

# Low-Resource Machine Translation based on Asynchronous Dynamic Programming

Xiaoning Jia, Hongxu Hou\*, Nier Wu, Haoran Li, Xin Chang

College of Computer Science-college of Software,  
Inner Mongolia University,China / Daxue West Road, Saihan District, Hohhot  
jiaxning@163.com, cshhx@imu.edu.cn, wunier04@126.com  
lihaoran009@126.com, changxin152@gmail.com

## Abstract

Reinforcement learning has been proved to be effective in handling low resource machine translation tasks and different sampling methods of reinforcement learning affect the performance of the model. The reward for generating translation is determined by the scalability and iteration of the sampling strategy, so it is difficult for the model to achieve bias-variance trade-off. Therefore, according to the poor ability of the model to analyze the structure of the sequence in low-resource tasks, this paper proposes a neural machine translation model parameter optimization method for asynchronous dynamic programming training strategies. In view of the experience priority situation under the current strategy, each selective sampling experience not only improves the value of the experience state, but also avoids the high computational resource consumption inherent in traditional valuation methods (such as dynamic programming). We verify the Mongolian-Chinese and Uyghur-Chinese tasks on CCMT2019. The result shows that our method has improved the quality of low-resource neural machine translation model compared with general reinforcement learning methods, which fully demonstrates the effectiveness of our method.

## 1 Introduction

Traditional machine translation (Ahmadnia and Dorr, 2019) uses cross entropy to measure the entropy of the generated token and reference token during the translation process. Its training process is usually to maximize the logarithmic likelihood of each token in the target sentence by taking the source sentence and the translated target token as input. This training method is called maximum likelihood estimation (MLE). Although it is easy to infer, the objective function at the token level in the training is not consistent with the evaluation indicators at the sequence level such as BLEU (Papineni et al., 2002) and the model uses a given real sample word as input to predict the next word. However, since there is no real distribution in the inference stage, the model can only generate the whole sequence by predicting one word at a time based on its own prediction output. The inconsistent methods of the prediction and inference stages are easy to cause the accumulation of errors. In order to solve the problem of inconsistency and accumulation of errors, a Minimum Risk Training (MRT) (Shen et al., 2016) algorithm is proposed to alleviate the problem. This training algorithm introduced evaluation index as the loss function, which aims to minimize the expected loss of training data and combines sentence-level BLEU into the loss function, so that the maximum likelihood estimation result of the machine translation model can be significantly improved. However, this function is not necessarily differentiable, which hinders the model from applying the BLEU score directly to the training process.

Aiming at the above problems, reinforcement learning (RL) (Ranzato et al., 2016) method has been applied to target optimization at the sequence level. This method directly optimises the BLEU value through Reinforcement training, which is a successful attempt of RL in strategy optimization at the sequence level. BR-CSGAN (Yang et al., 2018) is the first application of adversarial learning training to machine translation. The idea of reinforcement learning is cited in the generative adversarial network, which enables the neural machine translation model to obtain the BLEU value corresponding to each

possible sequence by the bootstrap method, and use this as the basis for optimizing the model parameters. In order to stabilize the training, the loss gradient of the reward is obtained by the reinforcement learning strategy through backpropagation, and finally the translation effect is steadily improved, it uses Monte Carlo for sampling and the sampling efficiency is improved to a certain extent. But due to the independent scalability of the Monte Carlo method, the sampled words are relatively independent of each other and the estimated state value does not depend on the previous state value, thereby reducing the learning accuracy.

