

Incorporating Word and Subword Units in Unsupervised Machine Translation Using Language Model Rescoring

Zihan Liu*, Yan Xu*, Genta Indra Winata, Pascale Fung

Center for Artificial Intelligence Research (CAiRE)

Department of Electronic and Computer Engineering

The Hong Kong University of Science and Technology, Clear Water Bay, Hong Kong

{zliurc, yxucb, giwinata}@connect.ust.hk, pascale@ece.ust.hk

Abstract

This paper describes CAiRE’s submission to the unsupervised machine translation track of the WMT’19 news shared task from German to Czech. We leverage a phrase-based statistical machine translation (PBSMT) model and a pre-trained language model to combine word-level neural machine translation (NMT) and subword-level NMT models without using any parallel data. We propose to solve the morphological richness problem of languages by training byte-pair encoding (BPE) embeddings for German and Czech separately, and they are aligned using MUSE (Conneau et al., 2018). To ensure the fluency and consistency of translations, a rescoring mechanism is proposed that reuses the pre-trained language model to select the translation candidates generated through beam search. Moreover, a series of pre-processing and post-processing approaches are applied to improve the quality of final translations.

1 Introduction

Machine translation (MT) has achieved huge advances in the past few years (Bahdanau et al., 2015; Gehring et al., 2017; Vaswani et al., 2017, 2018). However, the need for a large amount of manual parallel data obstructs its performance under low-resource conditions. Building an effective model on low resource data or even in an unsupervised way is always an interesting and challenging research topic (Gu et al., 2018; Radford et al., 2016; Lee et al., 2019). Recently, unsupervised MT (Artetxe et al., 2018b,a; Conneau et al., 2018; Lample et al., 2018b; Wu et al., 2019), which can immensely reduce the reliance on parallel corpora, has been gaining more and more interest.

Training cross-lingual word embeddings (Conneau et al., 2018; Artetxe et al., 2017) is always the

first step of the unsupervised MT models which produce a word-level shared embedding space for both the source and target, but the lexical coverage can be an intractable problem. To tackle this issue, Sennrich et al. (2016b) provided a subword-level solution to overcome the out-of-vocabulary (OOV) problem.

In this work, the systems we implement for the German-Czech language pair are built based on the previously proposed unsupervised MT systems, with some adaptations made to accommodate the morphologically rich characteristics of German and Czech (Tsarfaty et al., 2010). Both word-level and subword-level neural machine translation (NMT) models are applied in this task and further tuned by pseudo-parallel data generated from a phrase-based statistical machine translation (PBSMT) model, which is trained following the steps proposed in Lample et al. (2018b) without using any parallel data. We propose to train BPE embeddings for German and Czech separately and align those trained embeddings into a shared space with MUSE (Conneau et al., 2018) to reduce the combinatorial explosion of word forms for both languages. To ensure the fluency and consistency of translations, an additional Czech language model is trained to select the translation candidates generated through beam search by rescoring them. Besides the above, a series of post-processing steps are applied to improve the quality of final translations. Our contribution is two-fold:

- We propose a method to combine word and subword (BPE) pre-trained input representations aligned using MUSE (Conneau et al., 2018) as an NMT training initialization on a morphologically-rich language pair such as German and Czech.
- We study the effectiveness of language model

*These two authors contributed equally.

rescoring to choose the best sentences and unknown word replacement (UWR) procedure to reduce the drawback of OOV words.

This paper is organized as follows: in Section 2, we describe our approach to the unsupervised translation from German to Czech. Section 3 reports the training details and the results for each steps of our approach. More related work is provided in Section 4. Finally, we conclude our work in Section 5.

2 Methodology

In this section, we describe how we built our main unsupervised machine translation system, which is illustrated in Figure 1.

