

Self-Training for Unsupervised Neural Machine Translation in Unbalanced Training Data Scenarios

Haipeng Sun^{1,2*}, Rui Wang³, Kehai Chen⁴,
Masao Utiyama⁴, Eiichiro Sumita⁴, and Tiejun Zhao^{1†}

¹Harbin Institute of Technology, Harbin, China ²JD AI Research, Beijing, China

³Shanghai Jiao Tong University, Shanghai, China

⁴National Institute of Information and Communications Technology (NICT), Kyoto, Japan
sunhaipeng6@jd.com, wangrui.nlp@gmail.com, tjzhao@hit.edu.cn
{khchen, mutiyama, eiichiro.sumita}@nict.go.jp

Abstract

Unsupervised neural machine translation (UNMT) that relies solely on massive monolingual corpora has achieved remarkable results in several translation tasks. However, in real-world scenarios, massive monolingual corpora do not exist for some extremely low-resource languages such as Estonian, and UNMT systems usually perform poorly when there is not adequate training corpus for one language. In this paper, we first define and analyze the unbalanced training data scenario for UNMT. Based on this scenario, we propose UNMT self-training mechanisms to train a robust UNMT system and improve its performance in this case. Experimental results on several language pairs show that the proposed methods substantially outperform conventional UNMT systems.

1 Introduction

Recently, unsupervised neural machine translation (UNMT) that relies solely on massive monolingual corpora has attracted a high level of interest in the machine translation community (Artetxe et al., 2018; Lample et al., 2018a; Yang et al., 2018; Lample et al., 2018b; Wu et al., 2019; Sun et al., 2019, 2020b). With the help of cross-lingual language model pretraining (Lample and Conneau, 2019; Song et al., 2019; Sun et al., 2020a), the denoising auto-encoder (Vincent et al., 2010), and back-translation (Sennrich et al., 2016a), UNMT has achieved remarkable results in several translation tasks.

However, in real-world scenarios, in contrast to the many large corpora available for high-resource languages such as English and French, massive monolingual corpora do not exist for some extremely low-resource languages such as Estonian.

*Part of this work was done when Haipeng Sun and Rui Wang were an internship research fellow and a researcher at NICT, respectively.

† Corresponding author.

Data size (sentences)	En-Fr	Fr-En
50M En and 50M Fr (Baseline)	36.63	34.38
25M En and 25M Fr	36.59	34.34
50M En and 2M Fr	31.01	31.06
2M En and 50M Fr	31.84	30.21
2M En and 2M Fr	30.91	29.86

Table 1: UNMT performance (BLEU score) for different training data sizes on En–Fr language pairs.

The UNMT system usually performs poorly in a low-resource scenario when there is not an adequate training corpus for one language.

In this paper, we first define and analyze the unbalanced training data scenario for UNMT. Based on this scenario, we propose a self-training mechanism for UNMT. In detail, we propose self-training with unsupervised training (ST-UT) and self-training with pseudo-supervised training (ST-PT) strategies to train a robust UNMT system that performs better in this scenario. To the best of our knowledge, this paper is the first work to explore the unbalanced training data scenario problem in UNMT. Experimental results on several language pairs show that the proposed strategies substantially outperform conventional UNMT systems.

2 Unbalanced Training Data Scenario

In this section, we first define the unbalanced training data scenario according to training data size. Consider one monolingual corpus $\{X\}$ in high-resource language L_1 and another monolingual corpus $\{Y\}$ in low-resource language L_2 . The data size of $\{X\}$ and $\{Y\}$ are denoted by $|X|$ and $|Y|$, respectively. In an unbalanced training data scenario, $|X|$ is generally much larger than $|Y|$ so that training data $\{X\}$ is not fully utilized.

To investigate UNMT performance in an unbalanced training data scenario, we empirically chose English (En) – French (Fr) as the language pair. The detailed experimental settings for UNMT are

given in Section 5. We used a transformer based XLM toolkit and followed the settings of [Lample and Conneau \(2019\)](#). We randomly extracted 2 million sentences for each language from all 50 million sentences in the En and Fr training corpora to create small corpora and simulate unbalanced training data scenarios.

