

Greedy Search with Probabilistic N-gram Matching for Neural Machine Translation

Chenze Shao^{1,2} Yang Feng^{1,2*} Xilin Chen^{1,2}

¹Key Laboratory of Intelligent Information Processing
Institute of Computing Technology, Chinese Academy of Sciences (ICT/CAS)

²University of Chinese Academy of Science

{shaochenze18z, fengyang, xlchen}@ict.ac.cn

Abstract

Neural machine translation (NMT) models are usually trained with the word-level loss using the teacher forcing algorithm, which not only evaluates the translation improperly but also suffers from exposure bias. Sequence-level training under the reinforcement framework can mitigate the problems of the word-level loss, but its performance is unstable due to the high variance of the gradient estimation. On these grounds, we present a method with a differentiable sequence-level training objective based on probabilistic n-gram matching which can avoid the reinforcement framework. In addition, this method performs greedy search in the training which uses the predicted words as context just as at inference to alleviate the problem of exposure bias. Experiment results on the NIST Chinese-to-English translation tasks show that our method significantly outperforms the reinforcement-based algorithms and achieves an improvement of 1.5 BLEU points on average over a strong baseline system.

1 Introduction

Neural machine translation (NMT) (Kalchbrenner and Blunsom, 2013; Cho et al., 2014; Sutskever et al., 2014; Bahdanau et al., 2014) has now achieved impressive performance (Wu et al., 2016; Gehring et al., 2017; Vaswani et al., 2017; Hassan et al., 2018; Chen et al., 2018; Lample et al., 2018) and draws more attention. NMT models are built on the encoder-decoder framework where the encoder network encodes the source sentence to distributed representations and the decoder network reconstructs the target sentence form the representations word by word.

Currently, NMT models are usually trained with the word-level loss (i.e., cross-entropy) under the teacher forcing algorithm (Williams and Zipser,

1989), which forces the model to generate translation strictly matching the ground-truth at the word level. However, in practice it is impossible to generate translation totally the same as ground truth. Once different target words are generated, the word-level loss cannot evaluate the translation properly, usually under-estimating the translation. In addition, the teacher forcing algorithm suffers from the *exposure bias* (Ranzato et al., 2015) as it uses different inputs at training and inference, that is ground-truth words for the training and previously predicted words for the inference. Kim and Rush (2016) proposed a method of sequence-level knowledge distillation, which use teacher outputs to direct the training of student model, but the student model still have no access to its own predicted words. Scheduled sampling(SS) (Bengio et al., 2015; Venkatraman et al., 2015) attempts to alleviate the exposure bias problem through mixing ground-truth words and previously predicted words as inputs during training. However, the sequence generated by SS may not be aligned with the target sequence, which is inconsistent with the word-level loss.

In contrast, sequence-level objectives, such as BLEU (Papineni et al., 2002), GLEU (Wu et al., 2016), TER (Snover et al., 2006), and NIST (Doddington, 2002), evaluate translation at the sentence or n -gram level and allow for greater flexibility, and thus can mitigate the above problems of the word-level loss. However, due to the non-differentiable of sequence-level objectives, previous works on sequence-level training (Ranzato et al., 2015; Shen et al., 2016; Bahdanau et al., 2016; Wu et al., 2016; He et al., 2016; Wu et al., 2017; Yang et al., 2017) mainly rely on reinforcement learning algorithms (Williams, 1992; Sutton et al., 2000) to find an unbiased gradient estimator for the gradient update. Sparse rewards in this situation often cause the high variance of gradient estimation, which consequently leads to unstable

*Corresponding Author

training and limited improvements.

Lamb et al. (2016); Gu et al. (2017); Ma et al. (2018) respectively use the discriminator, critic and bag-of-words target as sequence-level training objectives, all of which are directly connected to the generation model and hence enable direct gradient update. However, these methods do not allow for direct optimization with respect to evaluation metrics.

