
Dynamic Phrase Tables for Machine Translation in an Interactive Post-editing Scenario

Ulrich Germann
University of Edinburgh
ugermann@inf.ed.ac.uk

Abstract

This paper presents a phrase table implementation for the *Moses* system that computes phrase table entries for phrase-based statistical machine translation (PBSMT) on demand by sampling an indexed bitext. While this approach has been used for years in hierarchical phrase-based translation, the PBSMT community has been slow to adopt this paradigm, due to concerns that this would be slow and lead to lower translation quality. The experiments conducted in the course of this work provide evidence to the contrary: without loss in translation quality, the sampling phrase table ranks second out of four in terms of speed, being slightly slower than hash table look-up (Junczys-Dowmunt, 2012) and considerably faster than current implementations of the approach suggested by Zens and Ney (2007). In addition, the underlying parallel corpus can be updated in real time, so that professionally produced translations can be used to improve the quality of the machine translation engine immediately.

1 Introduction

In recent years, there has been an increasing interest in integrating machine translation (MT) into the professional translator’s work flow. With translation memories (TM) firmly established as a productivity tool in the translation industry, it is a conceptually obvious extension of this paradigm to include machine translation engines as virtual TMs in the set-up.

One major obstacle to this integration is the static nature of most machine translation systems that are currently available for use in production. They cannot adapt easily to feedback from the post-editor, or integrate new data into their knowledge base on short notice. In other words, they do not learn interactively from corrections to their output. Their models and knowledge bases were originally developed and designed for a batch translation scenario, where resources are first built and then used to translate in a fully automatic fashion without further intervention. Training the model parameters is still a slow and computationally very expensive process.

This paper presents *dynamic* phrase tables as an alternative, implemented within the open-source statistical machine translation (SMT) system *Moses* (Koehn et al., 2007).¹ Rather than simply looking up pre-computed entries from a database, they construct their entries on the fly by sampling word-aligned parallel data. The underlying corpus can be amended dynamically with low latency, for example by feeding post-edited output back to the translation server. New additions to the corpus can be exploited for future translations immediately.

While the underlying mechanisms are not new (cf. Callison-Burch et al., 2005; Lopez, 2007), the work reported here eliminates two major concerns about the use of bitext sampling for phrase table entry construction on demand: translation speed and translation quality. The experimental evaluation shows that in terms of speed, the sampling phrase table clearly outperforms current implementations of the work by Zens and Ney (2007). It comes close to the translation speed achievable with the hash-based compact phrase table implementation of Junczys-Dowmunt (2012). It should be noted that if translation speed is a serious concern, it is easy to pre-compute and store or cache phrase table entries for frequently occurring phrases. In terms of translation quality, the performance of the sampling phrase table is on par with conventional phrase tables for phrase-based SMT. Among the phrase table implementations that were evaluated for this work, the sampling phrase table is the only one that allows dynamic updates to its knowledge base in real time.

2 Conventional phrase tables vs. bitext sampling

2.1 Background

Most machine translation systems used in production today follow the paradigm of phrase-based statistical machine translation (PBSMT; Koehn et al., 2003). PBSMT systems typically rely on three distinct models: a language model that judges target-language fluency of a proposed translation; a translation model that gauges the quality of the elementary translation pairs that the final translation is composed of; and a distortion model that models changes in word order between source text and translation.

The units of translation in PBSMT are contiguous sequences of words in the source text (“*phrases*”) that are translated into contiguous sequences of words on the target side. Producing the translation hypothesis left-to-right in the target language, the translation algorithm selects non-overlapping phrases in arbitrary order from the source and concatenates the corresponding translations (i.e., target phrases) to produce a translation hypothesis. Jumps between the source phrases are modelled by the distortion model.

Translation options for source phrases are conventionally stored in a pre-computed table, which is called the phrase table. Phrase translation scores are computed via a (log-)linear model over a number of feature values associated with the phrase pair $\langle s, t \rangle$ in question. In the typical set-up, phrase table entries are evaluated by four feature

¹ The code has been added to the *Moses* master branch at <https://github.com/moses-smt/mosesdecoder>.

functions. In the formulas below, $\mathcal{A}_{s,t}$ is the phrase-internal word alignment between s and t . The four feature functions are as follows.

