
Bayesian Segmentation (tuned on Bayesian phrase-table)	0.283	0.732	3372	2	2.60	2.75
Bayesian Segmentation (+agglomerated)	0.323	0.748	102507	57	5.54	4.83
Bayesian Segmentation (+integrated)	0.329	0.752	164258	100	5.46	4.81

text of the other characters in the phrase, all of which have occurred together in the same context in the training corpus. In the model with the unagglomerated sequence-pairs, this role is performed mainly by the language model. In spite of the fact that we used a 5-gram language model the system clearly benefited from a model that contained longer sequence-pairs as the basic translation unit. Looking at the phrase-table overlap figures in Fig. 2 it seems that the process of agglomeration produced a phrase table with a reasonable degree of overlap (57%) with that produced by using GIZA++ and the MOSES phrase extraction procedure.

In a final experiment we collected counts for the sequence-pairs over multiple iterations of the training process: from iteration 5 (the *burn-in*) to the final iteration. This resulted in a 37% larger phrase-table, but surprisingly did not realize any notable improvement in performance.

5.3 Decoding Consistency

We ran an experiment to investigate the reasons for the improvements in system performance. Our hypothesis was that the Bayesian system had produced a phrase table that led to a more consistent decoding process. This was based on the belief that the fact that the Dirichlet process model strongly encourages reuse of the bilingual sequence-pairs it discovers. This should result in a more compact phrase-table, and should entail that similar words in the corpus are likely to be decoded in more homogenous fashion. To test the hypothesis we modified the machine translation decoder to count the number of *types* of bilingual sequence-pair used to decode the evaluation data, and re-ran the English-Japanese transliteration experiment that showed the largest gain in performance. We found that the decoding process that used the phrase-table generated from our Bayesian model (with agglomerated sequence-pairs) used a total of 3496 unique sequence-pairs, whereas decoding using the phrase-table extracted using GIZA++ and *grow-diag-final-and* required a total of 3970 phrase pairs during the decoding process, supporting our hypothesis. The 3496 sequence-pairs from the Bayesian model's phrase-table, could be further analysed into 1289 component bilingual pairs that were present in the segmentation in the sample taken at the end of the training process.

5.4 Training Times

Bayesian methods are often criticized for their slow performance. The current implementation will require optimization to enable it to handle long sequences, but for transliteration data, where the sequence length is short, the process is practicable. The full training process completed on the NEWS2010 training data set in 15 minutes, averaging around 30 seconds for each iteration over the data.

6. Human Evaluation using Speech-to-Speech Translation Field Data

In this subsection, we report the results from experiments in machine translation carried out to evaluate the effectiveness of our transliteration method with real-word data collected from the application of mobile translation devices in the field.

The test set for this evaluation was extracted from user log data of a set of speech-to-speech translation field experiments that occurred in the fiscal year 2009 [25]. As shown in Fig. 5, the field experiments were undertaken in nationwide in Japan in five broad regions of Japan: Kanto, Kansai, Kyushu, Hokkaido and Chubu. The field experiments were undertaken as part of the Ministry of Internal Affairs and Communications initiative titled "Field Testing of Automatic Speech Translation Technology and its Contribution to Local Tourism." We sampled 100 sentences from manually transcribed user log data of the field experiments from each of the five areas. Then we selected only those sentences that contained untranslatable named entities, leaving a total of 74 sentences. All of the test sentences contained at least one proper noun written in *hiragana* or *kanji* characters.

The translation direction of the evaluation is from Japanese into English. As a baseline system we used a standard phrase-based statistical machine translation system trained on the BTEC corpus consisting of 691,829 Japanese and English sentence pairs. To see the upper bound of the machine translation performance, we manually constructed a bilingual dictionary which consists of word categories and English translations of all proper nouns in the test set. Prior to the translation, proper noun in source sentences are re-