In this paper, we propose a method to combine the strengths of the word-level and sequence-level training, that is the direct gradient update without gradient estimation from word-level training and the greater flexibility from sequence-level training. Our method introduces probabilistic n -gram matching which makes sequence-level objectives (e.g., BLEU, GLEU) differentiable. During training, it abandons teacher forcing and performs greedy search instead to take into consideration the predicted words. Experiment results show that our method significantly outperforms word-level training with the cross-entropy loss and sequence-level training under the reinforcement framework. The experiments also indicate that greedy search strategy indeed has superiority over teacher forcing.

2 Background

NMT is based on an end-to-end framework which directly models the translation probability from the source sentence \mathbf{x} to the target sentence $\hat{\mathbf{y}}$:

$$P(\hat{\mathbf{y}}|\mathbf{x}) = \prod_{j=1}^T p(\hat{y}_j | \hat{\mathbf{y}}_{<j}, \mathbf{x}, \theta), \quad (1)$$

where T is the target length and θ is the model parameters. Given the training set $D = \{X^M, Y^M\}$ with M sentences pairs, the training objective is to maximize the log-likelihood of the training data as

$$\begin{aligned} \theta &= \arg \max_{\theta} \{\mathcal{L}(\theta)\} \\ \mathcal{L}(\theta) &= \sum_{m=1}^M \sum_{j=1}^{l^m} \log(p(\hat{y}_j^m | \hat{\mathbf{y}}_{<j}^m, \mathbf{x}^m, \theta)), \end{aligned} \quad (2)$$

where the superior m indicates the m -th sentence in the dataset and l^m is the length of m -th target sentence.

In the above model, the probability of each target word $p(\hat{y}_j^m | \hat{\mathbf{y}}_{<j}^m, \mathbf{x}^m, \theta)$ is conditioned on the previous target words. The scenario is that in the

training time, the teacher forcing algorithm is employed and the ground truth words from the target sentence are fed as context, while during inference, the ground truth words are not available and the previous predicted words are instead fed as context. This discrepancy is called *exposure bias*.

3 Model

3.1 Sequence-Level Objectives

Many automatic evaluation metrics of machine translation, such as BLEU, GLEU and NIST, are based on the n -gram matching. Assuming that \mathbf{y} and $\hat{\mathbf{y}}$ are the output sentence and the ground truth sentence with length T and T' respectively, the count of an n -gram $\mathbf{g} = (g_1, \dots, g_n)$ in sentence \mathbf{y} is calculated as

$$C_{\mathbf{y}}(\mathbf{g}) = \sum_{t=0}^{T-n} \prod_{i=1}^n 1\{g_i = y_{t+i}\}, \quad (3)$$

where $1\{\cdot\}$ is the indicator function. The matching count of the n -gram \mathbf{g} between $\hat{\mathbf{y}}$ and \mathbf{y} is given by

$$C_{\hat{\mathbf{y}}}(\mathbf{g}) = \min(C_{\mathbf{y}}(\mathbf{g}), C_{\hat{\mathbf{y}}}(\mathbf{g})). \quad (4)$$

Then the precision p_n and the recall r_n of the predicted n -grams are calculated as follows

$$p_n = \frac{\sum_{\mathbf{g} \in \mathbf{y}} C_{\hat{\mathbf{y}}}(\mathbf{g})}{\sum_{\mathbf{g} \in \mathbf{y}} C_{\mathbf{y}}(\mathbf{g})}, \quad (5)$$

$$r_n = \frac{\sum_{\mathbf{g} \in \mathbf{y}} C_{\hat{\mathbf{y}}}(\mathbf{g})}{\sum_{\mathbf{g} \in \hat{\mathbf{y}}} C_{\hat{\mathbf{y}}}(\mathbf{g})}. \quad (6)$$

BLEU, the most widely used metric for machine translation evaluation, is defined based on the n -gram precision as follows

$$\text{BLEU} = \text{BP} \cdot \exp\left(\sum_{n=1}^N w_n \log p_n\right), \quad (7)$$

where BP stands for the brevity penalty and w_n is the weight for the n -gram. In contrast, GLEU is the minimum of recall and precision of 1-4 grams where 1-4 grams are counted together:

$$\text{GLEU} = \min(p_{1-4}, r_{1-4}). \quad (8)$$

3.2 probabilistic Sequence-Level Objectives

In the output sentence \mathbf{y} , the prediction probability varies among words. Some words are translated by the model with high confidence while some words are translated with high uncertainty.