- the conditional phrase-level ‘forward’ translation probability $p(t | s)$
- the conditional phrase-level ‘backward’ translation probability $p(s | t)$
- the joint ‘lexical forward’ probability of all target words, given the source phrase (and possibly a word alignment between the two phrases): $\prod_{k=0}^{|\mathbf{t}|} p(t_k | s, \mathcal{A}_{s,t})$.
- the corresponding joint ‘lexical backward’ probability $\prod_{k=0}^{|\mathbf{s}|} p(s_k | t, \mathcal{A}_{s,t})$.

In order to achieve better translations, phrase-level probabilities are typically smoothed by Good-Turing or Kneser-Ney smoothing (Foster et al., 2006). The underlying counts and smoothing parameters are computed based on a complete list of phrase pairs extracted from the word-aligned parallel training corpus.

2.2 Bitext sampling

Except for toy examples, pre-computed phrase tables are typically very large, with the exact size of course depending on the maximum phrase length chosen and the size of the underlying corpus. The phrase table used for the timing experiments reported in Section 3.2, for example, consists of over 90 million distinct pairs of phrases of up to 7 words extracted from a moderately sized parallel corpus of fewer than 2 million parallel sentences of German-English text.

The large sizes of phrase tables make it impractical to fully load them into memory at translation time. Fully loaded into memory in the *Moses* decoder, the phrase table of the aforementioned system requires well over 100 GB of RAM and takes far beyond an hour to load. Therefore, phrase tables are usually converted to a disk-based representation, with phrase table entries retrieved from disk when needed. There are several such representations (Zens and Ney, 2007; Germann et al., 2009; Junczys-Dowmunt, 2012), two of which (Zens and Ney, 2007; Junczys-Dowmunt, 2012) have been integrated into the *Moses* system.

As an alternative to pre-computed phrase tables, Callison-Burch et al. (2005) suggested to compute phrase table entries on the fly at runtime by extracting and scoring a sample of source phrase occurrences and their corresponding translations from a pre-indexed bitext. For indexing, they use *suffix arrays* (Manber and Myers, 1990). A suffix array is an array of all token positions in a given linear sequence of tokens (e.g., a text or a DNA sequence), sorted in lexicographic order of the sub-sequence of tokens starting at the respective position. The use of suffix-array-based bitext sampling in the context of MT has been explored at length by Lopez (2007) as well as Schwartz and Callison-Burch (2010), especially with respect to Hierarchical Phrase-based Translation (HPBSMT; Chiang, 2005, 2007).

A great advantage of the suffix-array-based approach is that it is relatively cheap and easy to augment the underlying corpus. To add a pair of sentences to the parallel corpus, all we need to do is to construct a suffix array for the added material ($O(n \log n)$, where n is the number of tokens in the added material), and then merge-sort the original suffix array (of length m) with the new suffix array ($O(n + m)$).

While corpus sampling is common practice in other branches of MT research (especially HPBSMT, due to the prohibitive size of pre-computed, general-purpose, wide-coverage rule bases), adoption in the PBSMT community has been slow, apparently² due to concerns about translation speed and quality.

In the following, I intend to dispel these concerns by presenting experimental results obtained with an implementation of suffix-array-based phrase tables that sample the underlying bitext at run time, yet outperform existing disk-based implementations of conventional phrase tables by a wide margin in terms of speed (despite the greater computational effort), without any loss in translation quality.

Much of the speed benefit is related to RAM vs. disk access. Word-aligned parallel corpora are much more compact than fully expanded phrase tables, so we can afford to keep more of the information in memory, benefiting from access times that can be several orders of magnitude faster than random access to data stored on disk (Jacobs, 2009).

