

# Handling Syntactic Divergence in Low-resource Machine Translation

Chunting Zhou, Xuezhe Ma, Junjie Hu, Graham Neubig

Language Technologies Institute

Carnegie Mellon University

{chuntinz, xuezhem, junjieh, gneubig}@cs.cmu.edu

## Abstract

Despite impressive empirical successes of neural machine translation (NMT) on standard benchmarks, limited parallel data impedes the application of NMT models to many language pairs. Data augmentation methods such as back-translation make it possible to use monolingual data to help alleviate these issues, but back-translation itself fails in extreme low-resource scenarios, especially for syntactically divergent languages. In this paper, we propose a simple yet effective solution, whereby target-language sentences are re-ordered to match the order of the source and used as an additional source of training-time supervision. Experiments with simulated low-resource Japanese-to-English, and real low-resource Uyghur-to-English scenarios find significant improvements over other semi-supervised alternatives<sup>1</sup>

## 1 Introduction

While neural machine translation (NMT; Bahdanau et al. (2015); Vaswani et al. (2017)) now represents the state of the art in the majority of large-scale MT benchmarks (Bojar et al., 2017), it is highly dependent on the availability of copious parallel resources; NMT under-performs previous phrase-based methods when the training data is small (Koehn and Knowles, 2017). Unfortunately, million-sentence parallel corpora are often unavailable for many language pairs. Conversely, monolingual sentences, particularly in English, are often much easier to find, making semi-supervised approaches that can use monolingual data a desirable solution to this problem.<sup>2</sup>

<sup>1</sup><https://github.com/violet-zct/pytorch-reorder-nmt>

<sup>2</sup>Unsupervised MT (Artetxe et al., 2018a) has achieved success on simulated low-resource scenarios with related lan-



Figure 1: An English sentence re-ordered into Japanese order using the rule-based method of Isozaki et al. (2010b), and its reference Japanese translation.

Semi-supervised approaches for NMT are often based on automatically creating pseudo-parallel sentences through methods such as back-translation (Irvine and Callison-Burch, 2013; Senrich et al., 2016) or adding an auxiliary auto-encoding task on monolingual data (Cheng et al., 2016; He et al., 2016; Currey et al., 2017). However, both methods have problems with low-resource and syntactically divergent language pairs. Back translation assumes enough data to create a functional NMT system, an unrealistic requirement in low-resource scenarios, while auto-encoding target sentences by definition will not be able to learn source-target word reordering.

This paper proposes a method to create pseudo-parallel sentences for NMT for language pairs with divergent syntactic structures. Prior to NMT, word reordering was a major challenge for statistical machine translation (SMT), and many techniques have emerged over the years to address this challenge (Xia and McCord, 2004; Bisazza and Federico, 2016). Importantly, even simple heuristic reordering methods with a few hand-created rules have been shown to be highly effective in closing syntactic gaps (Collins et al. (2005); Isozaki et al. (2010b); Fig. 1). Because these rules usually function solely in high-resourced languages such as English with high-quality syn-

guages, but limited success on real low-resource settings and syntactically divergent language pairs (Neubig and Hu, 2018; Guzmán et al., 2019). Hence we focus on semi-supervised methods in this paper.

tactic analysis tools, a linguist with rudimentary knowledge of the structure of the target language can create them in short order using these tools.

However, similar pre-ordering methods have not proven useful in NMT (Du and Way, 2017), largely because high-resource scenarios NMT is much more effective at learning reordering than previous SMT methods were (Bentivogli et al., 2016). However, in low-resource scenarios it is less realistic to expect that NMT could learn this reordering from scratch on its own.

Here we ask “how can we efficiently leverage the monolingual target data to improve the performance of the NMT system in low-resource, syntactically divergent language pairs?” We tackle this problem via a simple two-step data augmentation method: (1) we first reorder monolingual target sentences to create *source-ordered target sentences* as shown in Fig. 1, (2) we then replace the words in the reordered sentences with source words using a bilingual dictionary, and add them as the source side of a pseudo-parallel corpus. Experiments demonstrate the effectiveness of our approach on translation from Japanese and Uyghur to English, with a simple, linguistically motivated method of head finalization (HF; Isozaki et al. (2010b)) as our reordering method.

