

A Comparative Study on Vocabulary Reduction for Phrase Table Smoothing

Yunsu Kim, Andreas Guta, Joern Wuebker*, and Hermann Ney

Human Language Technology and Pattern Recognition Group

RWTH Aachen University, Aachen, Germany

{surname}@cs.rwth-aachen.de

*Lilt, Inc.

joern@lilt.com

Abstract

This work systematically analyzes the smoothing effect of vocabulary reduction for phrase translation models. We extensively compare various word-level vocabularies to show that the performance of smoothing is not significantly affected by the choice of vocabulary. This result provides empirical evidence that the standard phrase translation model is extremely sparse. Our experiments also reveal that vocabulary reduction is more effective for smoothing large-scale phrase tables.

1 Introduction

Phrase-based systems for statistical machine translation (SMT) (Zens et al., 2002; Koehn et al., 2003) have shown state-of-the-art performance over the last decade. However, due to the huge size of phrase vocabulary, it is difficult to collect robust statistics for lots of phrase pairs. The standard phrase translation model thus tends to be sparse (Koehn, 2010).

A fundamental solution to a sparsity problem in natural language processing is to reduce the vocabulary size. By mapping words onto a smaller label space, the models can be trained to have denser distributions (Brown et al., 1992; Miller et al., 2004; Koo et al., 2008). Examples of such labels are part-of-speech (POS) tags or lemmas.

In this work, we investigate the vocabulary reduction for phrase translation models with respect to various vocabulary choice. We evaluate two types of smoothing models for phrase translation probability using different kinds of word-level labels. In particular, we use automatically generated word classes (Brown et al., 1992) to obtain

label vocabularies with arbitrary sizes and structures. Our experiments reveal that the vocabulary of the smoothing model has no significant effect on the end-to-end translation quality. For example, a randomized label space also leads to a decent improvement of BLEU or TER scores by the presented smoothing models.

We also test vocabulary reduction in translation scenarios of different scales, showing that the smoothing works better with more parallel corpora.

2 Related Work

Koehn and Hoang (2007) propose integrating a label vocabulary as a factor into the phrase-based SMT pipeline, which consists of the following three steps: mapping from words to labels, label-to-label translation, and generation of words from labels. Rishøj and Søgaard (2011) verify the effectiveness of word classes as factors. Assuming probabilistic mappings between words and labels, the factorization implies a combinatorial expansion of the phrase table with regard to different vocabularies.

Wuebker et al. (2013) show a simplified case of the factored translation by adopting hard assignment from words to labels. In the end, they train the existing translation, language, and reordering models on word classes to build the corresponding smoothing models.

Other types of features are also trained on word-level labels, e.g. hierarchical reordering features (Cherry, 2013), an n -gram-based translation model (Durrani et al., 2014), and sparse word pair features (Haddow et al., 2015). The first and the third are trained with a large-scale discriminative training algorithm.

For all usages of word-level labels in SMT,

a common and important question is which label vocabulary maximizes the translation quality. Bisazza and Monz (2014) compare class-based language models with diverse kinds of labels in terms of their performance in translation into morphologically rich languages. To the best of our knowledge, there is no published work on systematic comparison between different label vocabularies, model forms, and training data size for smoothing phrase translation models—the most basic component in state-of-the-art SMT systems. Our work fulfills these needs with extensive translation experiments (Section 5) and quantitative analysis (Section 6) in a standard phrase-based SMT framework.

3 Word Classes

In this work, we mainly use unsupervised word classes by Brown et al. (1992) as the reduced vocabulary. This section briefly reviews the principle and properties of word classes.

A word-class mapping c is estimated by a clustering algorithm that maximizes the following objective (Brown et al., 1992):

$$\mathcal{L} := \sum_{e_1^I} \sum_{i=1}^I p(c(e_i)|c(e_{i-1})) \cdot p(e_i|c(e_i)) \quad (1)$$

for a given monolingual corpus $\{e_1^I\}$, where each e_1^I is a sentence of length I in the corpus. The objective guides c to prefer certain collocations of class sequences, e.g. an auxiliary verb class should succeed a class of pronouns or person names. Consequently, the resulting c groups words according to their syntactic or semantic similarity.