Moreover, the data structures are designed to be mapped directly into memory, so that we can rely on the system's virtual memory manager to transfer the data efficiently into memory when needed. This is much faster than regular file access. Two of the four implementations evaluated here store all the data on disk by default and load them on demand (PhraseDictionaryBinary, PhraseDictionaryOnDisk); the other two (PhraseDictionaryCompact and PhraseDictionaryBitextSampling (this work)) use memory-mapped files to ensure the fastest transfer possible between disk and memory. I attribute most of the speed benefits to these implementational choices (see also Sec. 3.2).

Last but not least, one can alleviate the impact of the computational overhead on overall translation time by caching frequently occurring entries, so that they must be computed only once, and perform phrase table look-up in parallel for all source phrases in a sentence submitted for translation, subject to the number of CPUs available.

The issue of translation quality is less obvious. Despite common misconceptions, it is not so much a matter of missing translation options due to sampling the bitext instead of taking into account every single source phrase occurrence. The vast majority of phrases occur so rarely that we can easily investigate every single occurrence. More frequent words and phrases will often be contained in longer, rarer phrases whose instances we also fully explore. And if there is a rare translation of a very frequent word that escapes our sampling, it is highly unlikely that this translation would survive the

²I base this statement on numerous conversations with practitioners in the field.

system’s hypothesis ranking process.

On the contrary, it is the rarity of most phrases that causes problems, as maximum likelihood estimates based on low counts are less reliable — they tend to over-estimate the true translation probability. As Foster et al. (2006) have shown, smoothing phrase-level conditional phrase probabilities improves translation performance. My experiments confirm this finding (Table 2).

Both standard methods for smoothing phrase-level translation probabilities in the phrase table, Good-Turing and Kneser-Ney, require global information about the entire set of phrasal translation relations contained in the parallel corpus. This information is not available when we sample. To take the amount of evidence available into account when estimating phrase translation probabilities, we therefore compute the lower bound of the confidence interval³ over the true translation probability, at some confidence level α , based on the observed counts. The more evidence is available, the narrower the confidence interval.

Another issue is the computation of the useful backward phrase-level translation probabilities $p(\textit{source phrase} | \textit{target phrase})$. Omitting this feature function seriously hurts performance (see Line 5 in Table 2). One could, of course, perform a full reverse look-up for each translation candidate to obtain the inverse translation probability. This would increase the number of full phrase look-ups operations necessary to construct a phrase table entry from scratch by a factor equal to the number of translation options considered for each source phrase (although again, these look-up operations could be cached). In practice, this is not necessary. To determine the denominator for the backward phrase-level translation probability, we simply scale the number of occurrences of each translation candidate in the bitext by the ratio of the source phrase sample size to the total number of source phrase occurrences in the corpus. Retrieving the total number of occurrences of the translation candidate in the corpus is trivial if we also index the target side of the corpus with a suffix array: we only need to measure the distance between the first and the occurrence of the phrase in the suffix array. Since the suffix array is sorted in lexicographic order of the corresponding suffixes, this distance is the total number of phrase occurrences.

3 Experiments

Two sets of experiments were conducted to compare bitext sampling to conventional phrase tables in terms of static performance (without updates), and a third one to assess the benefits of dynamically updating the phrase table as interactive translation progresses. The first experiment aimed at determining the quality of translation achievable with bitext sampling and the best parameter settings; the second focused on translation speed and resource requirements. Training, tuning and test data for these two experiments were taken from the data sets for the WMT 2014 shared translation task (cf. Table 1). The language model was a standard 5-gram model with Kneser-Ney smooth-

³ Specifically, the Clopper-Pearson interval (Clopper and Pearson, 1934) as implemented in the Boost C++ library.

Table 1: Corpus statistics for the training, development and test data. All corpora were part of the official data for the shared translation task at WMT 2014 and true-cased for processing.

	corpus	# of sentences	# of tokens	
			German	English
LM train	Europarl-v7	2,218,201		60,502,373
	News-Commentary-v9	304,174		7,676,138
TM train	Europarl-v7	1,920,209	50,960,730	53,586,045
	News-Commentary-v9	201,288	5,168,511	5,151,459
	total after alignment ^a	2,084,594	53,863,321	56,351,895
Tuning	Newstest-2013	3,000	64,251	65,602
Testing	Newstest-2014	3003	64,498	68,940

^a Some sentence pairs were discarded during word alignment

ing; the distortion model was a simple distance-based model without lexicalisation. The phrase table limit (i.e., the limit on the number of distinct translation hypotheses that will be considered during translation) was set to 20; the distortion limit to 6. Sampling was performed without replacement.