Table 1 shows the UNMT performance for different training data sizes. The performance with 25M training sentences for both French and English configuration is similar to the baseline (50M training sentences for both French and English configuration). However, the UNMT performance decreased substantially (4–5 BLEU points) when the size of the training data decreased rapidly. In the unbalanced training data scenario, when training data for one language was added, they were not fully utilized and only slightly improved the UNMT’s BLEU score. The performance (2M/50M) is similar with the UNMT system, configured 2M training sentences for both French and English. In short, Table 1 demonstrates that the UNMT performance is bounded by the smaller monolingual corpus. The UNMT model converges and even causes over-fitting in the low-resource language while the model in the high-resource language doesn’t converge. This observation motivates us to better use the larger monolingual corpus in the unbalanced training data scenario.

3 Background

We first briefly describe the three components of the UNMT model ([Lample and Conneau, 2019](#)): cross-lingual language model pre-training, the denoising auto-encoder ([Vincent et al., 2010](#)), and back-translation ([Sennrich et al., 2016a](#)). Cross-lingual language model pre-training provides a naive bilingual signal that enables the back-translation to generate pseudo-parallel corpora at the beginning of the training. The denoising auto-encoder acts as a language model to improve translation quality by randomly performing local substitutions and word reorderings.

Generally, back-translation plays an important role in achieving unsupervised translation across two languages. The pseudo-parallel sentence pairs produced by the model at the previous iteration are used to train the new translation model. The general back-translation probability is optimized

by maximizing

$$\begin{aligned} \mathcal{L}_{bt} = & \mathbb{E}_{X \sim P(X)} \mathbb{E}_{Y \sim P_{M^{U^*}}(Y|X)} \log P_{M^U}(X|Y) \\ & + \mathbb{E}_{Y \sim P(Y)} \mathbb{E}_{X \sim P_{M^{U^*}}(X|Y)} \log P_{M^U}(Y|X), \end{aligned} \quad (1)$$

where $P(X)$ and $P(Y)$ are the empirical data distribution from monolingual corpora $\{X\}$, $\{Y\}$, and $P_{M^U}(Y|X)$ and $P_{M^U}(X|Y)$ are the conditional distributions generated by the UNMT model. In addition, M^{U^*} denotes the model at the previous iteration for generating new pseudo-parallel sentence pairs to update the UNMT model.

Self-training proposed by [Scudder \(1965\)](#), is a semi-supervised approach that utilizes unannotated data to create better models. Self-training has been successfully applied to many natural language processing tasks ([Yarowsky, 1995](#); [McClosky et al., 2006](#); [Zhang and Zong, 2016](#); [He et al., 2020](#)). Recently, [He et al. \(2020\)](#) empirically found that noisy self-training could improve the performance of supervised machine translation and synthetic data could play a positive role, even as a target.

4 Self-training Mechanism for UNMT

Based on these previous empirical findings and analyses, we propose a self-training mechanism to generate synthetic training data for UNMT to alleviate poor performance in the unbalanced training data scenario. The synthetic data increases the diversity of low-resource language data, further enhancing the performance of the translation, even though the synthetic data may be noisy. As the UNMT model is trained, the quality of synthetic data becomes better, causing less and less noise. Compared with the original UNMT model that the synthetic data is just used as the source part, we also use the synthetic data as the target part in our proposed methods. Newly generated synthetic data, together with original monolingual data, are fully utilized to train a robust UNMT system in this scenario. According to the usage of the generated synthetic training data, our approach can be divided into two strategies: ST-UT (Algorithm 1) and ST-PT (Algorithm 2).

ST-UT: In this strategy, we first train a UNMT model on the existing monolingual training data. The final UNMT system is trained using the ST-UT strategy for k_1 epochs. For one epoch l in the ST-UT strategy, a subset $\{X^{sub}\}$ is selected randomly from monolingual training data $\{X\}$. The quantity of $\{X^{sub}\}$ is ϵ of $|X|$, ϵ is a quantity

Algorithm 1 ST-UT strategy

Input: Monolingual training data $\{X\}, \{Y\}$

- 1: Train a UNMT model M_0^U on monolingual training data $\{X\}, \{Y\}$
- 2: **while** epoch $l \leq \max \text{epoch } k_1$ **do**
- 3: Select a subset $\{X^{sub}\}$ randomly on monolingual training data $\{X\}$
- 4: Apply the last trained UNMT model M_{l-1}^U to this subset $\{X^{sub}\}$ to generate synthetic data $\{Y_M^{sub}\} = \{M_{l-1}^U(X^{sub})\}$
- 5: Train a new UNMT model M_l^U on monolingual data $\{X\}, \{Y\}$ and synthetic data $\{Y_M^{sub}\}$
- 6: **end while**