2.1 Unsupervised Machine Translation

2.1.1 Word-level Unsupervised NMT

We follow the unsupervised NMT in Lample et al. (2018b) by leveraging initialization, language modeling and back-translation. However, instead of using BPE, we use MUSE (Conneau et al., 2018) to align word-level embeddings of German and Czech, which are trained by FastText (Bojanowski et al., 2017) separately. We leverage the aligned word embeddings to initialize our unsupervised NMT model.

The language model is a denoising auto-encoder, which is trained by reconstructing original sentences from noisy sentences. The process of language modeling can be expressed as minimizing the following loss:

$$L^{lm} = \lambda * \{E_{x \sim S}[-\log P_{s \rightarrow s}(x|N(x))] + E_{y \sim T}[-\log P_{t \rightarrow t}(x|N(y))]\}, \quad (1)$$

where N is a noise model to drop and swap some words with a certain probability in the sentence x , $P_{s \rightarrow s}$ and $P_{t \rightarrow t}$ operate on the source and target sides separately, and λ acts as a weight to control the loss function of the language model. a Back-translation turns the unsupervised problem into a supervised learning task by leveraging the generated pseudo-parallel data. The process of back-translation can be expressed as minimizing the following loss:

$$L^{bt} = E_{x \sim S}[-\log P_{t \rightarrow s}(x|v^*(x))] + E_{y \sim T}[-\log P_{s \rightarrow t}(y|u^*(y))], \quad (2)$$

where $v^*(x)$ denotes sentences in the target language translated from source language sentences

S , $u^*(y)$ similarly denotes sentences in the source language translated from the target language sentences T and $P_{t \rightarrow s}$, and $P_{s \rightarrow t}$ denote the translation direction from target to source and from source to target respectively.

2.1.2 Subword-level Unsupervised NMT

We note that both German and Czech (Tsarfaty et al., 2010) are morphologically rich languages, which leads to a very large vocabulary size for both languages, but especially for Czech (more than one million unique words for German, but three million unique words for Czech). To overcome OOV issues, we leverage subword information, which can lead to better performance.

We employ subword units (Sennrich et al., 2016a) to tackle the morphological richness problem. There are two advantages of using the subword-level. First, we can alleviate the OOV issue by zeroing out the number of unknown words. Second, we can leverage the semantics of subword units from these languages. However, German and Czech are distant languages that originate from different roots, so they only share a small fraction of subword units. To tackle this problem, we train FastText word vectors (Bojanowski et al., 2017) separately for German and Czech, and apply MUSE (Conneau et al., 2018) to align these embeddings.

2.1.3 Unsupervised PBSMT

PBSMT models can outperform neural models in low-resource conditions. A PBSMT model utilizes a pre-trained language model and a phrase table with phrase-to-phrase translations from the source language to target languages, which provide a good initialization. The phrase table stores the probabilities of the possible target phrase translations corresponding to the source phrases, which can be referred to as $P(s|t)$, with s and t representing the source and target phrases. The source and target phrases are mapped according to inferred cross-lingual word embeddings, which are trained with monolingual corpora and aligned into a shared space without any parallel data (Artetxe et al., 2017; Conneau et al., 2018).

We use a pre-trained n-gram language model to score the phrase translation candidates by providing the relative likelihood estimation $P(t)$, so that the translation of a source phrase is derived from: $argmax_t P(t|s) = argmax_t P(s|t)P(t)$.