## 2 The Proposed Method

**Training Framework** We assume that there are two types of available resources: a small parallel corpus  $\mathcal{P} = \{(s, t)\}$  and a large monolingual target corpus  $\mathcal{Q}$ . The goal of our method is to create a pseudo-parallel corpus  $\hat{\mathcal{Q}} = \{(\hat{s}, t)\}$ , where  $\hat{s}$  is a pseudo-parallel sentence automatically created in two steps of (1) word reordering, and (2) word-by-word translation.

**Word Reordering** The first step reorders monolingual target sentences  $t \in \mathcal{Q}$  into the source order  $t_s$ . Instead of devising an entirely new word-ordering method, we can simply rely on methods that have already been widely studied and proven useful in SMT (Bisazza and Federico, 2016). Reordering can be done either using rules based on linguistic knowledge (Isozaki et al., 2010b; Collins et al., 2005) or learning from aligned parallel data (Xia and McCord, 2004; Habash, 2007), and in principle our pseudo-corpus creation paradigm is compatible with any of these methods.

Specifically, in this work we utilize rule-based methods, as our goal is to improve translation of

low-resource languages, where large quantities of high-quality parallel data do not exist and we posit that current data-driven reordering methods are unlikely to function well. Examples of rule-based methods include those to reorder English into German (Navratil et al., 2012), Arabic (Badr et al., 2009), or Japanese (Isozaki et al., 2010b). In experiments we use Isozaki et al. (2010b)’s method of reordering SVO languages (e.g. English) into the order of SOV languages (e.g. Japanese) by simply (1) applying a syntactic parser to English (Tsuruoka et al., 2004), (2) identifying the head constituent of each phrase and moving it to the end of the phrase, and (3) inserting special tokens after subjects and objects of predicates to mimic Japanese case markers.

**Word-by-word Translation** To generate data for training MT models, we next perform word-by-word translation of  $t_s$  into pseudo-source sentence  $\hat{s}$  using a bilingual dictionary (Xie et al., 2018).<sup>3</sup> There are many ways we can obtain this dictionary: even for many low-resource languages with a paucity of bilingual text, we can obtain manually-curated lexicons with reasonable coverage, or run unsupervised word alignment on whatever parallel data we have available. In addition, we can induce word translations for more words in target language using methods for bilingual lexicon induction over pre-trained word embeddings (e.g. Grave et al. (2018)).

## 3 Experiments

We evaluate our method on two language pairs: Japanese-to-English (ja-en) and Uyghur-to-English (ug-en). Japanese and Uyghur are phylogenetically distant languages, but they share similar SOV syntactic structure, which is greatly divergent from English SVO structure.

### 3.1 Experimental Setup

For both language pairs, we use an attention-based encoder-decoder NMT model with a one-layer bidirectional LSTM as the encoder and one-layer uni-directional LSTM as the decoder.<sup>4</sup> Embeddings and LSTM states were set to 300 and

<sup>3</sup>We also performed extensive preliminary experiments that learned bilingual word embeddings a-priori and froze them when training the NMT model, or continued to align the bilingual word embedding space during NMT training, but the word-by-word translation approach worked best.

<sup>4</sup>We experimented with small Transformers (Vaswani et al., 2017) but they under-performed LSTM-based models.

256 dimensions respectively. Target word embeddings are shared with the softmax weight matrix in the decoder. As noted above, we use HF (Isozaki et al., 2010b) as our re-ordering rule. HF was designed for transforming English into Japanese order, but we use it *as-is* for the Uyghur-English pair as well to demonstrate that simple, linguistically motivated rules can generalize across pairs with similar syntax with little or no modification. Further details regarding the experimental settings are in the supplementary material.