Output: The final translation model $M_{k_1}^U$

ratio hyper-parameter. The last trained UNMT model M_{l-1}^U is used to generate synthetic data $\{Y_M^{sub}\} = \{M_{l-1}^U(X^{sub})\}$. The synthetic data are used¹, together with the monolingual data to train a new UNMT model M_l^U . Therefore, the translation probability for the ST-UT strategy is optimized by maximizing

$$\begin{aligned} \mathcal{L}_{bt} = & \mathbb{E}_{X \sim P(X)} \mathbb{E}_{Y \sim P_{M_l^U}^*(Y|X)} \log P_{M_l^U}(X|Y) \\ & + \mathbb{E}_{Y \sim P(Y)} \mathbb{E}_{X \sim P_{M_l^U}^*(X|Y)} \log P_{M_l^U}(Y|X) \\ & + \mathbb{E}_{Y \sim P_{M_{l-1}^U}^*(Y|X)} \mathbb{E}_{X \sim P_{M_{l-1}^U}^*(X|Y)} \log P_{M_l^U}(Y|X), \end{aligned} \quad (2)$$

where $P_{M_l^U}(Y|X)$ and $P_{M_l^U}(X|Y)$ are the conditional distribution generated by the UNMT model on epoch l for the ST-UT strategy and $P_{M_{l-1}^U}^*(Y|X)$ is the conditional distribution generated by the UNMT model on epoch $l-1$ for the ST-UT strategy.

ST-PT: In this strategy, we first train a UNMT system on the existing monolingual training data and switch to a standard neural machine translation system from UNMT system with synthetic parallel data for both translation directions. The final translation system is trained using the ST-PT strategy for k_2 epochs. For one epoch q in the ST-PT strategy, a subset $\{X^{sub}\}$ is selected randomly from monolingual training data $\{X\}$, and all monolingual data $\{Y\}$ is selected. The quantity of $\{X^{sub}\}$ is ϵ of $|X|$, ϵ is a quantity ratio hyper-parameter. The last trained pseudo-supervised neu-

¹In contrast to using all synthetic data, we tried to train a language model and select more fluent synthetic data according to a language model perplexity score. This did not improve translation performance.

Algorithm 2 ST-PT strategy

Input: Monolingual training data $\{X\}, \{Y\}$

- 1: Train a UNMT model M_0^U on monolingual training data $\{X\}, \{Y\}$
- 2: **while** epoch $q \leq \max \text{epoch } k_2$ **do**
- 3: Select a subset $\{X^{sub}\}$ randomly on monolingual training data $\{X\}$ and all monolingual training data $\{Y^{all}\}$
- 4: Apply the last trained PNMT model $M_{q-1}^P(M_0^P = M_0^U)$ to generate $\{Y_M^{sub}\} = \{M_{q-1}^P(X^{sub})\}$ and $\{X_M^{all}\} = \{M_{q-1}^P(Y^{all})\}$
- 5: Train a new PNMT model M_q^P on synthetic parallel corpora $\{X^{sub}, Y_M^{sub}\}$ and $\{Y^{all}, X_M^{all}\}$
- 6: **end while**

Output: The final translation model $M_{k_2}^P$

ral machine translation (PNMT) model² M_{q-1}^P is used to generate $\{Y_M^{sub}\} = \{M_{q-1}^P(X^{sub})\}$ and $\{X_M^{all}\} = \{M_{q-1}^P(Y^{all})\}$ to create synthetic parallel data $\{X^{sub}, Y_M^{sub}\}$ and $\{Y^{all}, X_M^{all}\}$. Note that we use the UNMT model to generate synthetic parallel data during the first epoch of the ST-PT strategy. Synthetic parallel data $\{X^{sub}, Y_M^{sub}\}$ and $\{Y^{sub}, X_M^{sub}\}$ are selected to train a new PNMT model M_q^P that can generate translation in both directions. Therefore, the translation probability for ST-PT strategy is optimized by maximizing

$$\begin{aligned} \mathcal{L}_{bt} = & \mathbb{E}_{X \sim P(X)} \mathbb{E}_{Y \sim P_{M_{q-1}^P}^*(Y|X)} \log P_{M_q^P}(X|Y) \\ & + \mathbb{E}_{X \sim P(X)} \mathbb{E}_{Y \sim P_{M_{q-1}^P}^*(Y|X)} \log P_{M_q^P}(Y|X) \\ & + \mathbb{E}_{Y \sim P(Y)} \mathbb{E}_{X \sim P_{M_{q-1}^P}^*(X|Y)} \log P_{M_q^P}(Y|X) \\ & + \mathbb{E}_{Y \sim P(Y)} \mathbb{E}_{X \sim P_{M_{q-1}^P}^*(X|Y)} \log P_{M_q^P}(X|Y), \end{aligned} \quad (3)$$

