

The LMU Munich System for the WMT 2020 Unsupervised Machine Translation Shared Task

Alexandra Chronopoulou, Dario Stojanovski, Viktor Hangya, Alexander Fraser

Center for Information and Language Processing, LMU Munich, Germany
{achron, stojanovski, hangyav, fraser}@cis.lmu.de

Abstract

This paper describes the submission of LMU Munich to the WMT 2020 unsupervised shared task, in two language directions, German↔Upper Sorbian. Our core unsupervised neural machine translation (UNMT) system follows the strategy of Chronopoulou et al. (2020), using a monolingual pretrained language generation model (on German) and fine-tuning it on both German and Upper Sorbian, before initializing a UNMT model, which is trained with online backtranslation. Pseudo-parallel data obtained from an unsupervised statistical machine translation (USMT) system is used to fine-tune the UNMT model. We also apply BPE-Dropout to the low-resource (Upper Sorbian) data to obtain a more robust system. We additionally experiment with residual adapters and find them useful in the Upper Sorbian→German direction. We explore sampling during backtranslation and curriculum learning to use SMT translations in a more principled way. Finally, we ensemble our best-performing systems and reach a BLEU score of 32.4 on German→Upper Sorbian and 35.2 on Upper Sorbian→German.

1 Introduction

Neural machine translation achieves remarkable results (Bahdanau et al., 2015; Vaswani et al., 2017) when large parallel training corpora are available. However, such corpora are only available for a limited number of languages. UNMT addresses this issue by using monolingual data only (Artetxe et al., 2018c; Lample et al., 2018). The performance of UNMT models is further improved using transfer learning from a pretrained cross-lingual model (Lample and Conneau, 2019; Song et al., 2019). However, pretraining also demands large monolingual corpora for both languages. Without abundant data, UNMT methods are often ineffective (Guzmán et al., 2019). Therefore, effectively trans-

lating between a high-resource and a low-resource language, in terms of monolingual data, which is the target of this year’s unsupervised shared task, is challenging.

We participate in the WMT 2020 unsupervised machine translation shared task. The task includes two directions: German→Upper Sorbian (De→Hsb) and Upper Sorbian→German (Hsb→De). Our systems are constrained, using only the provided Hsb monolingual data and De NewsCrawl monolingual data released for WMT. We pretrain a monolingual encoder-decoder model on a language generation task with the Masked Sequence to Sequence model (MASS) (Song et al., 2019) and fine-tune it on both languages of interest, following Chronopoulou et al. (2020). We then train it on UNMT, using online backtranslation. We use our USMT system to backtranslate monolingual data in both languages. This pseudo-parallel corpus serves to fine-tune our UNMT model. Iterative offline backtranslation is later leveraged, yielding a performance boost. We use BPE-Dropout (Provilkov et al., 2020) as a data augmentation technique, sampling instead of greedy decoding in online backtranslation, and curriculum learning to best include the SMT pseudo-parallel data. We also use residual adapters (Houlsby et al., 2019) to translate to the low-resource language (Hsb).

Results Summary. The ensemble of our best-performing systems yields the best performance in terms of BLEU¹ among the participants of the unsupervised machine translation shared task. We release the code and our best models² in order to facilitate reproduction of our work and experimentation in this field. We note that we have built upon