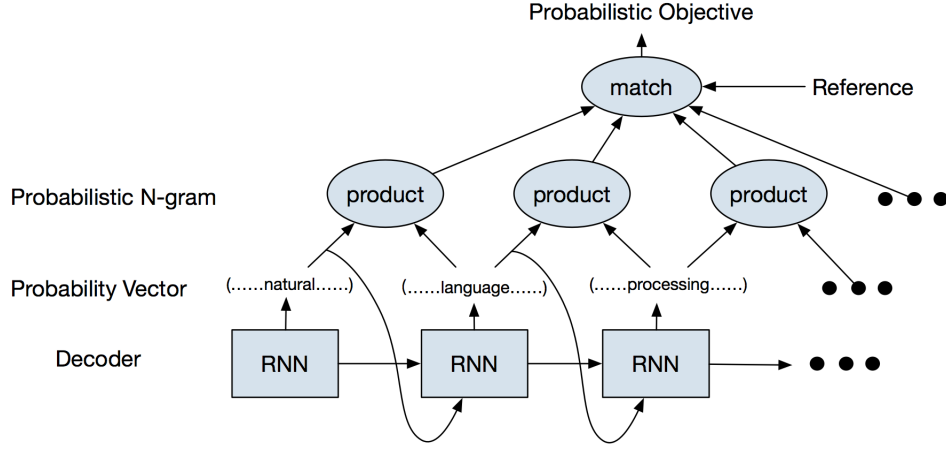


Figure 1: The overview of our model with greedy search. At each decoding step, the predicted word which has the highest probability in the probability vector is selected as context and fed into the RNN, and meanwhile this word and its probability are also used to calculate the probabilistic n -gram count.

However, when calculating the count of n -grams in Eq.(3), all the words in the output sentence are treated equally, regardless of their respective prediction probabilities.

To give a more precise description of n -gram counts which considers the variety of prediction probabilities, we use the prediction probability $p(y_j|\mathbf{y}_{<j}, \mathbf{x}, \theta)$ as the count of word y_j , and correspondingly the count of an n -gram is the product of these probabilistic counts of all the words in the n -gram, not *one* anymore. Then the probabilistic count of $\mathbf{g} = (g_1, \dots, g_n)$ is calculated by summing over the output sentence \mathbf{y} as

$$\tilde{C}_{\mathbf{y}}(\mathbf{g}) = \sum_{t=0}^{T-n} \prod_{i=1}^n 1\{g_i = y_{t+i}\} \cdot p(y_{t+i}|\mathbf{y}_{<t+i}, \mathbf{x}, \theta). \quad (9)$$

Now the probabilistic sequence-level objective can be got by replacing $C_{\mathbf{y}}(\mathbf{g})$ with $\tilde{C}_{\mathbf{y}}(\mathbf{g})$ (the tilde over the head indicates the probabilistic version) and keeping the rest unchanged. Here, we take BLEU as an example and show how the probabilistic BLEU (denoted as P-BLEU) is defined. From this purpose, the matching count of n -gram \mathbf{g} in Eq.(4) is modified as follows

$$\tilde{C}_{\mathbf{y}}^{\hat{\mathbf{y}}}(\mathbf{g}) = \min(\tilde{C}_{\mathbf{y}}(\mathbf{g}), C_{\hat{\mathbf{y}}}(\mathbf{g})). \quad (10)$$

and the predict precision of n -grams changes into

$$\tilde{p}_n = \frac{\sum_{\mathbf{g} \in \mathbf{y}} \tilde{C}_{\mathbf{y}}^{\hat{\mathbf{y}}}(\mathbf{g})}{\sum_{\mathbf{g} \in \mathbf{y}} \tilde{C}_{\mathbf{y}}(\mathbf{g})}. \quad (11)$$

Finally, the probabilistic BLEU (P-BLEU) is defined as

$$\text{P-BLEU} = \text{BP} \cdot \exp\left(\sum_{n=1}^N w_n \log \tilde{p}_n\right), \quad (12)$$

Probabilistic GLEU (P-GLEU) can be defined in a similar way. Specifically, we denote the probabilistic precision of n -grams as P-P $_n$. The probabilistic precision is more reasonable than recall since the denominator in Eq.(11) plays a normalization role, so we modify the definition in Eq.(8) and define P-GLEU as simply the probabilistic precision of 1-4 grams.