where $P_{M_q^P}(Y|X)$ and $P_{M_q^P}(X|Y)$ are the conditional distributions generated by the PNMT model on epoch q for the ST-PT strategy; $P_{M_{q-1}^P}^*(Y|X)$ and $P_{M_{q-1}^P}^*(X|Y)$ are the conditional distributions generated by the PNMT model on epoch $q-1$ for the ST-PT strategy.

5 Experiments

5.1 Datasets

We considered three language pairs in our simulation experiments: Fr–En, Romanian (Ro)–En and

²Only synthetic parallel data were used to train PNMT model.

Method	En-Fr	Fr-En	En-Ro	Ro-En	En-Et	Et-En
Lample et al. (2018a)	15.05	14.31	n/a	n/a	n/a	n/a
Artetxe et al. (2018)	15.13	15.56	n/a	n/a	n/a	n/a
Lample et al. (2018b)	27.60	27.68	25.13	23.90	n/a	n/a
Lample and Conneau (2019)	33.40	33.30	33.30	31.80	n/a	n/a
UNMT	31.01	31.06	33.63	31.89	14.89	20.61
+ST-UT	34.43++	33.56++	35.04++	32.94++	17.05++	22.60++
+ST-PT	35.58++	34.91++	35.96++	33.64++	17.97++	24.97++

Table 3: Performance (BLEU score) of UNMT on the unbalanced training data scenario. Note that only 2 million Fr monolingual training data were used for En–Fr. The quantity ratio ϵ was set to 10%. The number of epochs was set to two for both proposed strategies. “++” after a score indicates that the strategy was significantly better than the baseline at significance level $p < 0.01$.

Estonian (Et)–En translation tasks. The statistics of the data are presented in Table 2. We used the monolingual WMT news crawl datasets³ for each language. For the high-resource languages En and Fr, we randomly extracted 50M sentences. For the low-resource languages Ro and Et, we used all available monolingual news crawl training data. To make our experiments comparable with previous work (Lample and Conneau, 2019), we report the results on newstest2014 for Fr–En, newstest2016 for Ro–En, and newstest2018 for Et–En.

Language	Sentences	Words
En	50.00M	1.15B
Fr	50.00M	1.19B
Ro	8.92M	207.07M
Et	3.00M	51.39M

Table 2: Statistics of the monolingual corpora.

For preprocessing, we used the Moses tokenizer (Koehn et al., 2007). To clean the data, we only applied the Moses script `clean-corpus-n.perl` to remove lines from the monolingual data containing more than 50 words. We used a shared vocabulary for all language pairs, with 60,000 subword tokens based on BPE (Sennrich et al., 2016b).

5.2 UNMT Settings

We used a transformer-based XLM toolkit and followed the settings of Lample and Conneau (2019) for UNMT: six layers for the encoder and the decoder. The dimensions of the hidden layers were set to 1024. The batch size was set to 2000 tokens. The Adam optimizer (Kingma and Ba, 2015) was used to optimize the model parameters. The initial

³<http://data.statmt.org/news-crawl/>

learning rate was 0.0001, $\beta_1 = 0.9$, and $\beta_2 = 0.98$. We trained a specific cross-lingual language model for each different training dataset. The language model was used to initialize the full parameters of the UNMT system. Eight V100 GPUs were used to train all UNMT models. We used the case-sensitive 4-gram BLEU score computed by the *multi-bleu.perl* script from Moses (Koehn et al., 2007) to evaluate the test sets.

5.3 Main Results

Table 3 presents the detailed BLEU scores of the UNMT systems on the En–Fr, En–Ro, and En–Et test sets. Our re-implemented baseline performed similarly to the state-of-the-art method of Lample and Conneau (2019) on the En–Ro language pair. In particular, we used only 2 million Fr monolingual training data on the En–Fr language pair, so the re-implemented baseline performed slightly worse than Lample and Conneau (2019).