Word classes have a big advantage for our comparative study: The structure and size of the class vocabulary can be arbitrarily adjusted by the clustering parameters. This makes it possible to prepare easily an abundant set of label vocabularies that differ in linguistic coherence and degree of generalization.

4 Smoothing Models

In the standard phrase translation model, the translation probability for each segmented phrase pair (\tilde{f}, \tilde{e}) is estimated by relative frequencies:

$$p_{\text{std}}(\tilde{f}|\tilde{e}) = \frac{N(\tilde{f}, \tilde{e})}{N(\tilde{e})} \quad (2)$$

where N is the count of a phrase or a phrase pair in the training data. These counts are very low for many phrases due to a limited amount of bilingual training data.

Using a smaller vocabulary, we can aggregate the low counts and make the distribution smoother. We now define two types of smoothing models for Equation 2 using a general word-label mapping c .

4.1 Mapping All Words at Once (map-all)

For the phrase translation model, the simplest formulation of vocabulary reduction is obtained by replacing all words in the source and target phrases with the corresponding labels in a smaller space. Namely, we employ the following probability instead of Equation 2:

$$p_{\text{all}}(\tilde{f}|\tilde{e}) = \frac{N(c(\tilde{f}), c(\tilde{e}))}{N(c(\tilde{e}))} \quad (3)$$

which we call *map-all*. This model resembles the word class translation model of Wuebker et al. (2013) except that we allow any kind of word-level labels.

This model generalizes all words of a phrase without distinction between them. Also, the same formulation is applied to word-based lexicon models.

4.2 Mapping Each Word at a Time (map-each)

More elaborate smoothing can be achieved by generalizing only a sub-part of the phrase pair. The idea is to replace one source word at a time with its respective label. For each source position j , we also replace the target words aligned to the source word f_j . For this purpose, we let $a_j \subseteq \{1, \dots, |\tilde{e}|\}$ denote a set of target positions aligned to j . The resulting model takes a weighted average of the redefined translation probabilities over all source positions of \tilde{f} :

$$p_{\text{each}}(\tilde{f}|\tilde{e}) = \sum_{j=1}^{|\tilde{f}|} w_j \cdot \frac{N(c^{(j)}(\tilde{f}), c^{(a_j)}(\tilde{e}))}{N(c^{(a_j)}(\tilde{e}))} \quad (4)$$

where the superscripts of c indicate the positions that are mapped onto the label space. w_j is a weight for each source position, where $\sum_j w_j = 1$. We call this model *map-each*.

We illustrate this model with a pair of three-word phrases: $\tilde{f} = [f_1, f_2, f_3]$ and $\tilde{e} = [e_1, e_2, e_3]$ (see Figure 1 for the in-phrase word alignments). The map-each model score for this phrase pair is:




Figure 1: Word alignments of a pair of three-word phrases.

$$\begin{aligned}
 p_{\text{each}}([f_1, f_2, f_3] | [e_1, e_2, e_3]) = & \\
 w_1 \cdot \frac{N(\overbrace{[c(f_1), f_2, f_3]}, \overbrace{[c(e_1), e_2, e_3]})}{N([c(e_1), e_2, e_3])} & \\
 + w_2 \cdot \frac{N([f_1, c(f_2), f_3], [e_1, e_2, e_3])}{N([e_1, e_2, e_3])} & \\
 + w_3 \cdot \frac{N(\overbrace{[f_1, f_2, c(f_3)]}, \overbrace{[e_1, c(e_2), c(e_3)]})}{N([e_1, c(e_2), c(e_3)])} &
 \end{aligned}$$

where the alignments are depicted by line segments.

First of all, we replace f_1 and also e_1 , which is aligned to f_1 , with their corresponding labels. As f_2 has no alignment points, we do not replace any target word accordingly. f_3 triggers the class replacement of two target words at the same time. Note that the model implicitly encapsulates the alignment information.