Back-translation enables the PBSMT models

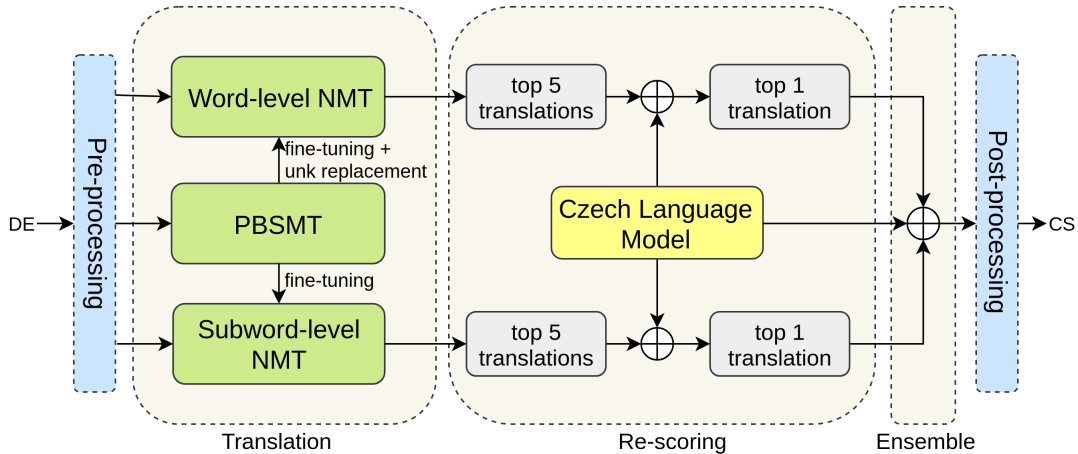


Figure 1: The illustration of our system. The translation procedure can be divided into five steps: (a) pre-processing, (b) translation generation (§2.1) from word-level NMT, subword-level NMT, and PBSMT. In the training, we fine-tune word-level and subword-level NMT models with pseudo-parallel data from NMT models and the best PBSMT model. Moreover, an unknown word replacement mechanism (§2.2) is applied to the translations generated from the word-level NMT model, (c) translation candidate rescoring, (d) construction of an ensemble of the translations from NMT models, and (e) post-processing.

to be trained in a supervised way by providing pseudo-parallel data from the translation in the reverse direction, which indicates that the PBSMT models need to be trained in dual directions so that the two models trained in the opposite directions can promote each other’s performance.

In this task, we follow the method proposed by Lample et al. (2018b) to initialize the phrase table, train the KenLM language models (Heafield, 2011)¹ and train a PBSMT model, but we make two changes. First, we only initialize a uni-gram phrase table because of the large vocabulary size of German and Czech and the limitation of computational resources. Second, instead of training the model in the *truecase* mode, we maintain the same pre-processing step (see more details in §3.1) as the NMT models.

2.1.4 Fine-tuning NMT

We further fine-tune the NMT models mentioned above on the pseudo-parallel data generated by a PBSMT model. We choose the best PBSMT model and mix the pseudo-parallel data from the NMT models and the PBSMT model, which are used for back-translation. The intuition is that we can use the pseudo-parallel data produced by the PBSMT model as the supplementary translations in our NMT model, and these can potentially boost the robustness of the NMT model by increasing the variety of back-translation data.

¹The code can be found at <https://github.com/kpu/kenlm>

2.2 Unknown Word Replacement

Around 10% of words found in our NMT training data are unknown words (<UNK>), which immensely limits the potential of the word-level NMT model. In this case, replacing unknown words with reasonable words can be a good remedy. Then, assuming the translations from the word-level NMT model and PBSMT model are roughly aligned in order, we can replace the unknown words in the NMT translations with the corresponding words in the PBSMT translations. Compared to the word-level NMT model, the PBSMT model ensures that every phrase will be translated without omitting any pieces from the sentences. We search for the word replacement by the following steps, which are also illustrated in Figure 2:

Step 1 For every unknown word, we can get the context words with a context window size of two.

Step 2 Each context word is searched for in the corresponding PBSMT translation. From our observation, the meanings of the words in Czech are highly likely to be the same if only the last few characters are different. Therefore, we allow the last two characters to be different between the context words and the words they match.