¹http://matrix.statmt.org/matrix/systems_list/1920

²<https://github.com/alexandra-chron/umt-lmu-wmt2020>

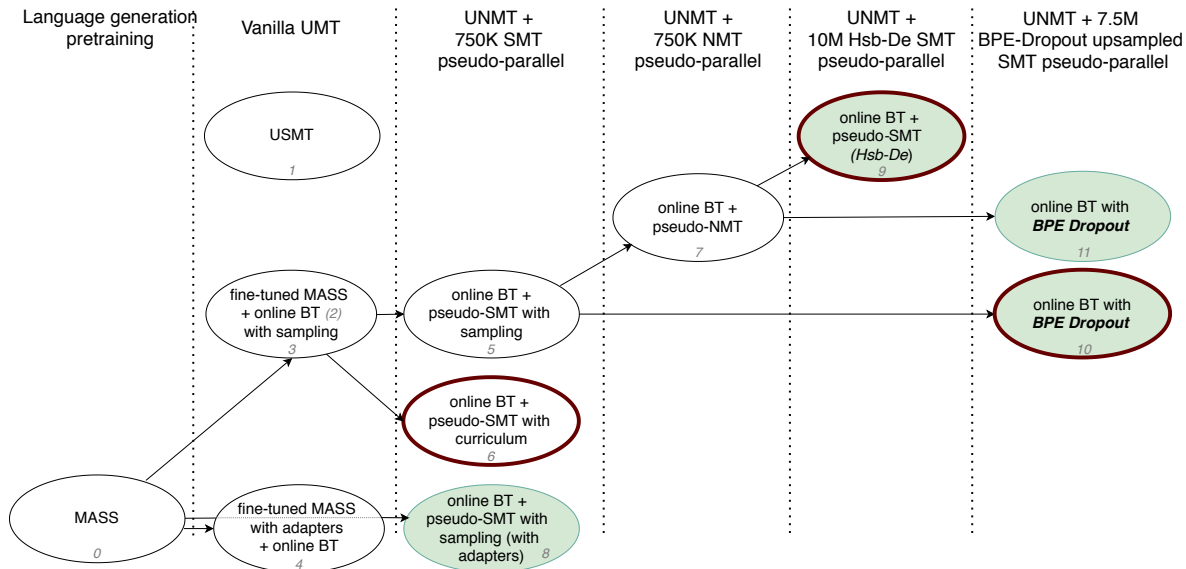


Figure 1: Illustration of our system. We denote with green the systems that were ensembled for the De→Hsb direction and with maroon the systems that were ensembled for the Hsb→De direction. Right arrows indicate transfer of weights. The numbers in gray correspond to the rows of Table 1. Online BT refers to the backtranslation of sentences with the actual model and updating it with the generated pseudo-parallel data. Pseudo-SMT refers to data obtained by backtranslating using the USMT baseline system while pseudo-NMT to our translations using system 5. The components of our approach are explained in Section 2.

the MASS codebase³ for our experiments.

2 Model Description

Figure 1 presents all the different components of our system and how they are connected to each other. We train both an unsupervised SMT (#1) and NMT (#2) model. The UNMT model is based on a pretrained MASS model (#0), which is *monolingual* (De). The model is later fine-tuned on both Hsb and De. We additionally explore fine-tuning only on Hsb using adapters. These models are used to initialize an NMT model (#2, #4) which is trained with online backtranslation. We additionally experiment with sampling (#3) during backtranslation. The USMT model is used to backtranslate Hsb and De data. This synthetic bi-text is used to fine-tune the baseline UNMT model (#5). We use the synthetic bi-text also to fine-tune directly the adapter-augmented MASS model, while employing online backtranslation and sampling (#8). We experiment with curriculum learning (#6) to estimate the optimal way to feed the model this pseudo-parallel data. We also use our UNMT model to generate backtranslations and fine-tune existing models (#7). Further USMT-backtranslated data is used in #9. Finally, some models are fine-tuned with monolingual data which is oversampled and segmented

with BPE-Dropout (#10, #11). The details of these components are outlined in the following.

2.1 Unsupervised SMT

First we describe the USMT system which we use to generate pseudo-parallel data to fine-tune our NMT system. We use *monoses* (Artetxe et al., 2018b), which builds unsupervised bilingual word embeddings (BWEs) and integrates them to Moses (Koehn et al., 2006), but apply some modifications to it.