3.1 Translation Quality

Table 2 shows the the quality of translation achieved by the various system configurations, as measured by the BLEU score Papineni et al. (2002). The system configurations were identical except for the method used for construction and scoring of phrase table entries.

Each system was tuned 10 times in independent tuning runs to gauge the influence of parameter initialisation on overall performance (cf. also Clark et al., 2011). The 95% confidence interval in the second-but-last column was computed with bootstrap resampling for the median system within the respective group.

The first four systems rely on conventional phrase tables with four feature functions as described in Sec. 2.1: forward and backward phrase-level conditional probabilities as well as forward and backward joint lexical translation probabilities. They differ in the smoothing method used, except for the system in Line 3, which shows that filtering the phrase table to include only the top 100 entries (according to the forward phrase-level probability $p(\mathbf{t} | \mathbf{s})$) has no effect on translation quality.

Lines 5 and below are based on bitext sampling. The poor performance in Line 5 illustrates the importance of the phrase-level backward probability. Without it, the performance suffers significantly. Lines 4 and 6 show the benefits of smoothing.

The parameter α in Lines 7 to 9 is the confidence level for which the Clopper-Pearson interval was computed. Notice the minuscule difference between lines 2/3

Table 2: BLEU scores with different phrase score computation methods.

#	method	low	high	median	mean	95% conf. interval ^a	runs
1	precomp., Kneser-Ney smoothing	18.36	18.50	18.45	18.43	17.93 – 18.95	10
2	precomp., Good-Turing smoothing	18.29	18.63	18.54	18.52	18.05 – 19.05	10
3	precomp., Good-Turing smoothing, filtered^b	18.43	18.61	18.53	18.53	18.04 – 19.08	10
4	precomp., no smoothing	17.86	18.12	18.07	18.05	17.58 – 18.61	10
5	max. 1000 smpl., no smoothing, no bwd. prob.	16.70	16.92	16.84	16.79	16.35 – 17.32	10
6	max. 1000 smpl., no smoothing, with bwd. prob.	17.61	17.72	17.69	17.68	17.14 – 18.22	8
7	max. 1000 smpl., $\alpha = .05$, with bwd. prob.^c	18.35	18.43	18.38	18.38	17.86 – 18.90	10
8	max. 1000 smpl., $\alpha = .01$, with bwd. prob.	18.43	18.65	18.53	18.52	18.03 – 19.12	10
9	max. 100 smpl., $\alpha = .01$, with bwd. prob.	18.40	18.55	18.46	18.46	17.94 – 19.00	10

^a Confidence intervals were computed via bootstrap resampling for the median system in the group.

^b Top 100 entries per source phrase selected according to $p(\mathbf{t} | \mathbf{s})$.

^c The parameter α is the one-sided confidence level of the Clopper-Pearson interval for the observed counts.

and 8! By replacing plain maximum likelihood estimates with the lower bound of the confidence interval over the respective underlying translation probability, we can make up for the lack of global information necessary for Good-Turing or Kneser-Ney smoothing.

3.2 Speed

Table 3 shows average translation times⁴ per sentence for four phrase table implementations in the *Moses* system. `PhraseDictionaryBinary` and `PhraseDictionaryOnDisk` are implementations of the method described in Zens and Ney (2007). `PhraseDictionaryCompact` (Junczys-Dowmunt, 2012) is a compressed phrase table that relies on a perfect minimum hash for look-up. `PhraseDictionaryBitextSampling` is the suffix array-based phrase table presented in this paper. Each system was run with 8 threads as the only processes on an 8-core machine with locally mounted disks, translating 3003 sentences from the WMT 2014 test set. Prior to each run, all file system caches in RAM were dropped.