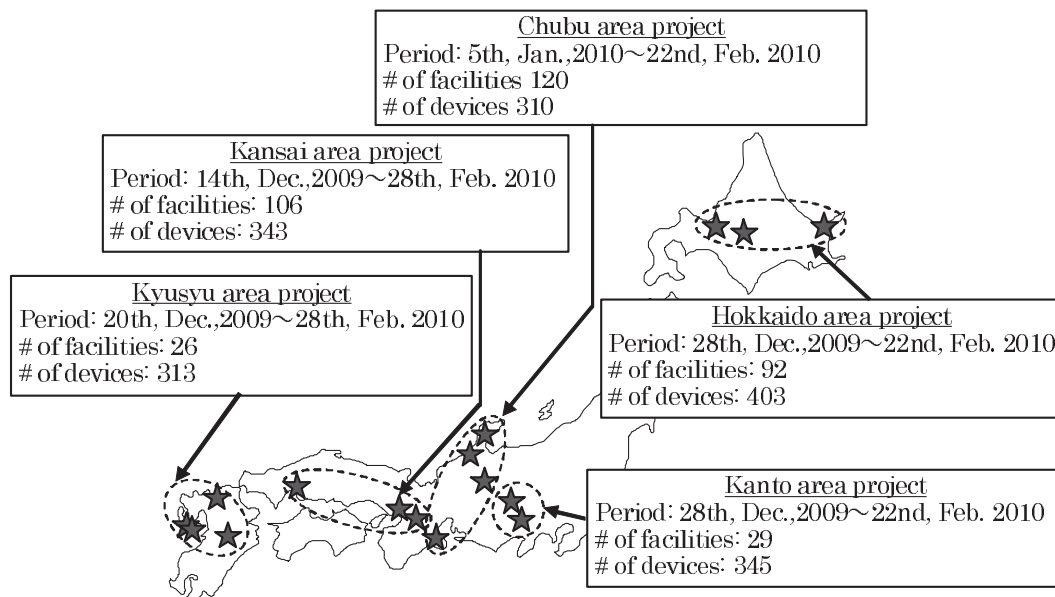


Fig. 5 Overview of the five local projects.

placed with high-frequency words of the same category in the training corpus [26]. The target sentences are then acquired by translating the modified source sentences. Finally, the high-frequency words in the target sentences are replaced with target words for the untrained words. The reason why we use high-frequency words is that we expect them to be already trained well, in other words, the high-frequency words may already appear frequently in phrase tables and therefore provide ample statistics. A single human judge was used to score the translation output, according to a five-level scale ranging from ‘perfect’ through ‘acceptable’ to ‘nonsense’.

We evaluated the transliteration component of the system on the transliteration task used in our field experiments. Each source word was transliterated from kanji into romaji, and the output evaluated against a reference set with up to 3 references per input sequence using the same scoring metrics as the NEWS2010 workshop. The transliteration accuracy on this task was 19.14%, this is lower than the NEWS2010 task scores but this is to be expected as our task contains both multi-word sequences and also sequences that are not transliterations of each other. Nonetheless the F-score, which is a character level score based on the length of the longest common subsequence, was 71.03, indicating the performance is quite respectable at the character level.

Visual inspection of the output, showed that the biggest single cause of errors in the system is due to the assumption that all words should be transliterated. In reality, expressions may need to be translated, or partially transliterated and partially translated. One such example from the test corpus being the expression in Kanji: 伊丹空港. This should be correctly transliterated as “Itami Airport”, however the transliteration system produced “itamikuukou” as output - a perfect transcription into kana, but incorrect nonetheless. Future research will need to address ways to identifying

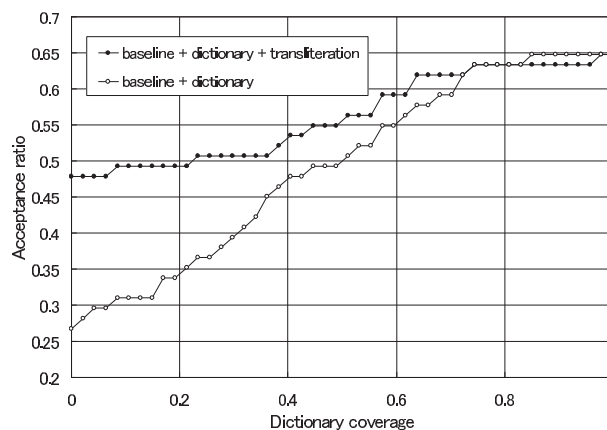


Fig. 6 Results of machine translation evaluation.

when to transliterate and when to translate. The problem of re-ordering was also an issue for our approach, although these errors were found in less than 4% of the sentences in the test corpus. An example of this type of error is 富士山. The system’s output in this case was “fujisan”, however the correct outputs were “Mount Fuji” or “Mt. Fuji”. In this example, both transliteration and translation are required, but in addition the order of the words is swapped. We believe that modeling this re-ordering process would give rise to improvements in system performance.