As a first step, we build unsupervised BWEs with *fastText* (Bojanowski et al., 2017) and *VecMap* (Artetxe et al., 2018a) containing representations of 1-, 2- and 3-grams. Since the size of the available monolingual Hsb data is low, mapping monolingual embeddings to BWEs without any bilingual signal fails, i.e., we find no meaningful translations by manually investigating the most similar cross-lingual pairs of a few words. Instead, we rely on identical words occurring in both De and Hsb corpora as the initial seed dictionary. The BWEs are then converted to phrase-tables using cosine similarity of words and a language model is trained on the available monolingual data. The shared task organizers released a validation set which we use to tune the parameters of the system with MERT, instead of running unsupervised tuning as described in Artetxe et al. (2018b). Finally, we run 4 itera-

³<https://github.com/microsoft/MASS>

tive refinement steps to further improve the system. Other than the above, all steps and parameters are unchanged.

We use this system in inference mode to back-translate 7M D_e and 750K H_{sb} sentences. We refer to this pseudo-parallel dataset as `7.7M SMT pseudo-parallel`. We also backtranslate 10M more D_e sentences. This dataset is later used to fine-tune one of our systems. We refer to it as `10M Hsb-De SMT pseudo-parallel`.

2.2 MASS

We initialize our UNMT systems with an encoder-decoder Transformer (Vaswani et al., 2017), which is pretrained using the MASS (Song et al., 2019) objective. The model is pretrained by trying to reconstruct a sentence fragment given the remaining part of the sentence. The encoder takes a randomly masked fragment as input, while the decoder tries to predict the masked fragment. MASS is inspired by BERT (Devlin et al., 2019), but is more suitable for machine translation, as it pre-trains the encoder-decoder and the attention mechanism, whereas BERT is an encoder Transformer. In order to pretrain the model, instead of training MASS on both D_e and H_{sb} , we initially train it on D_e . After this, we fine-tune it on both D_e and H_{sb} , following RE-LM (Chronopoulou et al., 2020). The intuition behind this is that, if we simultaneously train a cross-lingual model on unbalanced data, where X is much larger than Y , the model starts to overfit the low-resource side Y before being trained on all the high-resource language data (X). This results in poor translations. We refer to our pretrained model as FINE-TUNED MASS.

2.2.1 Vocabulary Extension for NMT

To fine-tune the pretrained D_e MASS model on H_{sb} , we need to overcome the following issue: the pretrained model uses BPE segmentation and vocabulary based only on D_e . To this end, we again follow RE-LM. We denote these BPE tokens as BPE_{D_e} and the resulting vocabulary as V_{D_e} . We aim to fine-tune the monolingual MASS model to H_{sb} . Splitting H_{sb} with BPE_{D_e} would result in heavy segmentation of H_{sb} words. To prevent this from happening, we learn BPEs on the joint D_e and H_{sb} corpus (BPE_{joint}). We then use BPE_{joint} tokens to split the H_{sb} data, resulting in a vocabulary $V_{H_{sb}}$. This method increases the number of shared tokens and enables cross-lingual transfer of the pretrained model. The final vocabulary is the union

of the V_{D_e} and $V_{H_{sb}}$ vocabularies. We extend the input and output embedding layer to account for the new vocabulary items. The new parameters are then learned during fine-tuning.

2.3 Adapters

Besides initializing our UNMT systems with FINE-TUNED MASS, we also experiment with pretraining MASS on D_e and fine-tuning *only* on H_{sb} . During fine-tuning, we freeze the encoder and decoder Transformer layers and add adapters (Houlsby et al., 2019) to each of the Transformer layers. Adapters can prevent catastrophic forgetting (Goodfellow et al., 2013) and show promising results in various tasks (Bapna and Firat, 2019; Artetxe et al., 2020). We fine-tune only the output layer, the embeddings and the decoder’s attention to the encoder as well as the lightweight adapter layers.