### Simulated Japanese to English Experiments

We first evaluate on a simulated low-resource ja-en translation task using the ASPEC dataset (Nakazawa et al., 2016). We randomly select 400k ja-en parallel sentence pairs to use as our full training data. We then randomly sub-sample low-resource datasets of 3k, 6k, 10k, and 20k parallel sentences, and use the remainder of the 400k English sentences as monolingual data. We duplicate the number of parallel sentences by 5 times in the training data augmented with the reordered pairs. For settings with supervised parallel sentences of 3k, 6k, 10k and 20k, we set the maximum vocabulary size of both Japanese and English to be 10k, 10k, 15k and 20k respectively.

To automatically learn a high-precision dictionary on the small amount of parallel data we have available for training, we use GIZA++ (Och and Ney, 2003) to learn alignments in both directions then take the intersection of alignments. We then learn the bilingual word embeddings with DeMaBWE (Zhou et al., 2019), an unsupervised method that has shown strong results on syntactically divergent language pairs. We give the more reliable alignments extracted from GIZA++ high priority by querying the alignment dictionary first, then follow by querying the embedding-induced dictionary. When an English word is not within any vocabulary, we output the English word *as-is* into the pseudo-source sentence.

**Real Uyghur to English Experiments** We also consider the harder case of Uyghur, a truly low-resource language. We create test and validation sets using the test data from the DARPA LORELEI corpus (Christianson et al., 2018) which contains 2,275 sentence pairs (after filtering out noisy ones) related to incidents that happened in the Uyghur area. We hold out 300 pairs as the validation data and use the rest as the test

Model	3k	6k	10k	20k	400k	ug
sup	2.17	7.86	11.67	15.98	26.56	0.58
sup-SMT	6.36	8.70	10.68	12.11	18.62	1.46
back	2.27	5.40	13.50	16.05	–	0.42
back-SMT	8.46	10.61	12.05	13.68	–	1.37
No-Reorder	6.46	9.73	12.57	15.56	–	3.24
Reorder	<b>9.94</b>	<b>12.42</b>	<b>14.98</b>	<b>17.58</b>	–	<b>4.17</b>

Table 1: BLEU of our approach (Reorder) with different amount of parallel sentences of ja-en and ug-en translation. Baselines are supervised learning from NMT and SMT (sup and sup-SMT), supervised learning with back translation from NMT and SMT (back and back-SMT) and data augmentation with translated original English sentences (No-Reorder).




Figure 2: Comparison of BLEU score w.r.t different sentence lengths.

set. The LORELEI language pack also contains the bilingual lexicons between Uyghur and English, and thousands of in-domain English sentences. We also use a large monolingual English corpus containing sentences related to various incidents occurring all over the world collected from ReliefWeb.<sup>5</sup> To sub-select a relevant subset of this corpus, we use the cross-entropy filtering (Moore and Lewis, 2010) to select 400k that are most like the in-domain English data.

For parallel data, like many low-resource languages, we only have access to data from the Bible<sup>6</sup> and Wikipedia language links (the total number of parallel Uyghur-English Wikipedia titles is 3,088), but no other in-domain parallel data. We run GIZA++ on this parallel data to obtain an alignment dictionary. We learn the bilingual word embeddings via the supervised Geometric approach (Jawanpuria et al., 2019) on FastText (Grave et al., 2018) pre-trained Uyghur and English monolingual embeddings.

<sup>5</sup><https://reliefweb.int>

<sup>6</sup><https://bible.is>

<b>source</b>	しかし、回転速度が大きすぎると、逆向きの変形が生じる
<b>reference</b>	the too high rotation speed produces the reverse deformation
<b>supervised</b>	however, the deformation of <unk> and the deformation of <unk> is caused by the dc rate
<b>ours</b>	however, the deformation of <unk> is generated when the rotation rate is large
<b>source</b>	2 - ئاينىڭ 12 - ك ، 3 : 29 پىنۇتا 3.3 كالىپ تېفرىگېن ، بىر تېفرىسە خۇنۇڭقۇرلىقى 8000 مېتىر .
<b>reference</b>	a 3.3 magnitude earthquake with the depth of 8000 meters hit feb 12 at 3:29 urumqi time
<b>supervised</b>	2 , on february 12 - 12 of darkness , urumqi time hit a 3.3 earthquake , the earthquake hit .
<b>ours</b>	2 - on february 12 - on , urumqi time 3:29minus 3.3 magnitude earthquake hit , the earthquake under depth of 8000 meters .