When the pre-computed phrase tables are not filtered, the bitext sampler outperforms even the hash-based phrase table of Junczys-Dowmunt (2012). This is due to the cost of ranking very long lists of translation candidates for very frequent source phrases. Filtering the phrase table off-line to include only the 100 most likely translation candidates for each phrase (based on $p(\mathbf{t} | \mathbf{s})$) leads to a significant speed-up without impact on translation quality (cf. Line 3 in Table 2).⁵ Similarly, the speed of the bitext sampler

⁴ The times shown were computed by dividing the total wall time of the system run by the number of sentences translated. Translations were performed in 8 parallel threads, so that the average actual translation time for a single sentence is about 8 times the time shown. Since the bitext sampler is inherently multi-threaded, the fairest form of comparison was to run the systems in a way that exhausts the host computer’s CPU capacity.

⁵ I thank M. Junczys-Dowmunt for pointing out to me that phrase tables must be filtered for optimal performance.

Table 3: Translation speed (wall time) with different phrase table implementations. The implementation names correspond to *Moses* configuration options. Translations were performed in multi-threaded mode with 8 parallel threads.

type	implementation	ave. sec./snt
static	PhraseDictionaryBinary (Zens and Ney, 2007)	0.879
static	PhraseDictionaryOnDisk (Zens and Ney, 2007)	0.717
static	PhraseDictionaryCompact (Junczys-Dowmunt, 2012)	0.366
static	PhraseDictionaryCompact (Junczys-Dowmunt, 2012), filtered ^a	0.214
dynamic	PhraseDictionaryBitextSampling, max. 1000 samples (this work)	0.256
dynamic	PhraseDictionaryBitextSampling, max. 100 samples (this work)	0.228

^a max 100 entries per source phrase

can be improved by reducing the maximum number of samples considered, although this slightly (but not significantly) reduces translation quality as measured by BLEU (cf. Line 9 in Table 2). Phrase table filtering has no impact on the speed of the other phrase table implementations.

3.3 Simulated Post-editing

The main goal of this work was to develop a phrase table that can incorporate user edits of raw machine translation output into its knowledge base at runtime. Since experiments involving real humans in the loop are expensive to conduct, I simulated the process by translating sentences from an earlier post-editing field trial in English-to-Italian translation in the legal domain. The training corpus consisted of ca. 2.5 million sentence pairs (English: ca. 44.6 million tokens, Italian: ca. 45.9 million). Due to the nature of such studies, the amount of data available for tuning and testing was fairly small: 564 sentence pairs with 17,869 English and 18,528 Italian tokens for tuning, and 472 segments with 10,829 tokens of English source text and 11,595 tokens of post-edited translation into Italian.

Several feature functions were added for use with dynamic updates to the underlying bitext. In the following, “background data” means parallel data available prior to the translation of the first sentence, and “foreground data” the parallel data that is successively added to the parallel corpus.

- Separate vs. pooled phrase-level conditional translation probabilities (forward and backward), i.e. the use of distinct feature functions for these probability estimates based on counts obtained *separately* from the background and the foreground corpus separately, or feature functions based on *pooled* counts from two corpora. Because of the small size of our tuning and test sets, counts were pooled in the experiments for this work.
- A provenance feature $\frac{n}{x+n}$, where n is the number of occurrences in the corpus

Table 4: Simulated post-editing vs. batch translation for English-to-Italian translation in the legal domain. For simulated post-editing, counts were pooled.

method	low	high	median	mean	95% conf. interval ^a	runs
conventional, Good-Turing smoothing	29.97	30.93	30.74	30.67	29.16 – 32.37	10
sampled, no updates, no smoothing, rarity pen.	29.84	30.97	30.52	30.43	28.97 – 32.25	10
simulated post-editing, pooled counts, no smoothing, rarity, provenance	30.63	33.05	31.96	31.88	30.19 – 33.77	10

^aConfidence intervals were computed via bootstrap resampling for the median system in the group.

and $x > 1$ an adjustable parameter that determines the slope of the provenance reward. The purpose of this feature is to boost the score of phrase pairs that occur in the foreground corpus.

