

Why Not Grab a Free Lunch? Mining Large Corpora for Parallel Sentences to Improve Translation Modeling

Ferhan Ture

Dept. of Computer Science,
University of Maryland
fture@cs.umd.edu

Jimmy Lin

The iSchool
University of Maryland
jimmylin@umd.edu

Abstract

It is well known that translation quality increases with more training data. Unlike target language text, parallel text for translation modeling is substantially more difficult to acquire. Researchers have turned to the web to mine parallel sentences, but most previous approaches have avoided the difficult problem of pairwise similarity on cross-lingual documents and instead rely on heuristics. In our work, we confront this challenge head on with a scalable MapReduce solution. On a modest cluster, our end-to-end processing pipeline was able to automatically gather 5.8m parallel sentence pairs from English and German Wikipedia. Augmenting existing bitext with these data yielded significant improvements on translation quality over a state-of-the-art baseline (2.39 BLEU points in the best case).

1 Introduction

It has been repeatedly shown that “throwing more data at the problem” is effective in increasing SMT output quality, both in terms of translation modeling (Dyer et al., 2008) and language modeling (Brants et al., 2007). In this paper, we bring together two related research threads to gather parallel sentences for better translation modeling: cross-lingual pairwise similarity to mine comparable documents and techniques to identify sentence pairs that are mutual translations.

Unlike previous work, which sidesteps the computationally-intensive task of pairwise comparisons to mine comparable documents and instead relies on heuristics, we tackle the challenge head on.

This paper describes a scalable MapReduce-based processing pipeline that is able to automatically generate large quantities of parallel sentences, and examines the impact these data have on a state-of-the-art SMT system.

While we acknowledge that different components of this work are not novel and the general principles behind “big data” MT are well known. However, when considered together with our previous work (Ture et al., 2011), to our knowledge this is the first exposition in which all the pieces have been put together, in an end-to-end pipeline that is accessible to most academic research groups, and with open-source code for the benefit of the community.

Starting from nothing more than two corpora in different languages (in German and English, in our case), we are able to extract bitext and improve translation quality by a significant margin (2.39 BLEU points), essentially “for free”. By varying both the quantity and quality of the bitext, we can characterize the relation between data and translation quality.

2 Related Work

The idea of mining parallel sentences, particularly on the web, is of course not new. Most adopt a two step process: 1. identify comparable documents and generate candidate sentence pairs, and 2. filter candidate pairs to retain parallel sentences.

The general solution to the first step involves computing pairwise similarities across multi-lingual corpora. As this is computationally intensive, most studies fall back to heuristics, e.g., comparing news articles close in time (Munteanu and

Marcu, 2005), exploiting manual “inter-wiki” links in Wikipedia (Smith et al., 2010), or bootstrapping off an existing search engine (Resnik and Smith, 2003).

In contrast, we adopt a more exhaustive approach by solving the cross-lingual pairwise similarity problem directly, on a modest cluster. We perform experiments on German and English Wikipedia (two largest available), but our technique is general and does not depend on sparse, manually-created inter-wiki links. Thus, compared to these approaches, we achieve much higher recall.

The second step (filtering candidate sentence pairs) is relatively straightforward, and we adopt the classification approach of Munteanu and Marcu (2005). However, unlike in previous work, we need to classify large volumes of data (due to higher recall in the first step), therefore we explicitly care about the relationship between classification accuracy and the speed of the classifier. Our two-stage classification approach allows us to trade off a little accuracy for large gains in efficiency.

A more recent study presented a solution that does not depend on any existing text structure, and can scale to web collections (Uszkoreit et al., 2010). In their approach, the authors use shingles to efficiently identify similar document pairs, and then filter non-parallel sentence pairs based on an word alignment score. Despite the similarities, our approach has several additional contributions. First of all, we explore the effect of dataset size on results. Our conclusions are a bit more nuanced than simply “more data is better”, since there is a tradeoff between quality and quantity. On one hand, even with very small amounts of text added, our approach can gain solid improvements (close to 2 BLEU points of increase). On the other hand, we learn that adding more text creates noise until a certain point, after which the additional benefits start to show up.