Table 2: Translation examples on ja-en (reorder with 6000 supervised pairs) and ug-en (reorder) from our model and the supervised counterpart.

### 3.2 Results and Comparison

**Baselines** In Tab. 1, we compare our models with baselines including regular supervised training (sup) and back-translation (Sennrich et al., 2016) (back).<sup>7</sup> To demonstrate the effectiveness of the reordering, we also compare our method against a copy-based data-augmentation method (No-reorder) where the original English sentences  $t \in Q$  rather than the reordered ones  $t_s$  are translated via the bilingual lexicon.<sup>8</sup> For each of the above settings, we also experimented with a phrased-based statistical machine translation (SMT) system (Dyer et al., 2010). In Tab. 1, we only show the results with supervised data and back-translation for SMT, since we observed that the data augmentation method performs poorly with SMT (complete results are presented in the Appendix).

**Main Results** In Tab. 1, we observe consistent improvements on both ja-en and ug-en translation tasks against other baseline methods. First, comparing our results with the NMT models trained using the same amount of parallel data, our word reordering-based semi-supervised models consistently outperform standard NMT models by a large margin. In the case that we have no access to in-domain parallel data at all, our method can still achieve some success in ug-en translation. Second, comparing our Reorder method with the No-Reorder one, reordering English sentences

into the source language order consistently brings large performance gains, which demonstrates the importance of reordering. These results are notable given previous reports that explicit reordering is not beneficial for NMT (Du and Way, 2017). Third, for ja-en translation, when gradually decreasing the amount of parallel data, the improvements of our model over the supervised NMT models become more significant, demonstrating the effectiveness of our approach in low-resource settings. Fourth, back-translation is not very beneficial or even harmful, likely because the back-translation system trained on limited supervised data can not provide high-quality translations to train the model. Finally, we also notice that although SMT performs better than NMT with less supervised training data (3k, 6k supervised data and Uyghur), the performance gain is not as remarkable as NMT when the amount of supervised data increases. Moreover, in the case of less supervised data, our data augmentation method with reordering still outperforms SMT.




Figure 3: Comparison of RIBES score on Ja-En translation with different amounts of supervised data.

<sup>7</sup>We also trained unsupervised NMT (Artetxe et al., 2018b) on 400k Japanese and English sentences from the ASPEC corpus or monolingual Uyghur and English corpus provided in the LORELEI data package, and it achieved BLEU of 0.6 and 0.0 respectively, corroborating previous results noting the difficulty of unsupervised NMT for low-resource scenarios (Neubig and Hu, 2018; Guzmán et al., 2019)

<sup>8</sup>This is similar to Currey et al. (2017) who simply copy target-language sentences to the source. But in our experiments, the source and target languages do not share the alphabet and thus we found translation necessary. This also increases consistency with our “Reorder” experiments.

We give two examples of the translation outputs of our model and a supervised NMT model for ug-en and ja-en (trained with 6k supervised pairs) in Tab. 2. In the first example from ja-en, our model is able to output terminology such as “rotation rate” thanks to the enlarged vocabulary while the supervised model can not. In the example from ug-en, our model can produce a more fluent sentence with better information coverage.

**Analysis** To investigate the effects of reordering, we compare our method with “No-Reorder” described in 3.2. First, we bucket the test data by sentence length and compute the BLEU score accordingly. We present the comparison results in Fig. 2, from which we observe that “Reorder” outperforms “No-Reorder” consistently under different sentence length buckets, and the improvement is larger when the sentence length is longer.

Second, we also evaluate the model outputs on the test data with RIBES (Isozaki et al., 2010a) which is an automatic evaluation metric of translation quality designed for distant languages that is especially sensitive to word order. From Fig. 3, we can see that “Reorder” consistently outperforms “No-Reorder” on ja-en translation especially when the amount of supervised data decreases. This suggests that with reordered pairs as the augmented training data, the model is able to output more syntactically correct sentences.