We investigate adapters as fine-tuning in this way is considerably more computationally efficient. We also experimented with freezing the decoder’s attention to the encoder as well as adding an adapter on top of it, but these architecture designs are worse in terms of perplexity during MASS fine-tuning as well as BLEU scores during UNMT.

We use the fine-tuned model to initialize an encoder-decoder Transformer, augmented with adapters. The adapter-augmented model is then trained in an unsupervised way, using online back-translation. All layers are trainable during unsupervised NMT training. We refer to this model as FINE-TUNED MASS + ADAPTERS.

2.4 Unsupervised NMT (online backtranslation)

We initialize our UNMT models with FINE-TUNED MASS. Following Song et al. (2019), we train the systems in an unsupervised manner, using online backtranslation (Sennrich et al., 2016a) of the monolingual H_{sb} and D_e data, that were also used for pretraining. As proposed in Song et al. (2019), we do not use denoising auto-encoding (Vincent et al., 2008). We use online backtranslation to generate pseudo bilingual data for training. We refer to the resulting model as UNMT BASELINE.

2.5 Sampling

We experiment with sampling instead of greedy decoding during online backtranslation. Edunov et al. (2018) show that sampling is beneficial for backtranslation compared to greedy decoding or beam search for systems trained on larger amounts

of parallel data. Although we do not use any parallel data, we assumed that our initial UNMT baseline is of reasonable quality and that sampling would be beneficial. However, in order to provide a balance, we randomly use either greedy decoding or sampling during training. The frequency with which sampling is used is a hyperparameter which we set to 0.5. Sampling temperature is set to 0.95.

2.6 Curriculum learning

Considering the high improvements achieved by including SMT backtranslated data, we conduct experiments to determine a more meaningful way to feed the data to the model using curriculum learning (Kocmi and Bojar, 2017; Platanios et al., 2019; Zhang et al., 2019). We learn the curriculum using Bayesian Optimization (BO) for which we use an open source implementation⁴. Similar work has been proposed for transfer learning (Ruder and Plank, 2017) and NMT (Wang et al., 2020). As we already have a reasonably trained NMT model, we use it to compute instance-level features for learning the curriculum. Each sentence pair from the SMT backtranslated data is represented with two features: the model scores for this pair in the *original* (backtranslation \rightarrow monolingual sentence) and *reverse* direction (monolingual \rightarrow backtranslation).

The weights that determine the importance of these features are learned separately for $De \rightarrow Hsb$ and $Hsb \rightarrow De$, so that we have 4 features in total. BO runs for 30 trials. The feature weights are constrained in the range $[-1, 1]$. Each trial runs 5.4K NMT updates. The curriculum optimizes the sum of $Hsb \rightarrow De$ and $De \rightarrow Hsb$ validation perplexity. For the optimization trials, we only use the SMT backtranslated data as pseudo-parallel data and do not use online backtranslation. Finally, based on the feature weights and the features for each sentence, we sort the pseudo-parallel data and fine-tune the UNMT BASELINE with SMT backtranslations and online backtranslation. It would be interesting to study if a similar approach can be used to estimate a more optimal loading of monolingual data during MASS pretraining and UNMT.

2.7 Offline Iterative Backtranslation

We also experiment with creating synthetic training data using offline backtranslation with one of our UNMT systems (#5 in Table 1). We translate 750K De sentences to Hsb and 750K Hsb sen-

⁴<https://ax.dev/>

tences to De . The resulting pseudo-parallel system is denoted as `750K NMT pseudo-parallel corpus` and is used to fine-tune the same system.

2.8 BPE-Dropout

BPE segmentation is useful in machine translation, as it efficiently addresses the open vocabulary problem. This approach keeps the most frequent words intact and splits the rare ones into multiple tokens. It builds a vocabulary of subwords and a merge table, specifying which subwords have to be merged and the priority of the merges. BPE segmentation always splits a word deterministically. Introducing stochasticity to the algorithm (Provilkov et al., 2020), by simply removing a merge from the merges with a pre-defined probability p , results in significant BLEU improvements for various languages in low- and medium-resource datasets.