In a study (Wu et al., 2018b), it is the first time that reinforcement learning methods were applied to machine translation, and reward shaping is used to improve the effectiveness of monolingual data in translation quality. Nevertheless, the training strategy using reinforcement learning is still restricted by the setting of rewards and the efficiency of sample sampling. As mentioned above, the traditional Monte Carlo method has independent malleability. Compared with the Monte Carlo method, the state value of each iteration of the agent training is determined by the previous state value and it is difficult to achieve the bias-variance trade-off. In order to achieve the trade-off, we hope that the model has the characteristics of low bias and low variance. However, the most serious challenge in the field of deep learning is the problem of high variance. The general solutions to solve the high variance problems are as follows: reduce the complexity of the model, reduce dimensionality of the data and denoise, increase the number of samples, use validation sets and regularization methods. Because of the scarcity of low-resource corpus data, our method involves the use of validation sets and regularization methods to reduce variance. That is, use the validation set to observe the generalization of the training parameters and update the new state through the latest K-step state difference. The Dynamic Programming method further solves the problem of correlation between words, but each updated action reward needs to refer to the previous round of experience, so all the learned experience will be stored. Based on the above analysis, we combined with the Dynamic Programming to propose an asynchronous training strategy suitable for the training of low-resource neural machine translation model. The bias-variance trade-off is achieved by selectively sampling the experience each time through the experience priority.

In this paper, we used a reinforcement learning strategy combined with Asynchronous Dynamic Programming (ADP) to train machine translation models. The Prioritised Sweeping (PS) algorithm is used to determine the experience priority. The higher the priority, the higher the priority. Ultimately, we maximize the priority to generate rewards to update the strategy and use the ADP method for sampling. Compared with the Monte Carlo method, we can pay more attention to the estimation Q of non-terminal state. Compared with the general Dynamic Programming method, the K-step update method avoid the calculation disaster and the real-time supervision of model training can also avoid the accumulation of errors. It improves the accuracy rate and effectively alleviates the sparse reward problem and has a higher universal computing capability.

We map the reinforcement learning algorithm to machine translation according to the Markov characteristics of reinforcement learning. Similar to the traditional machine translation model, we input the source language sentence into the translation model to generate the target sentence. Here we use the LSTM (Cui et al., 2016) and Transformer (Vaswani et al., 2017) frameworks of the machine translation model based on soft attention. The second chapter of this article is about the introduction of background knowledge. The third chapter is divided into framework, asynchronous strategy and training. The fourth chapter is the experiment and analysis part, and the fifth chapter is the conclusion.

## 2 Background

### 2.1 Neural machine translation

The typical Neural machine translation (NMT) (Wu et al., 2018a) model is an encoder-decoder framework based on attention mechanism and is usually using a cross-entropy function for training. The encoder first maps the source sentence to  $X = (x_1, \dots, x_n)$  vector form and then the decoder converts the set of vectors to generate the corresponding target sentence  $Y = (y_1, \dots, y_n)$ . And mark the words of the target statement one by one. In the time step  $t$ , the model maximizes the likelihood of the current time step  $t$  by using the source statement and the target statement before the time step  $t$  as input in the

training stage. The goal of the translation model is to maximize the maximum likelihood function, MLE is usually used to optimize the model:

$$L_{MLE} = \sum_{i=1}^S \log p(\hat{y}^i | x^i) = \sum_{i=1}^S \sum_{t=1}^m \log p(\hat{y}_t^i | \hat{y}_1^i \dots \hat{y}_{t-1}^i, x^i) \quad (1)$$

Where  $S$  represents the number of sentences in the training set,  $m$  is the length of sentence  $\hat{y}^i$ .

Among all the frameworks based on the encoder-decoder model, the recently popular Transformer architecture has achieved the best translation quality so far. The main difference between Transformer and the previous RNNSearch (Bahdanau et al., 2015) or Conv2Seq (Gehring et al., 2017) is that it completely relies on the self-attention mechanism (Yun et al., 2020), using its own structure to calculate the representation of the source and target sentences without recursion or convolution operations.

## 2.2 RL based NMT

Reinforcement learning solves the decision-making problem in complex state space by mapping the environment state to the action process to obtain the maximum cumulative reward and provides a new idea for natural language processing tasks. It reduces the gap between NMT training and inference by directly optimizing evaluation indicators (such as BLEU) during training. Specifically, the NMT model can be regarded as an agent interacting with the environment and the parameter definition strategy is the conditional probability  $p(y_t | x, y_{<t})$ . The agent will select a word from the vocabulary according to the strategy, that is, a candidate word. Once the agent generates a complete sequence, it will get a final reward. Reward is introduced to encourage the model to produce an output with a higher reward. In practice, the reward in machine translation is the BLEU value. Obtaining a higher reward means having a higher BLEU score. Here we denote the sentence-level reward as  $R$ , where  $\hat{y}$  represents the prediction result of the model,  $y^i$  represents the reference translation of sentence  $i$ ,  $R(\hat{y}, y^i)$  represents the reward value of sentence  $i$  and by sampling  $y$  from the strategy  $p(\hat{y} | x^i)$ . At this time, the goal of reinforcement learning training is to maximize the expected reward:

$$L_{RL} = \sum_{i=1}^S R(\hat{y}^i, y^i), \hat{y}^i \sim p(\hat{y} | x^i), \forall i \in [S]. \quad (2)$$

## 3 Approach

In this section, we describe in detail the asynchronous strategy of NMT model training, apply the reinforcement learning algorithm to the neural machine translation model, use the ADP strategy to train the gradient and update the network with parameterization.

### 3.1 Priority acquisition of experience

We abstract machine translation into a Markov process, where  $s$  represents the network hidden layer of a specific state of the model,  $a$  represents the model decoding prediction,  $P$  represents the state transition probability,  $reward$  represents the BLEU value reward that can be obtained when transferring from one state  $s$  to another state  $s'$ ,  $strategy \pi$  represents the probability distribution of any action that the agent may take next in a state  $s$ , it can also be said that our strategy is to directly select a certain action next in each state. In any case, the reinforcement learning method aims to learn an efficient strategy from the environment to enable agents to have certain model generalization. We can continuously improve our strategy so that we can finally get the maximum cumulative reward. In order to ensure the consistency and rationality of sampling in the training and inference stages and to avoid the problem of strictly traversing each state in Dynamic Programming. Different from the previous work, this paper uses an ADP to sampling, namely the important state priority update(PS) algorithm. The algorithm learns strategy during sampling and reinforcement training. Bellman error is used to calculate the experience priority based on the experience learned from the strategy. The experience that will be collected in the next round of training is determined by priority.




Figure 1: The illustration of the experience priority.

As shown in formula (3), we use the idea of Bellman Error to determine the experience priority situation and use the absolute value of the difference between the state value of the previous round and the current state value to judge. If the Bellman Error is large, it means that the state is unstable and it is necessary to give priority to updating. The experience priority is shown in Figure 1. We use translation prediction probability *action* instead of state value to judge and record the Bellman Error of each *action* in the first round of epoch and the second round of epoch as  $BE_{1\dots k}^{(1)}$ , and so on. After that,  $BE_{1\dots k}^{(1)}$  is stored in the buffer for the optimized call of the asynchronous strategy later. Compared with the sumtree method to extract data priority, the difference data obtained by our method is more dependent on the information content of the previous stage and has more correlation between contexts.

$$\text{Bellman Error} = \left| \max_{a \in A} (R_s^a + \gamma \sum_{s' \in S} P_{ss'}^a v(s') - v(s)) \right| \quad (3)$$

We store the sorted experience priorities in a specific experience pool, so that the subsequent asynchronous training strategy can extract experience from the experience pool in turn to optimize the model. This method is called experience playback technology. Experience playback is to store the experience of past agents in an experience playback pool and then repeatedly sample from the experience playback pool to optimize the strategy. By using past experience many times to train the current strategy, in order to improve the sample efficiency and training stability of the reinforcement learning algorithm.

### 3.2 Asynchronous Strategy

The action space of machine translation tasks based on reinforcement learning is huge and discrete and its size is the whole word list. In the previous algorithm, we will completely store and update all the learned experiences in each iteration, thus requiring more program resources. Therefore, we use the important state priority update algorithm in ADP, and trains the machine translation model through the reinforcement learning method. The experience determined by Bellman Error is stored in the experience pool, and the experience *value* with large error is extracted from the experience pool through the K-step update method for optimization. K-step to select the difference of all states accumulated and  $R(\sum_k)$  as the final reward. According to the reward guidance next state  $State_{1\dots k}^{(2)}$  update prediction and update until repeated Bellman Error tends to stabilize. This method not only improves the value of the experience state, but also avoids the problem of high computing resource consumption inherent in traditional Dynamic Programming. As shown in Figure 2. We verified the effectiveness of ADP in the experimental part.