## 4 Conclusion

This paper proposed a simple yet effective semi-supervised learning framework for low-resource machine translation that artificially creates source-ordered target sentences for data-augmentation. Experimental results on ja-en and ug-en translations show that our approach achieves significant improvements over baseline systems, demonstrating the effectiveness of the proposed approach on divergent language pairs.

## Acknowledgments

This work is sponsored by Defense Advanced Research Projects Agency Information Innovation Office (I2O), Program: Low Resource Languages for Emergent Incidents (LORELEI), issued by DARPA/I2O under Contract No. HR0011-15-C-0114. The authors would like to thank Shruti Rijhwani and Hiroaki Hayashi for their help when preparing the data sets.

## References

- Mikel Artetxe, Gorka Labaka, and Eneko Agirre. 2018a. Unsupervised statistical machine translation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, Brussels, Belgium. Association for Computational Linguistics.
- Mikel Artetxe, Gorka Labaka, Eneko Agirre, and Kyunghyun Cho. 2018b. Unsupervised neural machine translation. In *International Conference on Learning Representations (ICLR)*.
- Ibrahim Badr, Rabih Zbib, and James Glass. 2009. Syntactic phrase reordering for English-to-Arabic statistical machine translation. In *Proceedings of the 12th Conference of the European Chapter of the Association for Computational Linguistics*, pages 86–93. Association for Computational Linguistics.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *International Conference on Learning Representations (ICLR)*.
- Luisa Bentivogli, Arianna Bisazza, Mauro Cettolo, and Marcello Federico. 2016. Neural versus phrase-based machine translation quality: a case study. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 257–267, Austin, Texas. Association for Computational Linguistics.
- Arianna Bisazza and Marcello Federico. 2016. A survey of word reordering in statistical machine translation: Computational models and language phenomena. *Computational linguistics*, 42(2):163–205.
- Ondřej Bojar, Rajen Chatterjee, Christian Federmann, Yvette Graham, Barry Haddow, Shujian Huang, Matthias Huck, Philipp Koehn, Qun Liu, Varvara Logacheva, et al. 2017. Findings of the 2017 conference on machine translation (WMT17). In *Proceedings of the Second Conference on Machine Translation*, pages 169–214.
- Yong Cheng, Wei Xu, Zhongjun He, Wei He, Hua Wu, Maosong Sun, and Yang Liu. 2016. Semi-supervised learning for neural machine translation. *The 54th Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Caitlin Christianson, Jason Duncan, and Boyan Onyshkevych. 2018. [Overview of the DARPA LORELEI Program](#). *Machine Translation*, 32(1-2):3–9.
- Michael Collins, Philipp Koehn, and Ivona Kučerová. 2005. Clause restructuring for statistical machine translation. In *Proceedings of the 43rd annual meeting on association for computational linguistics*, pages 531–540. Association for Computational Linguistics.