We empirically found that the map-each model performs best with the following weight:

$$w_j = \frac{N(c^{(j)}(\tilde{f}), c^{(a_j)}(\tilde{e}))}{\sum_{j'=1}^{|\tilde{f}|} N(c^{(j')}(\tilde{f}), c^{(a_{j'})}(\tilde{e}))} \quad (5)$$

which is a normalized count of the generalized phrase pair itself. Here, the count is relatively large when f_j , the word to be backed off, is less frequent than other words in \tilde{f} . In contrast, if f_j is a very frequent word and one of the other words in \tilde{f} is rare, the count becomes low due to that rare word. The same logic holds for target words in \tilde{e} . After all, Equation 5 carries more weight when a rare word is replaced with its label. The intuition is that a rare word is the main reason for unstable counts and should be backed off above all. We use this weight for all experiments in the next section.

In contrast, the map-all model merely replace all words at one time and ignore alignments within phrase pairs.

5 Experiments

5.1 Setup

We evaluate how much the translation quality is improved by the smoothing models in Section 4. The two smoothing models are trained in both source-to-target and target-to-source directions, and integrated as additional features in the log-linear combination of a standard phrase-based SMT system (Koehn et al., 2003). We also test linear interpolation between the standard and smoothing models, but the results are generally worse than log-linear interpolation. Note that vocabulary reduction models by themselves cannot replace the corresponding standard models, since this leads to a considerable drop in translation quality (Wuebker et al., 2013).

Our baseline systems include phrase translation models in both directions, word-based lexicon models in both directions, word/phrase penalties, a distortion penalty, a hierarchical lexicalized reordering model (Galley and Manning, 2008), a 4-gram language model, and a 7-gram word class language model (Wuebker et al., 2013). The model weights are trained with minimum error rate training (Och, 2003). All experiments are conducted with an open source phrase-based SMT toolkit Jane 2 (Wuebker et al., 2012).

To validate our experimental results, we measure the statistical significance using the paired bootstrap resampling method of Koehn (2004). Every result in this section is marked with ‡ if it is statistically significantly better than the baseline with 95% confidence, or with † for 90% confidence.

5.2 Comparison of Vocabularies

The presented smoothing models are dependent on the label vocabulary, which is defined by the word-label mapping c . Here, we train the models with various label vocabularies and compare their smoothing performance.

The experiments are done on the IWSLT 2012 German→English shared translation task. To rapidly perform repetitive experiments, we train the translation models with the in-domain TED portion of the dataset (roughly 2.5M running words for each side). We run the monolingual word clustering algorithm of (Botros et al., 2015) on each side of the parallel training data to obtain class label vocabularies (Section 3).

We carry out comparative experiments regarding the three factors of the clustering algorithm:

1) Clustering iterations. It is shown that the number of iterations is the most influential factor in clustering quality (Och, 1995). We now verify its effect on translation quality when the clustering is used for phrase table smoothing.

As we run the clustering algorithm, we extract an intermediate class mapping for each iteration and train the smoothing models with it. The model weights are tuned for each iteration separately. The BLEU scores of the tuned systems are given in Figure 2. We use 100 classes on both source and target sides.

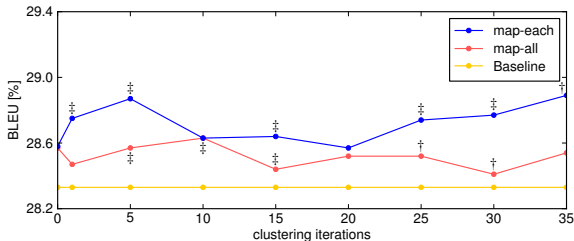


Figure 2: BLEU scores for clustering iterations when using individually tuned model weights for each iteration. Dots indicate those iterations in which the translation is performed.

The score does not consistently increase or decrease over the iterations; it is rather on a similar level ($\pm 0.2\%$ BLEU) for all settings with slight fluctuations. This is an important clue that the whole process of word clustering has no meaning in smoothing phrase translation models.

To see this more clearly, we keep the model weights fixed over different systems and run the same set of experiments. In this way, we focus only on the change of label vocabulary, removing the impact of non-deterministic model weight optimization. The results are given in Figure 3.

This time, the curves are even flatter, resulting in only $\pm 0.1\%$ BLEU difference over the iterations. More surprisingly, the models trained with the initial clustering, i.e. when the clustering algorithm has not even started yet, are on a par with those trained with more optimized classes in terms of translation quality.