Figure 6 shows the field test translation evaluation results. The vertical axis represents the acceptance ratio which is the ratio of better than acceptable translation sentences to the total number of test sentences. To see the relationship between the dictionary size and the translation performance, we controlled the dictionary size. The horizontal axis shows dictionary coverage, which is the ratio of the proper nouns covered by the dictionary to the total number of occurrences

of proper nouns in the evaluation set. \circ plots shows machine translation with conventional dictionary usage [26]. While, \bullet plots arised from translation strategy that uses the transliterated results *only if* there is no entry in the dictionary[†].

It is clear from the figure that the use of transliteration within a machine translation system can lead to improvements in translation quality from a user's perspective. In Fig. 6, the values at 0 on the x-axis represent the system performance without the ability to look up the correct translations for named entities in a dictionary. The translation performance is significantly higher at this point. The differences between the two lines on the graph gradually diminish as the dictionary is used more and more for the translation of named entities, and at a value of 1 on the x-axis of course the two lines meet, as the strategies are equivalent to each other at this point. When dictionary coverage is low, we can expect transliteration to have a substantial impact on system performance, however, when we look at the line "baseline + dictionary + transliteration" at both 0 and 1 on the x-axis we can see the effect of the errors arising from the transliteration system. Adding the dictionary to the transliteration system clearly improves the overall system performance. Therefore, further work is needed to improve the quality of machine transliteration systems.

7. Conclusion

In this paper we have presented a novel Bayesian bilingual co-segmentation scheme and applied it to the task of phrase-table generation for transliteration by phrase-based statistical machine translation. Furthermore, we have conducted a human evaluation to investigate the effectiveness of integrating a machine transliteration system with a machine translation system. Our results show that adding transliterations of unknown named entities into the machine translation output has a positive effect on translation quality from a user's perspective.

Our Bayesian model for bilingual segmentation was motivated by the fact that traditional models of phrasal alignment rely on maximum likelihood training coupled with the EM algorithm, but have serious issues with overfitting the training data. Because of these issues, alignment is typically performed in a one-to-many manner from source-to-target and from target-to-source and the phrase extraction process proceeds heuristically from an alignment table. Our approach offers the ability to align the training data in a many-to-many fashion directly using Bayesian techniques that offer a simple yet elegant solution to the issues inherent in maximum likelihood training. In addition, our approach is symmetrical with respect to source and target, and also with respect to the word order of the corpus.

We investigated the quality of the bilingual phrasal

[†]To incorporate transliterated results into machine translation, we used the same framework as the conventional dictionary-based technique which was proposed by [26]. In this method, category of the word must be known. In our experiments, even for the transliterated words, we used manually assigned category.

alignment achievable with unsupervised Bayesian co-segmentation, and designed experiments to compare directly to a standard GIZA++/grow-diag-final-and phrase extraction procedure by constructing a phrase-table from samples arising from the Gibbs sampling training procedure. Our experiments show that the Bayesian approach is able to produce a smaller phrase-table that can offer comparable or higher transliteration performance than the baseline system.

Furthermore, our technique offers other benefits: one example being that it provides a full co-segmentation of the training corpus at the end of training which can be used to directly train a joint sequence model. This contextual information is a key feature in joint sequence transliteration models such as [9], and is currently missing from the phrase-based SMT-based transliteration systems. Another virtue of our approach stems from the fact that the Dirichlet process model is able to assign a probability to any bilingual word pair. We believe this type of model in has considerable potential utility in transliteration mining and corpus filtering, since it provides a principled way of scoring any potential transliteration candidate.

In future research we would like to investigate the effect of introducing a joint sequence model feature into a phrase-based SMT-based transliteration system. We also plan to improve the underlying Dirichlet process model in order to better model the data, moving to higher-order and hierarchical models.