The general probabilistic loss function is:

$$\mathcal{L}(\theta) = - \sum_{m=1}^M \mathcal{P}(\mathbf{y}^m, \hat{\mathbf{y}}^m), \quad (13)$$

where \mathcal{P} represents the probabilistic sequence-level objectives, and \mathbf{y}^m and $\hat{\mathbf{y}}^m$ are the predicted translation and the ground truth for the m -th sentence respectively. The calculation of the probabilistic objective is illustrated in Figure 1. This probabilistic loss can work with decoding strategies such as greedy search and teacher forcing. In this paper we employ greedy search rather than teacher forcing so as to use the previously predicted words as context and alleviate the exposure bias problem.

System	Dev(MT02)	MT03	MT04	MT05	MT06	AVG
BaseNMT	36.72	33.95	37.44	33.96	33.09	34.61
MRT	37.17	34.89	37.90	34.62	33.78	35.30
RF	37.13	34.66	37.69	34.55	33.74	35.16
P-BLEU	37.26	34.54	38.05	34.30	34.11	35.25
P-GLEU	37.44	34.67	38.11	34.24	34.58	35.40
P-P2	38.03	35.45	39.30	35.10	34.59	36.11

Table 1: Results on NIST Chinese-to-English Translation Task. **AVG** = average BLEU scores for test sets. The bold number indicates the highest score in the column.

4 Experiment

4.1 Settings

We carry out experiments on Chinese-to-English translation.¹ The training data consists of 1.25M pairs of sentences extracted from LDC corpora². Sentence pairs with either side longer than 50 were dropped. We use NIST 2002 (MT 02) as the validation set and NIST 2003-2006 (MT 03-08) as the test sets. We use the case insensitive 4-gram NIST BLEU score (Papineni et al., 2002) for the translation task.

We apply our method to an attention-based NMT system (Bahdanau et al., 2014) implemented by Pytorch. Both source and target vocabularies are limited to 30K. All word embedding sizes are set to 512, and the sizes of hidden units in both encoder and decoder RNNs are also set to 512. All parameters are initialized by uniform distribution over $[-0.1, 0.1]$. The mini-batch stochastic gradient descent (SGD) algorithm is employed to train the model with batch size of 40. In addition, the learning rate is adjusted by adadelta optimizer (Zeiler, 2012) with $\rho = 0.95$ and $\epsilon = 1e-6$. Dropout is applied on the output layer with dropout rate of 0.5. The beam size is set to 10.

4.2 Performance

Systems We first pretrain the baseline model by maximum likelihood estimation (MLE) and then refine the model using probabilistic sequence-level objectives, including P-BLEU, P-GLEU and P-P2 (probabilistic 2-gram precision). In addition, we reproduce previous works which train the NMT model through minimum risk training (MRT) (Shen et al., 2016) and REINFORCE algo-

rithm (RF) (Ranzato et al., 2015). When reproducing their works, we set BLEU, GLEU and 2-gram precision as training objectives respectively and find out that GLEU yields the best performance. In the following, we only report the results with training objective GLEU.

Performance Table 1 shows the translation performance on test sets measured in BLEU score. Simply training NMT model by the probabilistic 2-gram precision achieves an improvement of 1.5 BLEU points, which significantly outperforms the reinforcement-based algorithms. We also test the precision of other n-grams and their combinations, but do not notice significant improvements over P-P2. Notice that our method only changes the loss function, without any modification on model structure and training data.