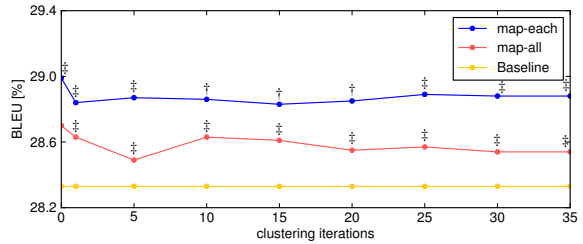


Figure 3: BLEU scores for clustering iterations when using a fixed set of model weights. The weights that produce the best results in Figure 2 are chosen.

2) Initialization of the clustering. Since the clustering process has no significant impact on the translation quality, we hypothesize that the initialization may dominate the clustering. We compare five different initial class mappings:

- random: randomly assign words to classes
- top-frequent (default): top-frequent words have their own classes, while all other words are in the last class
- same-countsum: each class has almost the same sum of word unigram counts
- same-#words: each class has almost the same number of words
- count-bins: each class represents a bin of the total count range

| | Initialization | BLEU [%] | TER [%] |
|------------|----------------|-------------------|-------------------|
| Baseline | | 28.3 | 52.2 |
| + map-each | random | 28.9 [‡] | 51.7 [‡] |
| | top-frequent | 29.0 [‡] | 51.5 [‡] |
| | same-countsum | 28.8 [‡] | 51.7 [‡] |
| | same-#words | 28.9 [‡] | 51.6 [‡] |
| | count-bins | 29.0 [‡] | 51.4 [‡] |

Table 1: Translation results for various initializations of the clustering. 100 classes on both sides.

Table 1 shows the translation results with the map-each model trained with these initializations—without running the clustering algorithm. We use the same set of model weights used in Figure 3. We find that the initialization method also does not affect the translation performance. As an extreme case,

random clustering is also a fine candidate for training the map-each model.

- 3) **Number of classes.** This determines the vocabulary size of a label space, which eventually adjusts the smoothing degree. Table 2 shows the translation performance of the map-each model with a varying number of classes. Similarly as before, there is no serious performance gap among different word classes, and POS tags and lemmas also conform to this trend.

However, we observe a slight but steady degradation of translation quality ($\approx -0.2\%$ BLEU) when the vocabulary size is larger than a few hundreds. We also lose statistical significance for BLEU in these cases. The reason could be: If the label space becomes larger, it gets closer to the original vocabulary and therefore the smoothing model provides less additional information to add to the standard phrase translation model.

| | #vocab (source) | BLEU [%] | TER [%] |
|--------------------|--------------------|-------------------|-------------------|
| Baseline | | 28.3 | 52.2 |
| + map-each | 100 | 29.0 [‡] | 51.5 [‡] |
| (word class) | 200 | 28.9 [†] | 51.6 [‡] |
| | 500 | 28.7 | 51.8 [‡] |
| | 1000 | 28.7 | 51.8 [‡] |
| | 10000 | 28.7 | 51.9 [†] |
| + map-each (POS) | 52 | 28.9 [†] | 51.5 [‡] |
| + map-each (lemma) | 26744 | 28.8 | 51.7 [‡] |

Table 2: Translation results for different vocabulary sizes.

The series of experiments show that the map-each model performs very similar across vocabulary size and its structure. From our internal experiments, this argument also holds for the map-all model. The results do not change even when we use a different clustering algorithm, e.g. bilingual clustering (Och, 1999). For the translation performance, the more important factor is the log-linear model training to find an optimal set of weights for the smoothing models.

5.3 Comparison of Smoothing Models

Next, we compare the two smoothing models by their performance in four different trans-

lation tasks: IWSLT 2012 German→English, WMT 2015 Finnish→English, WMT 2014 English→German, and WMT 2015 English→Czech. We train 100 classes on each side with 30 clustering iterations starting from the default (top-frequent) initialization.

Table 3 provides the corpus statistics of all datasets used. Note that a morphologically rich language is on the source side for the first two tasks, and on the target side for the last two tasks. According to the results (Table 4), the map-each model, which encourages backing off infrequent words, performs consistently better (maximum +0.5% BLEU, -0.6% TER) than the map-all model in all cases.