We use BPE-Dropout in the following way: we oversample the `Hsb` monolingual data by a factor of 10 and apply BPE-Dropout. In that way, we get different segmentations of the same sentences and feed this data to the model. We also oversample the `750K SMT pseudo-parallel corpus` in the same manner, but only apply BPE-Dropout to the `Hsb` side. These monolingual and pseudo-parallel oversampled datasets are used to fine-tune our models. These systems perform better than our other single systems.

2.9 Ensembling

For the final models, we perform ensemble decoding with the best training models obtained in our experiments. We evaluate several combinations of model ensembles. Based on BLEU scores on the test set provided during development, we decide on two separate ensembles for $De \rightarrow Hsb$ and $Hsb \rightarrow De$ for the final submission.

3 Experiments

3.1 Data Pre-processing

In line with the rules of the WMT 2020 unsupervised shared task⁵, we used 327M sentences from WMT monolingual News Crawl⁶ dataset for German, collected over the period of 2007 to 2019. We also used the Upper Sorbian side of the provided parallel data as well as all of the monolingual data, a total amount of 756K sentences, provided by the

⁵http://www.statmt.org/wmt20/unsup_and_very_low_res/

⁶<http://data.statmt.org/news-crawl/de/>

#	Methods	De→Hsb	Hsb→De
0	MASS	5.6	7.0
1	USMT	19.3	21.4
2	① UNMT baseline (fine-tuned MASS)	24.4	27.1
3	② UNMT baseline + sampling	25.4	27.4
4	① UNMT baseline (fine-tuned MASS with adapters)	18.8	21.7
5	③ + online BT + pseudo-SMT + sampling	29.9	31.9
6	③ + online BT + pseudo-SMT + curriculum	<u>30.0</u>	32.5
6*	③ + online BT + pseudo-SMT + curriculum + sampling	30.2	32.8
7	⑤ + online BT + pseudo-NMT	29.8	<u>33.2</u>
8	① + online BT + pseudo-SMT + sampling (with adapters)	29.0	32.3
9	⑦ + online BT + pseudo-SMT (Hsb-De)	<u>30.0</u>	32.7
Data oversampling with BPE-Dropout			
10	⑤ + BPE-Dropout	30.7	33.4
11	⑦ + BPE-Dropout	<u>31.8</u>	<u>34.0</u>
12	Model Ensemble (8, 9, 10, 11)	32.4	35.2
13	Model Ensemble (6, 9, 11)	31.9	34.8

Table 1: BLEU scores of UMT for De-Hsb and Hsb-De systems. The systems with the underlined results were ensembled and used in our primary submissions. #12 is our primary system submitted to the organizers in the De→Hsb direction, while #13 is our primary system submitted in the Hsb→De direction. 6* was trained after the shared task and is not used for the final submission.

organizers. We used the provided parallel data for validation/testing (2K/2K sentences). We normalized punctuation, tokenized and true-cased the data using standard scripts from the Moses toolkit (Koehn et al., 2006). We note that we tokenized Hsb data using Czech as the language of tokenization, since these two languages are very closely related and there are no tokenization rules for Hsb in Moses.

We used BPE (Sennrich et al., 2016b) segmentation for our neural system. Specifically, we learned 32K codes and computed the vocabulary using the De data. We then also learned the same amount of BPEs on the joint corpus (De, Hsb) and computed the joint vocabulary. We extended the initial vocabulary, adding to it unseen items. We used this augmented vocabulary to fine-tune the MASS model and run all the UNMT training experiments.

3.2 Data Post-processing

We fixed the quotes to be the same as in the source sentences (German-style). We also applied a recaser using Moses (Koehn et al., 2006) to convert the translations to mixed case.