Step 3 If several words in the PBSMT translation match a context word, the word that is closest to the position of the context word in the PBSMT translation will be selected and put into the can-

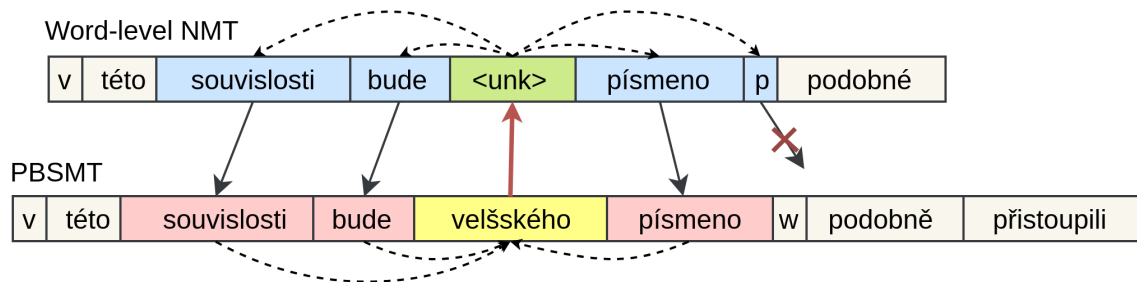


Figure 2: The illustration of the unknown word replacement (UWR) procedure for word-level NMT. The words of the PBSMT model translation in the pink boxes match the context words of the unknown word <UNK> in the word-level NMT model translation in the blue boxes. Finally, we choose a possible target word, in the yellow box, from the PBSMT model translation to replace the unknown word in the green box.

didate list to replace the corresponding <UNK> in the translation from the word-level NMT model.

Step 4 Step 2 and Step 3 are repeated until all the context words have been searched. After removing all the punctuation and the context words in the candidate list, the replacement word is the one that most frequently appears in the candidate list. If no candidate word is found, we just remove the <UNK> without adding a word.

2.3 Language Model Rescoring

Instead of direct translation with NMT models, we generate several translation candidates using beam search with a beam size of five. We build the language model proposed by Merity et al. (2018b,a) trained using a monolingual Czech dataset to rescore the generated translations. The scores are determined by the perplexity (PPL) of the generated sentences and the translation candidate with the lowest PPL will be selected as the final translation.

2.4 Model Ensemble

Ensemble methods have been shown very effective in many natural language processing tasks (Park et al., 2018; Winata et al., 2019). We apply an ensemble method by taking the top five translations from word-level and subword-level NMT, and rescore all translations using our pre-trained Czech language model mentioned in §2.3. Then, we select the best translation with the lowest perplexity.

3 Experiments

3.1 Data Pre-processing

We note that in the corpus, there are tokens representing quantity or date. Therefore, we delex-

icalize the tokens using two special tokens: (1) <NUMBER> to replace all the numbers that express a specific quantity, and (2) <DATE> to replace all the numbers that express a date. Then, we retrieve these numbers in the post-processing. There are two advantages of data pre-processing. First, replacing numbers with special tokens can reduce vocabulary size. Second, the special tokens are more easily processed by the model.

3.2 Data Post-processing

Special Token Replacement In the pre-processing, we use the special tokens <NUMBER> and <DATE> to replace numbers that express a specific quantity and date respectively. Therefore, in the post-processing, we need to restore those numbers. We simply detect the pattern <NUMBER> and <DATE> in the original source sentences and then replace the special tokens in the translated sentences with the corresponding numbers detected in the source sentences. In order to make the replacement more accurate, we will detect more complicated patterns like <NUMBER> / <NUMBER> in the original source sentences. If the translated sentences also have the pattern, we replace this pattern <NUMBER> / <NUMBER> with the corresponding numbers in the original source sentences.

Quotes Fixing The quotes are fixed to keep them the same as the source sentences.

Recaser For all the models mentioned above that work under a lower-case setting, a recaser implemented with Moses (Koehn et al., 2007) is applied to convert the translations to the real cases.

Patch-up From our observation, the ensemble NMT model lacks the ability to translate name entities correctly. We find that words with capital

characters are named entities, and those named entities in the source language may have the same form in the target language. Hence, we capture and copy these entities at the end of the translation if they does not exist in our translation.