Our proposed self-training mechanism substantially outperformed the corresponding baseline in all language pairs by 2–4 BLEU points. Regarding the two proposed strategies, the ST-PT strategy performed better than the ST-UT strategy by 1 BLEU point because the synthetic data are more directly integrated into the training. For ST-UT, the synthetic data was just used as the target part. In contrast, the synthetic data was used as the source and target part for ST-PT. The synthetic parallel data could improve translation performance. These results demonstrate that synthetic data improve translation performance in our proposed self-training mechanism. The detailed analyses of the hyperparameters such as quantity ratio ϵ and epoch number k_1, k_2 are provided in Appendix.

Input	Ma pole oma loomingust kunagi kaugel ja tööd ma ei karda .
Reference	I ’m never far from my work and I ’m not afraid of work .
Baseline	I am never far from my work and work I ’m not afraid of .
+ST-PT	I ’m never far from my work and I ’m not afraid of the job .
Input	Salvador Adame kadus läänepoolses Michoacani osariigis kolm päeva pärast Valdezi tapmmist .
Reference	Salvador Adame disappeared in the western state of Michoacan three days after Valdezi was killed .
Baseline	Salvador Adame disappeared in west Michoacan , Mexico , three days after Valdezi was killed .
+ST-PT	Salvador Adame disappeared in the western state of Michoacan three days after Valdezi was killed .

Table 4: Comparison of translation results of baseline and +ST-PT system on the Et-En dataset.

5.4 Case Study

Moreover, we analyze translation examples to further analyze the effectiveness of our proposed self-training mechanism. Table 4 shows two translation examples, which were generated by UNMT baseline system and +ST-PT system on the Et-En dataset, respectively. For the first example, +ST-PT method could make the translation more fluent, compared with the baseline system. For the second example, +ST-PT method could make the translation more accurate. These examples indicate that our proposed self-training mechanism could be widely applied to the unbalanced training data scenario.

6 Conclusion

UNMT has achieved remarkable results on massive monolingual corpora. However, a UNMT system usually does not perform well in a scenario where there is not an adequate training corpus for one language. Based on this unbalanced training data scenario, we proposed two self-training strategies for UNMT. Experimental results on several language pairs show that our proposed strategies substantially outperform UNMT baseline.

Acknowledgments

We are grateful to the anonymous reviewers and the area chair for their insightful comments and suggestions. Tiejun Zhao was partially supported by National Key Research and Development Program of China via grant 2020AAA0108001. Masao Utiyama and Eiichiro Sumita were partly supported by the commissioned research program “Research and Development of Advanced Multilingual Translation Technology” in the “R&D Project for Information and Communications Technology (JPMI00316)” of the Ministry of Internal Affairs and Communications (MIC), Japan.

References

- Mikel Artetxe, Gorka Labaka, Eneko Agirre, and Kyunghyun Cho. 2018. Unsupervised neural machine translation. In *Proceedings of the Sixth International Conference on Learning Representations*, Vancouver, Canada.
- Junxian He, Jiatao Gu, Jiajun Shen, and Marc’Aurelio Ranzato. 2020. Revisiting self-training for neural sequence generation. In *Proceedings of the 8th International Conference on Learning Representations*, Addis Ababa, Ethiopia.
- Diederik P Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *Proceedings of the Third International Conference on Learning Representations*, San Diego, California, USA.
- Philipp Koehn, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran, Richard Zens, Chris Dyer, Ondřej Bojar, Alexandra Constantin, and Evan Herbst. 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, Prague, Czech Republic.
- Guillaume Lample and Alexis Conneau. 2019. Cross-lingual language model pretraining. *CoRR*, abs/1901.07291.
- Guillaume Lample, Alexis Conneau, Ludovic Denoyer, and Marc’Aurelio Ranzato. 2018a. Unsupervised machine translation using monolingual corpora only. In *Proceedings of the Sixth International Conference on Learning Representations*, Vancouver, Canada.
- 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*, pages 5039–5049, Brussels, Belgium.
- David McClosky, Eugene Charniak, and Mark Johnson. 2006. Effective self-training for parsing. In *Proceedings of the Human Language Technology Conference of the NAACL, Main Conference*, pages 152–159, New York City, USA.