References

- [1] K. Papineni, S. Roukos, T. Ward, and W.J. Zhu, "Bleu: A method for automatic evaluation of machine translation," ACL '02: Proc. 40th Annual Meeting on Association for Computational Linguistics, pp.311-318, Morristown, NJ, USA, Association for Computational Linguistics, 2001.
- [2] G. Doddington, "Automatic evaluation of machine translation quality using N-gram co-occurrence statistics," Proc. HLT Conference, San Diego, California, 2002.
- [3] X. Duan, D. Xiong, H. Zhang, M. Zhang, and H. Li, "I²R's machine translation system for iwslt 2009," Proc. International Workshop on Spoken Language Translation, pp.50-54, 2009.
- [4] A. Finch and E. Sumita, "Phrase-based machine transliteration," Proc. Workshop on Technologies and Corpora for Asia-Pacific Speech Translation (TCAST), pp.13-18, Hyderabad, India, Jan. 2008.
- [5] T. Rama and K. Gali, "Modeling machine transliteration as a phrase based statistical machine translation problem," NEWS '09: Proc. 2009 Named Entities Workshop: Shared Task on Transliteration, pp.124-127, Morristown, NJ, USA, Association for Computational Linguistics, 2009.
- [6] S. Noeman, "Language independent transliteration system using phrase based smt approach on substrings," NEWS '09: Proc. 2009 Named Entities Workshop: Shared Task on Transliteration, pp.112-115, Morristown, NJ, USA, Association for Computational Linguistics, 2009.
- [7] F.J. Och and H. Ney, "A systematic comparison of various statistical alignment models," Computational Linguistics, vol.29, no.1, pp.19-51, 2003.
- [8] P. Koehn, H. Hoang, A. Birch, C. Callison-Burch, M. Federico, N. Bertoldi, B. Cova, W. Shen, C. Moran, R. Zens, C. Dyer, O. Bojar, A. Constantin, and E. Herbst, "Moses: Open source toolkit for statistical machine translation," ACL 2007: Proc. Demo and Poster Ses-

sions, pp.177–180, Prague, Czeck Republic, June 2007.

[9] H. Li, M. Zhang, and J. Su, “A joint source-channel model for machine transliteration,” *ACL ’04: Proc. 42nd Annual Meeting on Association for Computational Linguistics*, p.159, Morristown, NJ, USA, Association for Computational Linguistics, 2004.

[10] D. Yang, P. Dixon, Y.C. Pan, T. Oonishi, M. Nakamura, and S. Furui, “Combining a two-step conditional random field model and a joint source channel model for machine transliteration,” *NEWS ’09: Proc. 2009 Named Entities Workshop: Shared Task on Transliteration*, pp.72–75, Morristown, NJ, USA, Association for Computational Linguistics, 2009.

[11] D. Marcu and W. Wong, “A phrase-based, joint probability model for statistical machine translation,” *Proc. EMNLP*, pp.133–139, 2002.

[12] M. Bisani and H. Ney, “Joint-sequence models for grapheme-to-phoneme conversion,” *Speech Commun.*, vol.50, no.5, pp.434–451, 2008.

[13] P. Blunsom, T. Cohn, C. Dyer, and M. Osborne, “A gibbs sampler for phrasal synchronous grammar induction,” *Proc. Joint Conference of the 47th Annual Meeting of the ACL*, pp.782–790, Suntec, Singapore, Association for Computational Linguistics, Aug. 2009.

[14] J. DeNero, A. Bouchard-Côté, and D. Klein, “Sampling alignment structure under a bayesian translation model,” *Proc. Conference on Empirical Methods in Natural Language Processing, EMNLP ’08*, pp.314–323, Stroudsburg, PA, USA, Association for Computational Linguistics, 2008.

[15] G. Neubig, T. Watanabe, E. Sumita, S. Mori, and T. Kawahara, “An unsupervised model for joint phrase alignment and extraction,” *49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies (ACL-HLT)*, Portland, Oregon, USA, June 2011.

[16] J. Wuebker, A. Mauser, and H. Ney, “Training phrase translation models with leaving-one-out,” *Proc. 48th Annual Meeting of the Association for Computational Linguistics*, pp.475–484, Uppsala, Sweden, Association for Computational Linguistics, July 2010.

[17] J. Xu, J. Gao, K. Toutanova, and H. Ney, “Bayesian semi-supervised chinese word segmentation for statistical machine translation,” *COLING ’08: Proc. 22nd International Conference on Computational Linguistics*, pp.1017–1024, Morristown, NJ, USA, Association for Computational Linguistics, 2008.

[18] D. Mochihashi, T. Yamada, and N. Ueda, “Bayesian unsupervised word segmentation with nested pitman-yor language modeling,” *ACL-IJCNLP ’09: Proc. Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP: Volume 1*, pp.100–108, Morristown, NJ, USA, Association for Computational Linguistics, 2009.