Finally, the code and datasets are available as part of the Ivory, an open-source Hadoop toolkit for web-scale information retrieval (Lin et al., 2009).

3 Generating Candidate Sentences

We applied our approach on English Wikipedia (10.9m documents, 30.6GB) and German Wikipedia (2.4m articles, 8.5GB), using XML dumps from Jan.

2011. These were selected because they are the largest Wikipedia collections available, and we want to measure effects in a language for which we already have lots of bitext. In both collections redirect pages and stub articles were discarded.

To mine comparable documents, we used our previously described algorithm (Ture et al., 2011), based on the local-sensitive hashing technique, also implemented in MapReduce. The reader is referred to the paper for details. On a 16 node (96 core) cluster, we were able to extract 64m (d_e, d_f) document pairs (with cosine similarity ≥ 0.3) in 8.8 hours.

For each of (d_e, d_f) pairs, the next processing step involves generating the Cartesian product of sentences in both documents as candidate sentence pairs: this itself is a non-trivial problem. Although in this particular case it may be possible to load both document collections in memory, we envision scaling up to collections in the future for which this is not possible. Therefore, we devised a scalable, distributed, out-of-memory solution using Hadoop.

The algorithm works as follows: We map over (docid n , document d) pairs from both German and English collections. In each mapper all (d_e, d_f) similarity pairs are loaded in memory. If the input document is not found in any of these pairs, no work is performed. Otherwise, we extract all sentences and retain only those that have at least 5 terms and at least 3 unique terms. Sentences are converted into BM25-weighted vectors in the English term space; for German sentences, translation into English is accomplished using the technique proposed by Darwish and Oard (2003). For every (d_e, d_f) pair that the input document is found in, the mapper emits the list of weighted sentence vectors, with the (d_e, d_f) pair as the key. As all intermediate key-value pairs in MapReduce are grouped by their keys for reduce-side processing, the reducer receives the key (d_e, d_f) and weighted sentence vectors for both the German and English articles. From there, we generate the Cartesian product of sentences in both languages. As an initial filtering step, we discard all pairs where the ratio of sentence lengths is more than two, a heuristic proposed in (Munteanu and Marcu, 2005). Each of the remaining candidate sentences are then processed by two separate classifiers: a less accurate, fast classifier and a more accurate, slow classifier. This is described in the next section.

Classifier	Measure		
Simple	Recall @ P90	0.59	0.51
	Recall @ P80	0.69	0.63
	Best F-score	0.74	0.71
Complex	Recall @ P90	0.69	0.64
	Recall @ P80	0.79	0.70
	Best F-score	0.80	0.76

Table 1: Best F-score, and recall values at 80% and 90% precision, tested on Europarl data.

This algorithm is a variant of what is commonly known as a reduce-side join in MapReduce, where (d_e, d_f) serves as the join key. Note that in this algorithm, sentence vectors are emitted multiple times, one for each (d_e, d_f) pair that they participate in: this results in increased network traffic during the sort/shuffle phase. We experimented with an alternative algorithm that processes all foreign documents similar to the same English document together, e.g., processing $(d_e, [d_{f1}, d_{f2}, \dots])$ together. This approach, counter-intuitively, was actually slower despite reduced network traffic. The explanation is skew in the distribution of similar document pairs. In our experiments, half of the source collection was not linked to any target document, whereas 4% had more than 100 links. This results in reduce-side load imbalance, and while most of the reducers finish quickly, a few reducers end up performing substantially more computation, and these “stragglers” increase end-to-end running time.