5.4 Comparison of Training Data Size

Lastly, we analyze the smoothing performance for different training data sizes (Figure 4). The improvement of BLEU score over the baseline decreases drastically when the training data get smaller. We argue that this is because the smoothing models are only the additional scores for the phrases seen in the training data. For smaller training data, we have more out-of-vocabulary (OOV) words in the test set, which cannot be handled by the presented models.

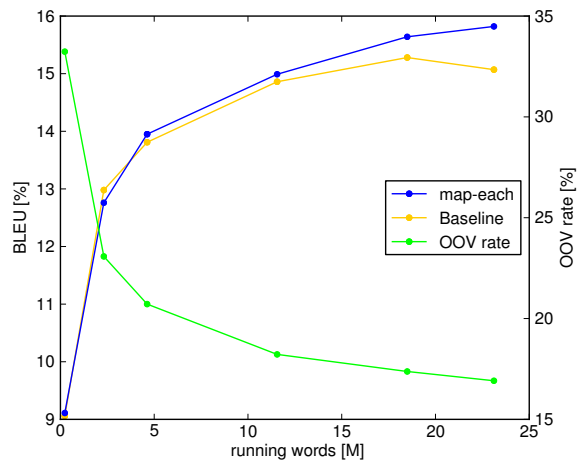


Figure 4: BLEU scores and OOV rates for the varying training data portion of WMT 2015 Finnish→English data.

6 Analysis

In Section 5.2, we have shown experimentally that more optimized or more fine-grained classes do not guarantee better smoothing performance. We now verify by examining translation outputs that

| | IWSLT 2012 | | WMT 2015 | | WMT 2014 | | WMT 2015 | |
|---------------|------------|---------|----------|---------|----------|--------|----------|-------|
| | German | English | Finnish | English | English | German | English | Czech |
| Sentences | 130k | | 1.1M | | 4M | | 0.9M | |
| Running Words | 2.5M | 2.5M | 23M | 32M | 104M | 105M | 23.9M | 21M |
| Vocabulary | 71k | 49k | 509k | 88k | 648k | 659k | 161k | 345k |

Table 3: Bilingual training data statistics for IWSLT 2012 German→English, WMT 2015 Finnish→English, WMT 2014 English→German, and WMT 2015 English→Czech tasks.

| | de-en | | fi-en | | en-de | | en-cs | |
|------------|-------------------------|-------------------------|-------------------------|-------------------------|-------------------------|-------------------------|-------------------------|-------------------------|
| | BLEU | TER | BLEU | TER | BLEU | TER | BLEU | TER |
| | [%] | [%] | [%] | [%] | [%] | [%] | [%] | [%] |
| Baseline | 28.3 | 52.2 | 15.1 | 72.6 | 14.6 | 69.8 | 15.3 | 68.7 |
| + map-all | 28.6 [‡] | 51.6 [‡] | 15.3 [‡] | 72.5 | 14.8 [‡] | 69.4 [‡] | 15.4 [‡] | 68.2 [‡] |
| + map-each | 29.0[‡] | 51.4[‡] | 15.8[‡] | 72.0[‡] | 15.1[‡] | 69.0[‡] | 15.8[‡] | 67.6[‡] |

Table 4: Translation results for IWSLT 2012 German→English, WMT 2015 Finnish→English, WMT 2014 English→German, and WMT 2015 English→Czech tasks.

| Model | Classes | #vocab | Top 200 TER-improved Sentences | |
|----------|---------------|--------|--------------------------------|------------------|
| | | | Common Input | Same Translation |
| | | | [%] | [%] |
| map-each | optimized | 100 | - | - |
| | non-optimized | 100 | 89.5 | 89.9 |
| | random | 100 | 88.5 | 89.8 |
| | lemma | 26744 | 87.0 | 92.6 |
| map-all | optimized | 100 | 56.0 | 54.5 |

Table 5: Comparison of translation outputs for the smoothing models with different vocabularies. “optimized” denotes 30 iterations of the clustering algorithm, whereas “non-optimized” means the initial (default) clustering.

the same level of performance is not by chance but due to similar hypothesis scoring across different systems.