- A global rarity penalty $\frac{x}{x+n}$ (where x and n mean the same as above) that can penalise phrase pairs that co-occur only rarely overall.

Results are shown in Table 4. None of the differences are statistically significant. In light of the small size of the test set, this is hardly surprising. In general, we should expect the benefit of adding post-edited data immediately to the knowledge base of the SMT system to vary widely depending on the repetitiveness of the source text, and on how well the translation domain is already covered by the background corpus.

4 Related Work

User-adaptive MT has received considerable research interest in recent years. Due to space limitations, we can only briefly mention a few closely related efforts here. A survey of recent work can be found, for example, in the recent journal article by Bertoldi et al. (2014b). Ortiz-Martínez et al. (2010), Bertoldi et al. (2014b), and Denkowski et al. (2014) all present systems that can be updated incrementally. Ortiz-Martínez et al. (2010) present a system that can be trained incrementally from scratch with translations that are produced in an interactive computer-aided translation scenario. The work by Bertoldi et al. (2014b) relies on cache-based models that keep track of how recently phrase pairs in the translation model and n -grams in the language models have been used in the translation pipeline and give higher scores to recently used items. They also augment the phrase table with entries extracted from post-edited translations. The work by Denkowski et al. (2014) is the closest to the work presented in this paper.⁶ Working with the *cdec* decoder (Dyer et al., 2010), they also use suffix arrays to construct phrase table entries on demand. In addition, they provide mechanisms to update the language model and re-tune the system parameters.

Focusing on dynamic adjustment of system parameters (feature function values and combination weights), Martínez-Gómez et al. (2012) investigate various online learning algorithms for this purpose. Blain et al. (2012) and Bertoldi et al. (2014a) describe

⁶Incidentally, Denkowski (personal communication) is using the implementation presented here to port the work of Denkowski et al. (2014) to the *Moses* framework.

online word alignment algorithms that can produce the word alignments necessary for phrase extraction.

5 Conclusions

I have presented a new phrase table for the *Moses* system that computes phrase table entries on the fly. It outperforms existing phrase table implementations in *Moses* in terms of speed, without sacrificing translation quality. This is accomplished by a new way of computing phrase-level conditional probabilities that takes the amount of evidence available into account and discounts probabilities whose estimates are based on little evidence. Unlike static conventional phrase tables, sampling-based phrase tables allow for rapid updates of the underlying parallel corpus and therefore lend themselves to use in an interactive and dynamic machine translation scenario.

Acknowledgements

This work was supported by the European Union’s 7th Framework Programme (FP7/2007-2013) under grant agreements 287576 (CASMAT), 287688 (MATECAT), and 288769 (ACCEPT). I thank the anonymous reviewers for numerous helpful suggestions.

References

- Bertoldi, Nicola, Amin Farajian, and Marcello Federico. 2014a. “Online word alignment for online adaptive machine translation.” *Workshop on Humans and Computer-assisted Translation*, 84–92. Gothenburg, Sweden.
- Bertoldi, Nicola, Patrick Simianer, Mauro Cettolo, Katharina Wäschle, Marcello Federico, and Stefan Riezler. 2014b. “Online adaptation to post-edits for phrase-based statistical machine translation.” *Machine Translation*. Accepted for publication. Preprint.
- Blain, Frédéric, Holger Schwenk, and Jean Senellart. 2012. “Incremental adaptation using translation information and post-editing analysis.” *9th International Workshop on Spoken Language Translation*, 229–236. Hong Kong.
- Callison-Burch, Chris, Colin Bannard, and Josh Schroeder. 2005. “Scaling phrase-based statistical machine translation to larger corpora and longer phrases.” *43rd Annual Meeting of the Association for Computational Linguistics (ACL ’05)*, 255–262. Ann Arbor, Michigan.
- Chiang, David. 2005. “A hierarchical phrase-based model for statistical machine translation.” *43rd Annual Meeting of the Association for Computational Linguistics (ACL ’05)*, 263–270. Ann Arbor, Michigan.
- Chiang, David. 2007. “Hierarchical phrase-based translation.” *Computational Linguistics*, 33(2):1–28.
- Clark, Jonathan H., Chris Dyer, Alon Lavie, and Noah A. Smith. 2011. “Better hypothesis testing for statistical machine translation: Controlling for optimizer instability.” *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies: Short Papers - Volume 2*, 176–181. Stroudsburg, PA, USA.