4.3 Why Pretraining

We use the probabilistic loss to finetune the baseline model rather than training from scratch. This is in line with our motivation: to alleviate the exposure bias and make the model exposed to its own output during training. In the very beginning of the training, the model’s translation capability is nearly zero and the generated sentences are often meaningless and do not contain useful information for the training, so it is unreasonable to directly apply the greedy search strategy. Therefore, we first apply the teacher forcing algorithm to pretrain the model, and then we let the model generate the sentences itself and learn from its own outputs.

Another reason favoring pretraining is that pretraining can lower the training cost. The training cost of the introduced probabilistic loss is about three times higher than the cost of cross entropy. Without pretraining, the training time will be much higher than usual. Otherwise, the training cost is acceptable if the probabilistic loss is only for finetuning.

¹Experiment code: <https://github.com/ictnlp/GS4NMT>

²The corpora include LDC2002E18, LDC2003E07, LDC2003E14, Hansards portion of LDC2004T07, LDC2004T08 and LDC2005T06.




Figure 2: learning curves of different decoding strategies with training objective P-P2.

4.4 Effect of Decoding Strategy

The probabilistic loss, defined in Eq.(13), is computed from the model output \mathbf{y} and reference $\hat{\mathbf{y}}$. In this section, we apply two different decoding strategies to generate \mathbf{y} : 1. teacher forcing, which uses the ground truth as decoder input. 2. greedy search, which feeds the word with maximum probability. By conducting this experiment, we attempt to figure out where the improvements come from: the modification of loss or the mitigation of exposure bias?

Figure 2 shows the learning curves of the two decoding strategies with training objective P-P2. Teacher forcing raises about 0.5 BLEU improvements and greedy search outperform the teacher forcing algorithm by nearly 1 BLEU point. We conclude that the probabilistic loss has its own advantage even when trained by the teacher forcing algorithm, and greedy search is effective in alleviating the exposure bias.

Notice that the greedy search strategy highly relies on the probabilistic loss and can not be conducted independently. Greedy search together with the word-level loss is very similar with the scheduled sampling(SS). However, SS is inconsistent with the word-level loss since the word-level loss requires strict alignment between hypothesis and reference, which can only be accomplished by the teacher forcing algorithm.

4.5 Correlation with Evaluation Metrics

In this section, we explore how the probabilistic objective correlates with the real evaluation metric. We randomly sample 100 pairs of sentences




Figure 3: P-GLEU and GLEU scores on 100 pairs of sentences.

from the training set and compute their P-GLEU and GLEU scores (Wu et al. (2016) indicates that GLEU have better performance in the sentence-level evaluation than BLEU).

Directly computing the correlation between GLEU and P-GLEU gives the correlation coefficient 0.86, which indicates strong correlation. In addition, we draw the scatter diagram of the 100 pairs of sentences in Figure 3 with GLEU as x-axis and P-GLEU as y-axis. Figure 3 shows that P-GLEU correlates well with GLEU, suggesting that it is reasonable to directly train the NMT model with P-GLEU.

5 Conclusion

Word-level loss cannot evaluate the translation properly and suffers from the exposure bias, and sequence-level objectives are usually indifferntiable and require gradient estimation. We propose probabilistic sequence-level objectives based on n-gram matching, which relieve the dependence on gradient estimation and can directly train the NMT model. Experiment results show that our method significantly outperforms previous sequence-level training works and successfully alleviates the exposure bias through performing greedy search.

6 Acknowledgments

We thank the anonymous reviewers for their insightful comments. This work was supported by the National Natural Science Foundation of China (NSFC) under the project NO.61472428 and the project NO. 61662077.