Given a test set, we compare its translations generated from different systems as follows. First, for each translated set, we sort the sentences by how much the sentence-level TER is improved over the baseline translation. Then, we select the top 200 sentences from this sorted list, which represent the main contribution to the decrease of TER. In Table 5, we compare the top 200 TER-improved translations of the map-each model setups with different vocabularies.

In the fourth column, we trace the input sentences that are translated by the top 200 lists, and count how many of those inputs are overlapped across given systems. Here, a large overlap indi-

cates that two systems are particularly effective in a large common part of the test set, showing that they behaved analogously in the search process. The numbers in this column are computed against the map-each model setup trained with 100 optimized word classes (first row). For all map-each settings, the overlap is very large—around 90%.

To investigate further, we count how often the two translations of a single input are identical (the last column). This is normalized by the number of common input sentences in the top 200 lists between two systems. It is a straightforward measure to see if two systems discriminate translation hypotheses in a similar manner. Remarkably, all systems equipped with the map-each model produce exactly the same translations for the most part of the top 200 TER-improved sentences.

We can see from this analysis that, even though a smoothing model is trained with essentially different vocabularies, it helps the translation process in basically the same manner. For comparison, we also compute the measures for a map-all model, which are far behind the high similarity among the map-each models. Indeed, for smoothing phrase translation models, changing the model structure for vocabulary reduction exerts a strong influence in the hypothesis scoring, yet changing the vocabulary does not.

7 Conclusion

Reducing vocabulary using word-label mapping is a simple and effective way of smoothing phrase translation models. By mapping each word in a phrase at a time, the translation quality can be improved by up to +0.7% BLEU and -0.8% TER over a standard phrase-based SMT baseline, which is superior to Wuebker et al. (2013).

Our extensive comparison among various vocabularies shows that different word-label mappings are almost equally effective for smoothing phrase translation models. This allows us to use any type of word-level label, e.g. a randomized vocabulary, for the smoothing, which saves a considerable amount of effort in optimizing the structure and granularity of the label vocabulary. Our analysis on sentence-level TER demonstrates that the same level of performance stems from the analogous hypothesis scoring.

We claim that this result emphasizes the fundamental sparsity of the standard phrase translation model. Too many target phrase candidates are originally undervalued, so giving them any reasonable amount of extra probability mass, e.g. by smoothing with random classes, is enough to broaden the search space and improve translation quality. Even if we change a single parameter in estimating the label space, it does not have a significant effect on scoring hypotheses, where many other models than the smoothed translation model, e.g. language models, are involved with large weights. Nevertheless, an exact linguistic explanation is still to be discovered.

Our results on varying training data show that vocabulary reduction is more suitable for large-scale translation setups. This implies that OOV handling is more crucial than smoothing phrase translation models for low-resource translation tasks.

For future work, we plan to perform a similar set of comparative experiments on neural machine translation systems.

Acknowledgments

This paper has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement n° 645452 (QT21).

References

- Arianna Bisazza and Christof Monz. 2014. Class-based language modeling for translating into morphologically rich languages. In *Proceedings of 25th International Conference on Computational Linguistics (COLING 2014)*, pages 1918–1927, Dublin, Ireland, August.
- Rami Botros, Kazuki Irie, Martin Sundermeyer, and Hermann Ney. 2015. On efficient training of word classes and their application to recurrent neural network language models. In *Proceedings of 16th Annual Conference of the International Speech Communication Association (Interspeech 2015)*, pages 1443–1447, Dresden, Germany, September.
- Peter F. Brown, Peter V. deSouza, Robert L. Mercer, Vincent J. Della Pietra, and Jenifer C. Lai. 1992. Class-based n-gram models of natural language. *Computational Linguistics*, 18(4):467–479, December.
- Colin Cherry. 2013. Improved reordering for phrase-based translation using sparse features. In *Proceedings of 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT 2013)*, pages 22–31, Atlanta, GA, USA, June.
- Nadir Durrani, Philipp Koehn, Helmut Schmid, and Alexander Fraser. 2014. Investigating the usefulness of generalized word representations in smt. In *Proceedings of 25th Annual Conference on Computational Linguistics (COLING 2014)*, pages 421–432, Dublin, Ireland, August.
- Michel Galley and Christopher D. Manning. 2008. A simple and effective hierarchical phrase reordering model. In *Proceedings of 2008 Conference on Empirical Methods in Natural Language Processing (EMNLP 2008)*, pages 848–856, Honolulu, HI, USA, October.
- Barry Haddow, Matthias Huck, Alexandra Birch, Nikolay Bogoychev, and Philipp Koehn. 2015. The edinburgh/jhu phrase-based machine translation systems for wmt 2015. In *Proceedings of 2016 EMNLP 10th Workshop on Statistical Machine Translation (WMT 2016)*, pages 126–133, Lisbon, Portugal, September.