- Clopper, C.J. and E.S. Pearson. 1934. “The use of confidence or fiducial limits illustrated in the case of the binomial.” *Biometrika*.
- Denkowski, Michael, Chris Dyer, and Alon Lavie. 2014. “Learning from post-editing: Online model adaptation for statistical machine translation.” *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics*, 395–404. Gothenburg, Sweden.
- Dyer, Chris, Adam Lopez, Juri Ganitkevitch, Jonathan Weese, Ferhan Ture, Phil Blunsom, Hendra Setiawan, Vladimir Eidelman, and Philip Resnik. 2010. “cdec: A decoder, alignment, and learning framework for finite-state and context-free translation models.” *Proceedings of the ACL 2010 System Demonstrations*, 7–12. Uppsala, Sweden.
- Foster, George F., Roland Kuhn, and Howard Johnson. 2006. “Phrasetable smoothing for statistical machine translation.” *EMNLP*, 53–61.
- Germann, Ulrich, Eric Joanis, and Samuel Larkin. 2009. “Tightly packed tries: How to fit large models into memory, and make them load fast, too.” *Workshop on Software Engineering, Testing, and Quality Assurance for Natural Language Processing (SETQA-NLP 2009)*, 31–39. Boulder, CO, USA.
- Jacobs, Adam. 2009. “The pathologies of big data.” *Queue*, 7(6):10:10–10:19.
- Junczys-Dowmunt, Marcin. 2012. “Phrasal rank-encoding: Exploiting phrase redundancy and translational relations for phrase table compression.” *Prague Bull. Math. Linguistics*, 98:63–74.
- Koehn, Philipp, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran, Richard Zens, Chris Dyer, Ondrej Bojar, Alexandra Constantin, and Evan Herbst. 2007. “Moses: Open source toolkit for statistical machine translation.” *45th Annual Meeting of the Association for Computational Linguistics (ACL '07): Demonstration Session*. Prague, Czech Republic.
- Koehn, Philipp, Franz Josef Och, and Daniel Marcu. 2003. “Statistical phrase-based translation.” *Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics (HLT-NAACL '03)*, 48–54. Edmonton, AB, Canada.
- Lopez, Adam. 2007. “Hierarchical phrase-based translation with suffix arrays.” *EMNLP-CoNLL*, 976–985.
- Manber, Udi and Gene Myers. 1990. “Suffix arrays: A new method for on-line string searches.” *Proceedings of the First Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '90*, 319–327. Philadelphia, PA, USA.
- Martínez-Gómez, Pascual, Germán Sanchis-Trilles, and Francisco Casacuberta. 2012. “Online adaptation strategies for statistical machine translation in post-editing scenarios.” *Pattern Recognition*, 45(9):3193–3203.
- Ortiz-Martínez, Daniel, Ismael García-Varea, and Francisco Casacuberta. 2010. “Online learning for interactive statistical machine translation.” *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics, HLT '10*, 546–554. Stroudsburg, PA, USA.
- Papineni, Kishore, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. “Bleu: a method for automatic evaluation of machine translation.” *ACL 2002*, 311–318. Philadelphia, PA.

- Schwartz, Lane and Chris Callison-Burch. 2010. "Hierarchical phrase-based grammar extraction in joshua: Suffix arrays and prefix tree." *The Prague Bulletin of Mathematical Linguistics*, 93:157–166.
- Zens, Richard and Hermann Ney. 2007. "Efficient phrase-table representation for machine translation with applications to online MT and speech translation." *Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics (HLT-NAACL '07)*, 492–499. Rochester, New York.