- Anna Currey, Antonio Valerio Miceli Barone, and Kenneth Heafield. 2017. Copied monolingual data improves low-resource neural machine translation. In *Proceedings of the Second Conference on Machine Translation*, pages 148–156.
- Jinhua Du and Andy Way. 2017. Pre-Reordering for Neural Machine Translation: Helpful or Harmful? *The Prague Bulletin of Mathematical Linguistics*, 108(1):171–182.
- Chris Dyer, Jonathan Weese, Hendra Setiawan, Adam Lopez, Ferhan Ture, Vladimir Eidelman, Juri Ganitkevitch, Phil Blunsom, and Philip Resnik. 2010. cdec: A decoder, alignment, and learning framework for finite-state and context-free translation models. In *Proceedings of the ACL 2010 System Demonstrations*, pages 7–12. Association for Computational Linguistics.
- Edouard Grave, Piotr Bojanowski, Prakhar Gupta, Armand Joulin, and Tomas Mikolov. 2018. [Learning word vectors for 157 languages](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC-2018)*.
- Francisco Guzmán, Peng-Jen Chen, Myle Ott, Juan Pino, Guillaume Lample, Philipp Koehn, Vishrav Chaudhary, and Marc’Aurelio Ranzato. 2019. Two New Evaluation Datasets for Low-Resource Machine Translation: Nepali-English and Sinhala-English. *arXiv preprint arXiv:1902.01382*.
- Nizar Habash. 2007. Syntactic preprocessing for statistical machine translation. *Proceedings of the 11th MT Summit*, 10.
- Di He, Yingce Xia, Tao Qin, Liwei Wang, Nenghai Yu, Tiejun Liu, and Wei-Ying Ma. 2016. Dual learning for machine translation. In *Advances in Neural Information Processing Systems (NIPS)*, pages 820–828.
- Ann Irvine and Chris Callison-Burch. 2013. Combining bilingual and comparable corpora for low resource machine translation. In *Proceedings of the eighth workshop on statistical machine translation*, pages 262–270.
- Hideki Isozaki, Tsutomu Hirao, Kevin Duh, Katsuhito Sudoh, and Hajime Tsukada. 2010a. Automatic evaluation of translation quality for distant language pairs. In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*, pages 944–952. Association for Computational Linguistics.
- Hideki Isozaki, Katsuhito Sudoh, Hajime Tsukada, and Kevin Duh. 2010b. Head finalization: A simple reordering rule for SOV languages. In *Proceedings of the Joint Fifth Workshop on Statistical Machine Translation and MetricsMATR*, pages 244–251. Association for Computational Linguistics.
- Pratik Jawanpuria, Arjun Balgovid, Anoop Kunchukuttan, and Bamdev Mishra. 2019. Learning multilingual word embeddings in latent metric space: a geometric approach. *Transactions of the Association for Computational Linguistics (TACL)*.
- Philipp Koehn and Rebecca Knowles. 2017. Six challenges for neural machine translation. In *Proceedings of the First Workshop on Neural Machine Translation*, pages 28–39, Vancouver. Association for Computational Linguistics.
- Robert C Moore and William Lewis. 2010. Intelligent selection of language model training data. In *Proceedings of the ACL 2010 conference short papers*, pages 220–224. Association for Computational Linguistics.
- Toshiaki Nakazawa, Manabu Yaguchi, Kiyotaka Uchimoto, Masao Utiyama, Eiichiro Sumita, Sadao Kurohashi, and Hitoshi Isahara. 2016. ASPEC: Asian Scientific Paper Excerpt Corpus. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC 2016)*, pages 2204–2208, Portorož, Slovenia.
- Jiri Navratil, Karthik Visweswariah, and Ananthkrishnan Ramanathan. 2012. A comparison of syntactic reordering methods for English-German machine translation. *Proceedings of COLING 2012*, pages 2043–2058.
- Graham Neubig and Junjie Hu. 2018. Rapid adaptation of neural machine translation to new languages. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 875–880. Association for Computational Linguistics.
- Franz Josef Och and Hermann Ney. 2003. A systematic comparison of various statistical alignment models. *Computational Linguistics*, 29(1):19–51.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Improving Neural Machine Translation Models with Monolingual Data. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 86–96.
- Yoshimasa Tsuruoka, Yusuke Miyao, and Jun’ichi Tsujii. 2004. Towards efficient probabilistic HPSG parsing: integrating semantic and syntactic preference to guide the parsing. In *Proceedings of the IJCNLP-04 Workshop on Beyond Shallow Analyses*. ACL.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008.
- Fei Xia and Michael McCord. 2004. Improving a Statistical MT System with Automatically Learned Rewrite Patterns. In *Proceedings of Coling 2004*, pages 508–514, Geneva, Switzerland. COLING.

Jiateng Xie, Zhilin Yang, Graham Neubig, Noah A Smith, and Jaime Carbonell. 2018. Neural cross-lingual named entity recognition with minimal resources. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 369–379.

Chunting Zhou, Xuezhe Ma, Di Wang, and Graham Neubig. 2019. Density Matching for Bilingual Word Embedding. In *Meeting of the North American Chapter of the Association for Computational Linguistics (NAACL)*, Minneapolis, USA.