3.3 Training

Unsupervised SMT. As mentioned before, we used *fastText* (Bojanowski et al., 2017) to build 300 dimensional embeddings on the available monolingual data. We build BWEs with *VecMap* (Artetxe et al., 2018a) using identical words as the seed

dictionary and restricting the vocabulary to the most frequent 200K, 400K and 400K 1-, 2- and 3-grams respectively. We used *monoses* (Artetxe et al., 2018b) as the USMT pipeline but used the available validation data for parameter tuning and ran 4 iterative refinement steps.

MASS. We use a Transformer, which consists of 6-layer encoder and 6-layer decoder with 1024 embedding/hidden size, 4096 feed-forward network size and 8 attention heads. We pretrain MASS on De monolingual data, using Adam (Kingma and Ba, 2015) optimizer with inverse square root learning rate scheduling and a learning rate of 10^{-4} . We used a per-GPU batch size of 32. We trained the model for approximately 2 weeks on 8 NVIDIA GTX 1080 Ti 11 GB GPUs. The rest of the hyperparameters follows the original MASS paper. We fine-tune MASS on both De and Hsb using the same setup, but on 4 GPUs of the same type. Fine-tuning was performed for 2 days.

Unsupervised NMT. For unsupervised NMT, we further train the fine-tuned MASS using online backtranslation. We use 4 GPUs to train each one of our UNMT models. We report BLEU using SacreBLEU (Post, 2018)⁷ on the provided test set.

Unsupervised NMT + Pseudo-parallel MT. We train our UNMT systems using a pseudo-parallel supervised translation loss, in addition to the online backtranslation objective. We found out that aug-

⁷BLEU+case.mixed+numrefs.1+smooth.exp+tok.13a+version.1.4.13

menting UNMT systems with pseudo-parallel data obtained by USMT leads to major improvements in translation quality, as previous work has showed (Artetxe et al., 2018b; Stojanovski et al., 2019).

4 Results

The results of our systems on the test set provided during development are presented in Table 1. Our USMT model (#1) performs competitively, but is largely outperformed by the UNMT baseline (#2). These results are interesting considering that both systems are trained using small amounts of monolingual H_{sb} data. We believe that the performance of the UNMT model is largely due to the MASS fine-tuning scheme which allowed us to obtain a strong pretrained model for both languages. We also observe (#3) that mixing greedy decoding and sampling during backtranslation is beneficial compared to always using greedy decoding (#2), especially for $De \rightarrow H_{sb}$ which improved by 1.0 BLEU. However, it is likely that sampling is useful only if the model is of reasonable quality. We note that the adapter-augmented model (#4) is worse than the UNMT baseline.

After these initial experiments, we use the USMT model (#1) to backtranslate all H_{sb} monolingual data and 7M De sentences. This pseudo-parallel data is leveraged to fine-tune our UNMT models alongside online backtranslation. This approach, denoted as model #5, improves the UNMT baseline (#3) by more than 5.5 BLEU for $De \rightarrow H_{sb}$ and 4.5 BLEU for $H_{sb} \rightarrow De$. The curriculum learning approach (#6) yields a small improvement of 0.6 BLEU for $H_{sb} \rightarrow De$. Unfortunately, the curriculum learning model ran without the use of sampling. We later train the model with sampling (#6*) and obtain slight improvements in both directions.

Using NMT backtranslations in an offline manner (#7) provides for a large improvement in the $H_{sb} \rightarrow De$ direction, obtaining 33.2 BLEU. Further training our high scoring model #7 on USMT backtranslations, depicted as model #9, degrades performance on $H_{sb} \rightarrow De$. This might indicate that USMT backtranslations alone are not very important for high performance, but simply adding any kind of pseudo-parallel data during training.

The adapter-augmented model with USMT backtranslations (#8) manages to close the gap to the baseline model. Comparing #5 and #8, we can see that the model with adapters is worse by 0.9 BLEU on $De \rightarrow H_{sb}$, but better by 0.4 on $H_{sb} \rightarrow De$. Due

to time constraints, we train #4 and #8 in parallel and #8 is not fine-tuned from #4. Overall, adapters are a promising research direction as they lead to faster MASS fine-tuning and comparable performance.