4 Parallel Sentence Classification

We built two MaxEnt parallel sentence classifiers using the OpenNLP package, using data from a sample of the Europarl corpus of European parliament speeches. For training, we sampled 1000 parallel sentences from the German-English subset of the corpus as positive instances, and 5000 non-parallel sentence pairs as negative instances. For testing, we sampled another 1000 parallel pairs and generated all possible non-parallel pairs by Cartesian product from these samples. This provides a better estimate of the task we’re interested in, since most of the candidate sentence pairs will be non-parallel in a comparable corpus. We report precision, recall, and the F-score, using different classifier confidence scores as the decision threshold (see Table 1).

Our first, *simple* classifier, which uses cosine score between the sentences as the only feature, achieved a maximum F-score of 74%, with 80% precision and 69% recall. Following previous work (Smith et al., 2010), we also report recall when precision was 80% and 90% in Table 1, and these values match the results of Smith et al. (2010)’s CRF approach exactly. The second, *complex* classifier uses the following additional features: ratio of sentence lengths, number of source-side tokens that have translations on target side, ratio of target-side tokens that have translations on source side. We also experimented with features using the word alignment output, but there was no change in accuracy. The complex classifier showed better performance: recall of 79% at 80% precision and 69% at precision of 90%, with a maximum F-score of 80%.

Due to the large amounts of data involved in our experiments, we were interested in speed/accuracy tradeoffs between the two classifiers. Microbenchmarks were performed on a commodity laptop running Mac OS X on a 2.26GHz Intel Core Duo CPU, measuring per-instance classification speed (including feature computation time). The complex classifier took 100 μs per instance, about 4 times slower than the simple one, which took 27 μs .

The initial input of 64m similar document pairs translates into 400b raw candidate sentence pairs, which is first reduced to 214b by the per-sentence length filter, and then to 132b by enforcing a maximum sentence length ratio of 2. The simple classifier is applied to the remaining pairs, with different confidence thresholds. We adjusted the threshold to obtain different amounts of bitext, to see the effect on translation quality (this condition is called S_1 hereafter). The positive results of the first classifier is then processed by the second classifier (this two-level approach is called S_2 hereafter).

Candidate generation was completed in 2.4 hours on our cluster with 96 cores. These candidates went through the MapReduce shuffle-and-sort process in 0.75 hours, which were then classified in 4 hours. Processing by the more complex classifier in S_2 was an additional 0.52 hours.

5 End-to-End MT Experiments

In all experiments, our MT system learned a synchronous context-free grammar (Chiang, 2007), using GIZA++ for word alignments, MIRA for parameter tuning (Crammer et al., 2006), cdec for decoding (Dyer et al., 2010), a 5-gram SRILM for language modeling, and single-reference BLEU for evaluation. The baseline system was trained on the German-English WMT10 training data, consisting of 3.1m sentence pairs. For development and testing, we used the newswire datasets provided for WMT10, including 2525 sentences for tuning and 2489 sentences for testing.

Our baseline system includes all standard features, including phrase translation probabilities in both directions, word and arity penalties, and language model scores. It achieved a BLEU score of 21.37 on the test set, which would place it 5th out of 9 systems that reported comparable results in WMT10 (only three systems achieved a BLEU score over 22). Many of these systems used novel techniques that exploited the specific aspects of the task. In contrast, we present a knowledge-impooverished, entirely data-driven approach, by simply looking for more data in large web collections.

For both experimental conditions (one-step classification, S_1 , and two-step classification, S_2) we varied the decision threshold to generate new bitext collections of varying sizes. Each of these collections are added to the baseline training data, to train an entirely new translation model (note that GIZA additionally filters out some of the pairs based on length). The final dataset sizes, along with BLEU scores on the test data are shown in Fig. 1. In S_1 , we observe that increasing the amount of data (by lowering decision threshold) initially leads to lower BLEU scores (due to increased noise), but there is a threshold after which the improvement coming from the added data supersedes the noise. The S_2 condition increases the quality of bitext by reducing this noise: the best run, with 5.8m pairs added to the baseline (final dataset has 8.1m pairs), yields 23.76 BLEU (labeled as P on figure), 2.39 points above the baseline (and higher than the best WMT10 result). These results show that the two-step classification process, while slower, is worth the additional processing time.