3.3 Training

Unsupervised NMT The settings of the word-level NMT and subword-level NMT are the same, except the vocabulary size. We use a vocabulary size of 50k in the word-level NMT setting and 40k in the subword-level NMT setting for both German and Czech. In the encoder and decoder, we use a transformer (Vaswani et al., 2017) with four layers and a hidden size of 512. We share all encoder parameters and only share the first decoder layer across two languages to ensure that the latent representation of the source sentence is robust to the source language. We train auto-encoding and back-translation during each iteration. As the training goes on, the importance of language modeling become a less important compared to back-translation. Therefore the weight of auto-encoding (λ in equation (1)) is decreasing during training.

Unsupervised PBSMT The PBSMT is implemented with Moses using the same settings as those in Lample et al. (2018b). The PBSMT model is trained iteratively. Both monolingual datasets for the source and target languages consist of 12 million sentences, which are taken from the latest parts of the WMT monolingual dataset. At each iteration, two out of 12 million sentences are randomly selected from the the monolingual dataset.

Language Model According to the findings in Cotterell et al. (2018), the morphological richness of a language is closely related to the performance of the model, which indicates that the language models will be extremely hard to train for Czech, as it is one of the most complex languages. We train the QRNN model with 12 million sentences randomly sampled from the original WMT Czech monolingual dataset,² which is also pre-processed in the way mentioned in §3.1. To maintain the quality of the language model, we enlarge the vocabulary size to three million by including all the words that appear more than 15 times. Finally, the PPL of the language model on the test set achieves 93.54.

²<http://www.statmt.org/wmt19/>

Recaser We use the recaser model provided in Moses and train the model with the two million latest sentences in the Czech monolingual dataset. After the training procedure, the recaser can restore words to the form in which the maximum probability occurs.

3.4 PBSMT Model Selection

The BLEU (cased) score of the initialized phrase table and models after training at different iterations are shown in Table 1. From comparing the results, we observe that back-translation can improve the quality of the phrase table significantly, but after five iterations, the phrase table has hardly improved. The PBSMT model at the sixth iteration is selected as the final PBSMT model.

Model	BLEU Cased
Unsupervised PBSMT	
Unsupervised Phrase Table	3.8
+ Back-translation Iter. 1	6.6
+ Back-translation Iter. 2	7.3
+ Back-translation Iter. 3	7.5
+ Back-translation Iter. 4	7.6
+ Back-translation Iter. 5	7.7
+ Back-translation Iter. 6	7.7

Table 1: Results of PBSMT at different iterations

3.5 Results

The performances of our final model and other baseline models are illustrated in Table 2. In the baseline unsupervised NMT models, subword-level NMT outperforms word-level NMT by around a 1.5 BLEU score. Although the unsupervised PBSMT model is worse than the subword-level NMT model, leveraging generated pseudo-parallel data from the PBSMT model to fine-tune the subword-level NMT model can still boost its performance. However, this pseudo-parallel data from the PBSMT model can not improve the word-level NMT model since the large percentage of OOV words limits its performance. After applying unknown words replacement to the word-level NMT model, the performance improves by a BLEU score of around 2. Using the Czech language model to re-score helps the model improve by around a 0.3 BLEU score each time. We also use this language model to create an ensemble of the best word-level and subword-level NMT model and achieve the best performance.