[19] S. Goldwater, T.L. Griffiths, and M. Johnson, “Contextual dependencies in unsupervised word segmentation,” *ACL-44: Proc. 21st International Conference on Computational Linguistics and the 44th annual meeting of the Association for Computational Linguistics*, pp.673–680, Morristown, NJ, USA, Association for Computational Linguistics, 2006.

[20] D.J. Aldous, “Exchangeability and related topics,” in *École d’été de probabilités de Saint-Flour, XIII—1983, Lecture Notes in Math.*, vol.1117, pp.1–198, 1985.

[21] F.J. Och and H. Ney, “Discriminative training and maximum entropy models for statistical machine translation,” *Proc. 40th Annual Meeting of the Association for Computational Linguistics (ACL 2002)*, pp.295–302, 2002.

[22] F.J. Och, “Minimum error rate training for statistical machine translation,” *Proc. ACL*, 2003.

[23] T. Rama and K. Gali, “Modeling machine transliteration as a phrase based statistical machine translation problem,” *Proc. ACL/IJCNLP Named Entities Workshop Shared Task*, 2009.

[24] M.Z. Haizhou Li, A. Kumaran, and V. Pervouchine, “Whitepaper of news 2010 shared task on transliteration generation,” *Proc. ACL Named Entities Workshop Shared Task*, 2010.

[25] H. Kawai, R. Isotani, K. Yasuda, E. Sunita, U. Masao, S. Matsuda, Y. Ashikari, and S. Nakamura, “An overview of a nation-wide field experiment of speech-to-speech translation in fiscal year 2009 (Japanese only),” *Proc. 2010 Autumn Meeting of Acoustical Society of Japan*, pp.99–102, 2010.

[26] H. Okuma, H. Yamamoto, and E. Sumita, “Introducint a translation dictionary into phrase-based SMT,” *IEICE Trans. Inf. & Syst.*, vol.E91-D, no.7, pp.2051–2057, July 2008.



Andrew Finch received the B.S. degree in mathematics and the M.Sc. degree in cognition, computing and psychology, both from the University of Warwick, England, in 1984 and 1990, respectively. He received a Ph.D. in computer science in 1995 from the University of York, England. From 1995 to 1997 he worked as a postdoctoral researcher in the computer vision research group at the University of York. In 1997 he received an Honorable Mention of the Pattern Recognition Society Award for Outstanding Contribution to the Pattern Recognition journal. He is currently an expert researcher at NICT in Kyoto. His current research interests include Bayesian methods, machine translation, automatic paraphrasing, and transliteration.



Keiji Yasuda was born in Osaka, Japan, in 1976. He received his M.E. and Ph.D. from Doshisha University in 2001 and 2004. He is currently a researcher at the National Institute of Communications Technology. His research interests include speech recognition, natural language processing, and e-Learning. He received a society paper award from the IEICE-ISS in 2006. He is a member of IPSJ and ASJ.



Hideo Okuma received the B.S. degree in mathematics from the Tokyo University of Science 1987. Mr. Okuma is currently a researcher at National Institute of Communications Technology. His research interests include machine translation.



Eiichiro Sumita received the M.S. degree in computer science from the University of Electro-Communications in 1982 and the doctor degree in engineering from Kyoto University in 1999. Dr. Sumita is the director of NICT/MASTAR Project and the multilingual translation lab. He is a visiting professor of Kobe University. His research interests include machine translation and e-Learning.



Satoshi Nakamura received his B.S. from Kyoto Institute of Technology in 1981 and Ph.D. from Kyoto University in 1992. He was an Associate Professor of the Graduate School of Information Science at Nara Institute of Science and Technology in 1994–2000. He was Director of ATR Spoken Language Communication Research Laboratories in 2000–2008, and Director General of Keihanna Research Laboratories and the Executive Director of Knowledge Creating Communication Research Center, National Institute of Information and Communications Technology, Japan in 2009–2010. He is ATR Fellow. He is currently Professor of Graduate School of Information Science at Nara Institute of Science and Technology. He received Yamashita Research Award, Kiyasu Award from the Information Processing Society of Japan, Telecom System Award, AAMT Nagao Award, Docomo Mobile Science Award in 2007, ASJ Award for Distinguished Achievements in Acoustics, and the Commendation for Science and Technology by the Minister of Education, Culture, Sports, Science and Technology.