Figure 1: Evaluation results on WMT10 test set

Our approach yields solid improvements even with less data added: with only 382k pairs added to the baseline, the BLEU score increases by 1.84 points. In order to better examine the effect of data size alone, we created partial datasets from P by randomly sampling sentence pairs, and then repeated experiments, also shown in Fig. 1. As expected, we see an increasing trend of BLEU scores with respect to data size. By comparing the three plots, we see that S_2 or random sampling from P work better than S_1 . Also, random sampling is not always worse than S_2 , since some pairs that receive low classifier confidence turn out to be helpful.

6 Conclusions

In this paper, we described a scalable MapReduce implementation for automatically mining parallel sentences from arbitrary comparable corpora. We show, at least for German-English MT, that an impoverished, data-driven approach is more effective than one that relies on language-specific features and tedious hand engineering. With the distributed bitext mining machinery described in this paper, improvements come basically “for free” (the only cost is a modest amount of cluster resources). We are currently applying the same technique to other language pairs, but anticipate that similar improvements will be observed as well. Given the availability of data and computing power, there is simply no reason why MT researchers should not ride the large data “tide” that lifts all boats. For the benefit of the community, we intend to open source the algorithms described in the paper and share the bitext we gathered.

References

- T Brants, A Papat, P Xu, F Och, and J Dean. 2007. Large language models in machine translation. *EMNLP*.
- D Chiang. 2007. Hierarchical phrase-based translation. *CL*, 33:201–228.
- K Crammer, O Dekel, J Keshet, S Shalev-Shwartz, and Y Singer. 2006. Online passive-aggressive algorithms. *JMLR*, 7:551–585.
- K Darwish and D Oard. 2003. Probabilistic structured query methods. *SIGIR*.
- C Dyer, A Cordova, A Mont, and J Lin. 2008. Fast, easy, and cheap: Construction of statistical machine translation models with MapReduce. *SMT Workshop*.
- C Dyer, J Weese, H Setiawan, A Lopez, F Ture, V Eidelman, J Ganitkevitch, P Blunsom, and P Resnik. 2010. cdec: a decoder, alignment, and learning framework for finite-state and context-free translation models. *ACL demos*.
- V Eidelman, C Dyer, and P Resnik. 2010. The University of Maryland statistical machine translation system for the fifth workshop on machine translation. *MATR*.
- G Hanneman, J Clark, and A Lavie. 2010. Improved features and grammar selection for syntax-based MT. *MATR*.
- Jimmy Lin, Donald Metzler, Tamer Elsayed, and Lidan Wang. 2009. Of Ivory and Smurfs: Loxodontan MapReduce experiments for web search. *TREC*.
- D Munteanu and D Marcu. 2005. Improving machine translation performance by exploiting non-parallel corpora. *CL*, 31(4):477–504.
- P Resnik and N Smith. 2003. The web as a parallel corpus. *CL*, 29(3):349–380.
- J Smith, C Quirk, and K Toutanova. 2010. Extracting parallel sentences from comparable corpora using document level alignment. *HLT/NAACL*.
- F Ture, T Elsayed, and J Lin. 2011. No free lunch: Brute force vs. locality-sensitive hashing for cross-lingual pairwise similarity. *SIGIR*.
- Jakob Uszkoreit, Jay M. Ponte, Ashok C. Papat, and Moshe Dubiner. 2010. Large scale parallel document mining for machine translation. *Proceedings of the 23rd International Conference on Computational Linguistics, COLING '10*, pages 1101–1109, Stroudsburg, PA, USA. Association for Computational Linguistics.