References

- Dzmitry Bahdanau, Philemon Brakel, Kelvin Xu, Anirudh Goyal, Ryan Lowe, Joelle Pineau, Aaron Courville, and Yoshua Bengio. 2016. An actor-critic algorithm for sequence prediction. *arXiv preprint arXiv:1607.07086*.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Samy Bengio, Oriol Vinyals, Navdeep Jaitly, and Noam Shazeer. 2015. Scheduled sampling for sequence prediction with recurrent neural networks. In *Advances in Neural Information Processing Systems*, pages 1171–1179.
- Mia Xu Chen, Orhan Firat, Ankur Bapna, Melvin Johnson, Wolfgang Macherey, George Foster, Llion Jones, Niki Parmar, Mike Schuster, Zhifeng Chen, et al. 2018. The best of both worlds: Combining recent advances in neural machine translation. *arXiv preprint arXiv:1804.09849*.
- Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*.
- George Doddington. 2002. Automatic evaluation of machine translation quality using n-gram co-occurrence statistics. In *Proceedings of the second international conference on Human Language Technology Research*, pages 138–145. Morgan Kaufmann Publishers Inc.
- Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N Dauphin. 2017. Convolutional sequence to sequence learning. *arXiv preprint arXiv:1705.03122*.
- Jiatao Gu, Kyunghyun Cho, and Victor OK Li. 2017. Trainable greedy decoding for neural machine translation. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1968–1978.
- Hany Hassan, Anthony Aue, Chang Chen, Vishal Chowdhary, Jonathan Clark, Christian Federmann, Xuedong Huang, Marcin Junczys-Dowmunt, William Lewis, Mu Li, et al. 2018. Achieving human parity on automatic chinese to english news translation. *arXiv preprint arXiv:1803.05567*.
- 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*, pages 820–828.
- Nal Kalchbrenner and Phil Blunsom. 2013. Recurrent continuous translation models. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1700–1709.
- Yoon Kim and Alexander M Rush. 2016. Sequence-level knowledge distillation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1317–1327.
- Alex M Lamb, Anirudh Goyal ALIAS PARTH GOYAL, Ying Zhang, Saizheng Zhang, Aaron C Courville, and Yoshua Bengio. 2016. Professor forcing: A new algorithm for training recurrent networks. In *Advances In Neural Information Processing Systems*, pages 4601–4609.
- Guillaume Lample, Myle Ott, Alexis Conneau, Ludovic Denoyer, and Marc’Aurelio Ranzato. 2018. Phrase-based & neural unsupervised machine translation. *arXiv preprint arXiv:1804.07755*.
- Shuming Ma, Xu Sun, Yizhong Wang, and Junyang Lin. 2018. Bag-of-words as target for neural machine translation. *arXiv preprint arXiv:1805.04871*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting on association for computational linguistics*, pages 311–318. Association for Computational Linguistics.
- Marc’Aurelio Ranzato, Sumit Chopra, Michael Auli, and Wojciech Zaremba. 2015. Sequence level training with recurrent neural networks. *arXiv preprint arXiv:1511.06732*.
- Shiqi Shen, Yong Cheng, Zhongjun He, Wei He, Hua Wu, Maosong Sun, and Yang Liu. 2016. Minimum risk training for neural machine translation. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 1683–1692.
- Matthew Snover, Bonnie Dorr, Richard Schwartz, Linnea Micciulla, and John Makhoul. 2006. A study of translation edit rate with targeted human annotation. In *Proceedings of association for machine translation in the Americas*, volume 200.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pages 3104–3112.
- Richard S Sutton, David A McAllester, Satinder P Singh, and Yishay Mansour. 2000. Policy gradient methods for reinforcement learning with function approximation. In *Advances in neural information processing systems*, pages 1057–1063.
- 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 6000–6010.
- Arun Venkatraman, Martial Hebert, and J Andrew Bagnell. 2015. Improving multi-step prediction of learned time series models. In *AAAI*, pages 3024–3030.

- Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. In *Reinforcement Learning*, pages 5–32. Springer.
- Ronald J Williams and David Zipser. 1989. A learning algorithm for continually running fully recurrent neural networks. *Neural computation*, 1(2):270–280.
- Lijun Wu, Yingce Xia, Li Zhao, Fei Tian, Tao Qin, Jianhuang Lai, and Tie-Yan Liu. 2017. Adversarial neural machine translation. *arXiv preprint arXiv:1704.06933*.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*.
- Zhen Yang, Wei Chen, Feng Wang, and Bo Xu. 2017. Improving neural machine translation with conditional sequence generative adversarial nets. *arXiv preprint arXiv:1703.04887*.
- Matthew D Zeiler. 2012. Adadelta: an adaptive learning rate method. *arXiv preprint arXiv:1212.5701*.