We observe considerable improvements using BPE-Dropout. As noted before, we oversample the parallel and H_{sb} monolingual data and apply BPE-Dropout only on H_{sb} . We use this data to fine-tune some of our already trained models, specifically #5 and #7 which results in models #10 and #11, respectively. This approach improves the $H_{sb} \rightarrow De$ direction by up to 1.5 BLEU and up to 1.0 BLEU for $De \rightarrow H_{sb}$. System #11 proved to be our best single system in both translation directions. We hypothesize that using BPE-Dropout while simultaneously oversampling the data provides for a data augmentation effect. In future work, it would be interesting to decouple these two steps and measure their effect separately.

Ensembling further boosts performance. Ensemble #12 is used for $De \rightarrow H_{sb}$ and #13 for $H_{sb} \rightarrow De$. We note that while computing ensemble BLEU scores during development, we did not fix the issue with German-style quotes. This resulted in ensemble #13 obtaining better scores on $H_{sb} \rightarrow De$. We later fix the quotes issue and find out that ensemble #12 is better on both translation directions and is the best system overall.

5 Conclusion

In this paper, we present the LMU Munich system for the WMT 2020 unsupervised shared task for translation between German and Upper Sorbian. Our system is a combination of an SMT and an NMT model trained in an unsupervised way. The UNMT model is trained by fine-tuning a MASS model, according to the recently proposed RE-LM approach. The experiments show that the MASS fine-tuning technique is efficient even if little monolingual data is available for one language and results in a strong UNMT model. We also show that using pseudo-parallel data from USMT and UNMT backtranslations improves performance considerably. Furthermore, we show that oversampling the low-resource Upper Sorbian and applying BPE-Dropout, which can effectively be seen as data augmentation, results in further improvements. Adapters in MASS fine-tuning provided for a balance between performance and computational efficiency. Finally, smaller but noticeable gains are obtained from us-

ing curriculum learning and sampling during decoding in backtranslation.

Acknowledgments

This work was supported by the European Research Council (ERC) under the European Unions Horizon 2020 research and innovation programme (grant agreement No. 640550) and by the German Research Foundation (DFG; grant FR 2829/4-1). We would like to thank Jindřich Libovický for fruitful discussions regarding the use of BPE-Dropout as a data augmentation technique.