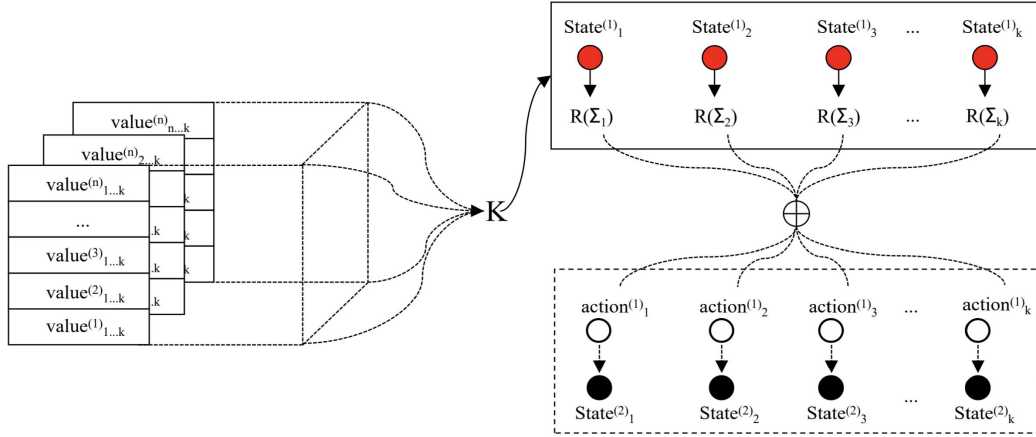


Figure 2: The illustration of the Asynchronous Dynamic Programming.

We know that establishing an appropriate reward mechanism is crucial to obtaining high-quality translation. In the process of machine translation, the NMT system acts as an agent of reinforcement learning, obtains the status information of the current moment through continuous interaction with the environment. After each action prediction probability step, the translation model uses the difference between the translation prediction probability of the previous round and the current time as the empirical priority evaluation standard in the current environment. In the  $K$ -step, the optimization is updated to the model convergence by selecting the priority to maximize.

Compared with the general Dynamic Programming method, our method uses ADP method for sampling and uses the  $K$ -step update method to select the priority to maximize the update to avoid the consumption of computing resources. Moreover, the real-time supervision of model training also avoids error accumulation, improves the accuracy rate while effectively alleviating the problem of reward sparseness and has higher ubiquitous computing capabilities. Drawing on Bellman Error idea, the sum of rewards for all updates in the  $K$ -step is used as the final feedback. The formula is as follows:

$$R(\hat{y}_k) = \sum_{k=1}^K Bellman_{1\dots k} value \quad (4)$$

### 3.3 Training

In this paper, ADP strategy is used as a sampling strategy for low-resource reinforcement learning to improve the translation model. The important state with higher experience priority is iteratively updated within  $K$ -step, the operation is selected by using the current strategy and the reward is determined by comparing the value of the last moment with the current moment. In the training process, the best sequence comes from the distribution of standard translations. The goal of the training is to find the model parameters that maximize the expected reward, so the objective function is set to:

$$L_K = \sum_{i=1}^S R(\hat{y}_k^i, y^i), \hat{y}_k^i \sim p(\hat{y}|x^i), \forall_i \in [S]. \quad (5)$$

Finally, in order to further stabilize the RL training process and alleviate the large variance caused by reinforcement learning, we combine the MLE training target with the RL target. The cross-entropy loss function of traditional machine translation is retained in the loss function and it is linearly combined with the reinforcement learning training target. The loss function after mixing is:

$$L_{COM} = \lambda \times L_{MLE} + (1 - \lambda) \times L_K \quad (6)$$

where  $\lambda$  is the hyperparameter to control the balance between MLE and RL.  $L_{COM}$  is the strategy to stabilize RL training progress.



## 4 Experiment and Analysis

We verified the effectiveness of our method on two low-resource NMT tasks, including Mongolian-Chinese (Mo-Zh) and Uyghur-Chinese (Ug-Zh).