- H. J. Scudder. 1965. Probability of error of some adaptive pattern-recognition machines. *IEEE Trans. Information Theory*, 11(3):363–371.
- 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)*, pages 86–96, Berlin, Germany.
- 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)*, pages 1715–1725, Berlin, Germany.
- Kaitao Song, Xu Tan, Tao Qin, Jianfeng Lu, and Tiejun Zhao. 2019. Mass: Masked sequence to sequence pre-training for language generation. In *Proceedings of the 36th International Conference on Machine Learning*, pages 5926–5936, Long Beach, California, USA.
- Haipeng Sun, Rui Wang, Kehai Chen, Masao Utiyama, Eiichiro Sumita, and Tiejun Zhao. 2019. Unsupervised bilingual word embedding agreement for unsupervised neural machine translation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1235–1245, Florence, Italy.
- Haipeng Sun, Rui Wang, Kehai Chen, Masao Utiyama, Eiichiro Sumita, and Tiejun Zhao. 2020a. Knowledge distillation for multilingual unsupervised neural machine translation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3525–3535, Online.
- Haipeng Sun, Rui Wang, Kehai Chen, Masao Utiyama, Eiichiro Sumita, and Tiejun Zhao. 2020b. Unsupervised neural machine translation with cross-lingual language representation agreement. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 28:1170–1182.
- Pascal Vincent, Hugo Larochelle, Isabelle Lajoie, Yoshua Bengio, and Pierre-Antoine Manzagol. 2010. Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion. *Journal of Machine Learning Research*, 11:3371–3408.
- 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, Minneapolis, Minnesota, USA.
- Zhen Yang, Wei Chen, Feng Wang, and Bo Xu. 2018. Unsupervised neural machine translation with weight sharing. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 46–55, Melbourne, Australia.
- David Yarowsky. 1995. Unsupervised word sense disambiguation rivaling supervised methods. In *33rd Annual Meeting of the Association for Computational Linguistics*, pages 189–196, Cambridge, Massachusetts, USA.
- Jiajun Zhang and Chengqing Zong. 2016. Exploiting source-side monolingual data in neural machine translation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1535–1545, Austin, Texas, USA.

A Appendix

A.1 Quantity Ratio Analysis

We investigated the effect of quantity ratio ϵ on UNMT performance for the En–Fr translation task during the first epoch of our proposed self-training methods. As shown in Fig. 1, ϵ ranging from 1% to 100% all enhanced UNMT performance and the performance was similar when the quantity ratio ϵ was greater than 10%. The UNMT model converged faster with less data. Therefore, we selected 10% as the quantity ratio ϵ for our proposed self-training methods.

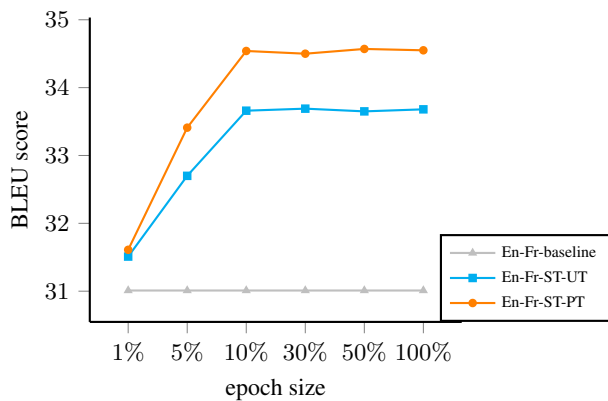


Figure 1: Effect of the quantity ratio ϵ on UNMT performance for the En–Fr translation tasks.

A.2 Epoch Number Analysis

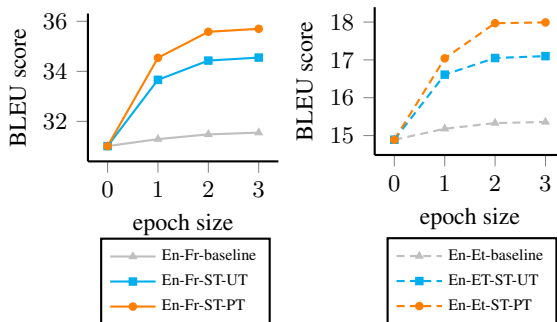


Figure 2: Effect of the number of epochs on UNMT performance for the En–Fr and En–Et translation tasks.

In Figure 2, we empirically demonstrate how the number of epochs affects the UNMT performance on the En–Fr and En–Et translation tasks. We found that the use of additional epochs has little influence on the baseline system. In contrast, increasing the number of epochs for our proposed strategies can improve performance because the quality of the synthetic data used by the UNMT model is better after more epochs; however, the

improvement decreases as additional epochs are added. Considering the computational cost of synthetic data generation, we trained the UNMT model for only two epochs.