References

- Mikel Artetxe, Gorka Labaka, and Eneko Agirre. 2018a. [A robust self-learning method for fully unsupervised cross-lingual mappings of word embeddings](#). In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pages 789–798.
- Mikel Artetxe, Gorka Labaka, and Eneko Agirre. 2018b. [Unsupervised statistical machine translation](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 3632–3642.
- Mikel Artetxe, Gorka Labaka, Eneko Agirre, and Kyunghyun Cho. 2018c. [Unsupervised neural machine translation](#). In *International Conference on Learning Representations*.
- Mikel Artetxe, Sebastian Ruder, and Dani Yogatama. 2020. [On the cross-lingual transferability of monolingual representations](#). In *Proceedings of the Annual Meeting of the 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*.
- Ankur Bapna and Orhan Firat. 2019. [Simple, scalable adaptation for neural machine translation](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing and the International Joint Conference on Natural Language Processing*, pages 1538–1548.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. [Enriching Word Vectors with Subword Information](#). *Transactions of the Association for Computational Linguistics*, pages 135–146.
- Alexandra Chronopoulou, Dario Stojanovski, and Alexander Fraser. 2020. [Reusing a Pretrained Language Model on Languages with Limited corpora for Unsupervised NMT](#). *arXiv preprint arXiv:2009.07610*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4171–4186.
- Sergey Edunov, Myle Ott, Michael Auli, and David Grangier. 2018. [Understanding back-translation at scale](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 489–500.
- Ian J Goodfellow, Mehdi Mirza, Da Xiao, Aaron Courville, and Yoshua Bengio. 2013. [An empirical investigation of catastrophic forgetting in gradient-based neural networks](#). *arXiv preprint arXiv:1312.6211*.
- Francisco Guzmán, Peng-Jen Chen, Myle Ott, Juan Pino, Guillaume Lample, Philipp Koehn, Vishrav Chaudhary, and MarcAurelio Ranzato. 2019. [The FLORES evaluation datasets for low-resource machine translation: Nepali–English and Sinhala–English](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing and the International Joint Conference on Natural Language Processing*, pages 6100–6113.
- Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin de Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019. [Parameter-efficient transfer learning for NLP](#). In *Proceedings of the International Conference on Machine Learning*.
- Diederick P Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *International Conference on Learning Representations*.
- Tom Kocmi and Ondřej Bojar. 2017. [Curriculum learning and minibatch bucketing in neural machine translation](#). In *Proceedings of the International Conference Recent Advances in Natural Language Processing*, pages 379–386.
- Philipp Koehn, Marcello Federico, Wade Shen, Nicola Bertoldi, Ondřej Bojar, Chris Callison-Burch, Brooke Cowan, Chris Dyer, Hieu Hoang, Richard Zens, et al. 2006. [Open source toolkit for statistical machine translation: Factored translation models and confusion network decoding](#). In *Final Report of the 2006 JHU Summer Workshop*.
- Guillaume Lample and Alexis Conneau. 2019. [Cross-lingual language model pretraining](#). In *Advances in Neural Information Processing Systems*, page 70577067.
- Guillaume Lample, Myle Ott, Alexis Conneau, Ludovic Denoyer, and Marc Aurelio Ranzato. 2018. [Phrase-based & neural unsupervised machine translation](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 5039–5049.

- Emmanouil Antonios Platanios, Otilia Stretcu, Graham Neubig, Barnabas Poczos, and Tom Mitchell. 2019. [Competence-based curriculum learning for neural machine translation](#). In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1162–1172.
- Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proceedings of the Conference on Machine Translation: Research Papers*, pages 186–191.
- Ivan Provilkov, Dmitrii Emelianenko, and Elena Voita. 2020. [BPE-Dropout: Simple and effective subword regularization](#). In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pages 1882–1892.
- Sebastian Ruder and Barbara Plank. 2017. [Learning to select data for transfer learning with Bayesian optimization](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 372–382.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016a. [Improving neural machine translation models with monolingual data](#). In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pages 86–96.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016b. [Neural machine translation of rare words with subword units](#). In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pages 1715–1725.
- Kaitao Song, Xu Tan, Tao Qin, Jianfeng Lu, and Tie-Yan Liu. 2019. [MASS: Masked Sequence to Sequence pre-training for language generation](#). In *Proceedings of the International Conference on Machine Learning*, pages 5926–5936.
- Dario Stojanovski, Viktor Hangya, Matthias Huck, and Alexander Fraser. 2019. [The LMU Munich unsupervised machine translation system for WMT19](#). In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 393–399.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, ukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, page 59986008.
- Pascal Vincent, Hugo Larochelle, Yoshua Bengio, and Pierre-Antoine Manzagol. 2008. [Extracting and composing robust features with denoising autoencoders](#). In *Proceedings of the International Conference on Machine Learning*, pages 1096–1103.
- Wei Wang, Ye Tian, Jiquan Ngiam, Yinfei Yang, Isaac Caswell, and Zarana Parekh. 2020. [Learning a multi-domain curriculum for neural machine translation](#). In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pages 7711–7723.
- Xuan Zhang, Pamela Shapiro, Gaurav Kumar, Paul McNamee, Marine Carpuat, and Kevin Duh. 2019. [Curriculum learning for domain adaptation in neural machine translation](#). In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1903–1915.