- Philipp Koehn and Hieu Hoang. 2007. Factored translation models. In *Proceedings of 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL 2007)*, pages 868–876, Prague, Czech Republic, June.
- Philipp Koehn, Franz Josef Och, and Daniel Marcu. 2003. Statistical phrase-based translation. In *Proceedings of 2003 Conference of the North American Chapter of the Association for Computational Linguistics on Human Language Technology (NAACL-HLT 2003)*, pages 48–54, Edmonton, Canada, May.
- Philipp Koehn. 2004. Statistical significance tests for machine translation evaluation. In *Proceedings of 2004 Conference on Empirical Methods in Natural Language Processing (EMNLP 2004)*, pages 388–395, Barcelona, Spain, July.
- Philipp Koehn. 2010. *Statistical Machine Translation*. Cambridge University Press, New York, NY, USA.
- Terry Koo, Xavier Carreras, and Michael Collins. 2008. Simple semi-supervised dependency parsing. In *Proceedings of 46th Annual Meeting of the Association for Computational Linguistics (ACL 2008)*, pages 595–603, Columbus, OH, USA, June.
- Scott Miller, Jethran Guinness, and Alex Zamanian. 2004. Name tagging with word clusters and discriminative training. In *Proceedings of 2004 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT 2004)*, pages 337–342, Boston, MA, USA, May.
- Franz Josef Och. 1995. Maximum-likelihood-schätzung von wortkategorien mit verfahren der kombinatorischen optimierung. Studienarbeit, Friedrich-Alexander-Universität Erlangen-Nürnberg, Erlangen, Germany, May.
- Franz Josef Och. 1999. An efficient method for determining bilingual word classes. In *Proceedings of 9th Conference on European Chapter of Association for Computational Linguistics (EACL 1999)*, pages 71–76, Bergen, Norway, June.
- Franz Josef Och. 2003. Minimum error rate training in statistical machine translation. In *Proceedings of 41st Annual Meeting of the Association for Computational Linguistics (ACL 2003)*, pages 160–167, Sapporo, Japan, July.
- Christian Rishøj and Anders Søgaard. 2011. Factored translation with unsupervised word clusters. In *Proceedings of 2011 EMNLP 6th Workshop on Statistical Machine Translation (WMT 2011)*, pages 447–451, Edinburgh, Scotland, July.
- Joern Wuebker, Matthias Huck, Stephan Peitz, Malte Nuhn, Markus Freitag, Jan-Thorsten Peter, Saab Mansour, and Hermann Ney. 2012. June 2: Open source phrase-based and hierarchical statistical machine translation. In *Proceedings of 24th International Conference on Computational Linguistics (COLING 2012)*, pages 483–492, Mumbai, India, December.
- Joern Wuebker, Stephan Peitz, Felix Rietig, and Hermann Ney. 2013. Improving statistical machine translation with word class models. In *Proceedings of 2013 Conference on Empirical Methods in Natural Language Processing (EMNLP 2013)*, pages 1377–1381, Seattle, USA, October.
- Richard Zens, Franz Josef Och, and Hermann Ney. 2002. Phrase-based statistical machine translation. In Matthias Jarke, Jana Koehler, and Gerhard Lakemeyer, editors, *25th German Conference on Artificial Intelligence (KI2002)*, volume 2479 of *Lecture Notes in Artificial Intelligence (LNAI)*, pages 18–32, Aachen, Germany, September. Springer Verlag.