Model	BLEU	BLEU Cased	TER	BEER 2.0	CharacterTER
Unsupervised PBSMT					
Unsupervised phrase table	4	3.8	-	0.384	0.773
+ Back-translation Iter. 6	8.3	7.7	0.887	0.429	0.743
Unsupervised NMT					
Subword-level NMT	9.4	9.1	-	0.419	0.756
+ fine-tuning	9.8	9.5	0.832	0.424	0.756
+ fine-tuning + rescoring	10.3	10	0.833	0.426	0.749
Word-level NMT	7.9	7.6	-	0.412	0.823
+ fine-tuning	7.9	7.7	-	0.413	0.819
+ fine-tuning + UWR	10.1	9.6	0.829	0.432	0.766
+ fine-tuning + UWR + rescoring	10.4	9.9	0.829	0.432	0.764
Model Ensemble					
Best Word-level + Subword-level	10.6	10.2	0.829	0.429	0.755
+ patch-up	10.6	10.2	0.833	0.430	0.757

Table 2: Unsupervised translation results. We report the scores of several evaluation methods for every step of our approach. Except the result that is listed on the last line, all results are under the condition that the translations are post-processed without patch-up.

4 Related Work

4.1 Unsupervised Cross-lingual Embeddings

Cross-lingual word embeddings can provide a good initialization for both the NMT and SMT models. In the unsupervised scenario, Artetxe et al. (2017) independently trained embeddings in different languages using monolingual corpora, and then learned a linear mapping to align them in a shared space based on a bilingual dictionary of a negligibly small size. Conneau et al. (2018) proposed a fully unsupervised learning method to build a bilingual dictionary without using any foregone word pairs, but by considering words from two languages that are near each other as pseudo word pairs. Lample and Conneau (2019) showed that cross-lingual language model pre-training can learn a better cross-lingual embeddings to initialize an unsupervised machine translation model.

4.2 Unsupervised Machine Translation

In Artetxe et al. (2018b) and Lample et al. (2018a), the authors proposed the first unsupervised machine translation models which combines an auto-encoding language model and back-translation in the training procedure. Lample et al. (2018b) illustrated that initialization, language modeling, and back-translation are key for both unsupervised neural and statistical machine translation. Artetxe et al. (2018a) combined back-translation

and MERT (Och, 2003) to iteratively refine the SMT model. Wu et al. (2019) proposed to discard back-translation. Instead, they extracted and edited the nearest sentences in the target language to construct pseudo-parallel data, which was used as a supervision signal.

5 Conclusion

In this paper, we propose to combine word-level and subword-level input representation in unsupervised NMT training on a morphologically rich language pair, German-Czech, without using any parallel data. Our results show the effectiveness of using language model rescoring to choose more fluent translation candidates. A series of pre-processing and post-processing approaches improve the quality of final translations, particularly to replace unknown words with possible relevant target words.

Acknowledgments

We would like to thank our colleagues Jamin Shin, Andrea Madotto, and Peng Xu for insightful discussions. This work has been partially funded by ITF/319/16FP and MRP/055/18 of the Innovation Technology Commission, the Hong Kong SAR Government.

References

- Mikel Artetxe, Gorka Labaka, and Eneko Agirre. 2017. Learning bilingual word embeddings with (almost) no bilingual data. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 451–462.
- 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*, pages 3632–3642.
- Mikel Artetxe, Gorka Labaka, Eneko Agirre, and Kyunghyun Cho. 2018b. [Unsupervised neural machine translation](#). In *International Conference on Learning Representations*.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. Enriching word vectors with subword information. *Transactions of the Association for Computational Linguistics*, 5:135–146.
- Alexis Conneau, Guillaume Lample, Marc’Aurelio Ranzato, Ludovic Denoyer, and Hervé Jégou. 2018. Word translation without parallel data. In *International Conference on Learning Representations (ICLR)*.
- Ryan Cotterell, Sebastian J Mielke, Jason Eisner, and Brian Roark. 2018. Are all languages equally hard to language-model? In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 536–541.
- Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N Dauphin. 2017. Convolutional sequence to sequence learning. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1243–1252. JMLR. org.
- Jiatao Gu, Hany Hassan, Jacob Devlin, and Victor OK Li. 2018. Universal neural machine translation for extremely low resource languages. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 344–354.
- Kenneth Heafield. 2011. Kenlm: Faster and smaller language model queries. In *Proceedings of the sixth workshop on statistical machine translation*, pages 187–197. Association for Computational Linguistics.
- Philipp Koehn, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran, Richard Zens, et al. 2007. Moses: Open source toolkit for statistical machine translation. In *Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics Companion Volume Proceedings of the Demo and Poster Sessions*, pages 177–180.
- Guillaume Lample and Alexis Conneau. 2019. Cross-lingual language model pretraining. *arXiv preprint arXiv:1901.07291*.
- Guillaume Lample, Alexis Conneau, Ludovic Denoyer, and Marc’Aurelio Ranzato. 2018a. [Unsupervised machine translation using monolingual corpora only](#). In *International Conference on Learning Representations*.
- Guillaume Lample, Myle Ott, Alexis Conneau, Ludovic Denoyer, and Marc’Aurelio Ranzato. 2018b. Phrase-based & neural unsupervised machine translation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Nayeon Lee, Zihan Liu, and Pascale Fung. 2019. Team yeon-zi at semeval-2019 task 4: Hyperpartisan news detection by de-noising weakly-labeled data. In *Proceedings of the 13th International Workshop on Semantic Evaluation*, pages 1052–1056.
- Stephen Merity, Nitish Shirish Keskar, and Richard Socher. 2018a. An Analysis of Neural Language Modeling at Multiple Scales. *arXiv preprint arXiv:1803.08240*.
- Stephen Merity, Nitish Shirish Keskar, and Richard Socher. 2018b. [Regularizing and optimizing LSTM language models](#). In *International Conference on Learning Representations*.
- Franz Josef Och. 2003. Minimum error rate training in statistical machine translation. In *Proceedings of the 41st Annual Meeting on Association for Computational Linguistics-Volume 1*, pages 160–167. Association for Computational Linguistics.
- Ji Ho Park, Peng Xu, and Pascale Fung. 2018. Plusemo2vec at semeval-2018 task 1: Exploiting emotion knowledge from emoji and# hashtags. In *Proceedings of The 12th International Workshop on Semantic Evaluation*, pages 264–272.
- Alec Radford, Luke Metz, and Soumith Chintala. 2016. Unsupervised representation learning with deep convolutional generative adversarial networks. *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016a. 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.

Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016b. Neural machine translation of rare words with subword units. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 1715–1725.

Reut Tsarfaty, Djamé Seddah, Yoav Goldberg, Sandra Kübler, Marie Candito, Jennifer Foster, Yannick Versley, Ines Rehbein, and Lamia Tounsi. 2010. Statistical parsing of morphologically rich languages (spmrl): what, how and whither. In *Proceedings of the NAACL HLT 2010 First Workshop on Statistical Parsing of Morphologically-Rich Languages*, pages 1–12. Association for Computational Linguistics.

Ashish Vaswani, Samy Bengio, Eugene Brevdo, François Chollet, Aidan Gomez, Stephan Gouws, Llion Jones, Łukasz Kaiser, Nal Kalchbrenner, Niki Parmar, et al. 2018. Tensor2tensor for neural machine translation. In *Proceedings of the 13th Conference of the Association for Machine Translation in the Americas (Volume 1: Research Papers)*, volume 1, pages 193–199.

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 *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 6000–6010. Curran Associates Inc.

Genta Indra Winata, Andrea Madotto, Zhaojiang Lin, Jamin Shin, Yan Xu, Peng Xu, and Pascale Fung. 2019. CAiRE_HKUST at SemEval-2019 task 3: Hierarchical attention for dialogue emotion classification. In *Proceedings of the 13th International Workshop on Semantic Evaluation*, pages 142–147, Minneapolis, Minnesota, USA. Association for Computational Linguistics.

Jiawei Wu, Xin Wang, and William Yang Wang. 2019. Extract and edit: An alternative to back-translation for unsupervised neural machine translation. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1173–1183.