### 4.1 Datasets and preprocessing

For the Mongolian-Chinese (Mo-Zh) experiment, the data resources used to train the translation model and the translation quality evaluation model come from the Mongolian Intelligent Information Processing Key Laboratory of the Inner Mongolia Autonomous Region. Nearly 250,000 bilingual parallel corpora provided by CCMT2019 are used to train translation models. And 1,000 parallel corpora in CWMT2017 are used as the test set and 1,000 as the verification set. For Uyghur-Chinese (Ug-Zh), 330,000 Uyghur-Chinese machine translation data published by CWMT2017 are used for experiments and 1,000 pieces of data were selected as the test set and the verification set. To avoid the model allocating too much training time to long sentences, all sentence pairs that are longer than 50 at the source or target end will be discarded.

The Stanford Word Segmentation Tool<sup>0</sup> is used to segment the words in Chinese. Because Mongolian and Uyghur are low-resource languages and too many low-frequency words in the vocabulary. So Byte Pair Encoder(BPE) (Sennrich et al., 2016) is used to process to alleviate the data sparse problem of low-resource languages to a certain extent in this paper.

### 4.2 Setting

Our model is improved based on the self-attention Transformer<sup>1</sup> model, use the traditional LSTM<sup>2</sup> and the latest Transformer model as the baseline model. For LSTM, we follow the basic settings and embed the word with a dimension of 512, the learning rate to 0.1, and the beam search size to 4. The settings for Transformer are exactly the same as the original paper, dropout=0.1, word embedding dimension is 512, header count is 8, encoder and decoder both have 6-layer stack. The quality evaluation is based on the sentence level evaluation standard BLEU. The gradient optimization algorithm uses Adam (Kingma and Ba, 2015). During reinforcement learning training, MLE model is used to initialize parameter. The learning rate is set to 0.0001 and the beam search width is set to 6.

### 4.3 Main results and Analysis

#### Comparison results of benchmark experiments

In order to verify the translation performance of the model, we conduct low-resource machine translation comparison experiments with several common representative neural machine translation models under the same hardware conditions and corpus scale. Calculating the BLEU value of each model on the test set separately. It can be seen from the table that our method can surpass the baseline model in both Mongolian-Chinese and Uyghur-Chinese translation tasks. As shown in Table 1.

As can be seen from the translation results, the baseline method of Mongolian-Chinese Machine Translation Transformer is 27.58, Transformer + RL<sup>3</sup> translation method is 31.37, up 3.79. The LSTM baseline method is 25.28, the LSTM + RL<sup>4</sup> is 30.95, and the addition of RL is also 5.67 higher than the baseline, indicating that the introduction of reinforcement learning plays an important role in improving

Model	Mo-Zh	Ug-Zh
Transformer	27.58	36.83
Transformer+RL	31.37	39.71
<b>Our Method</b>	<b>34.52</b>	<b>44.35</b>
LSTM	25.28	35.67
LSTM+RL	30.95	38.00
<b>Our Method</b>	<b>33.27</b>	<b>42.97</b>

Table 1: The BLEU scores of different NMT systems on Mongolian-Chinese and Uyghur-Chinese.

<sup>0</sup><https://nlp.stanford.edu/software/segmenter.html>

<sup>1</sup><https://github.com/tensorflow/tensor2tensor>

<sup>2</sup><https://github.com/xwgeng/RNNSearch>

<sup>3</sup><https://github.com/apeterswu/RL4NMT>

<sup>4</sup><https://github.com/facebookarchive/MIXER>






Figure 4: The length of source sentence.

From the figure 3, we can see that our method is closer to the reference translation than the traditional Transformer+RL translation model. And the translation obtained on the two low-resource corpora of Mongolian-Chinese and Uyghur-Chinese is relatively accurate. The fluency and fidelity of the translation conform to the target language specification, and the quality of the translation is significantly better than its baseline model, which proves that our method can effectively improve the performance of the low-resource neural machine translation model.

**Results of sentence length analysis** We conduct a set of typical experiments using Transformer on the Mongolian-Chinese task to verify the performance of this method on long sentences. We divided the development set data and test set data of the Mongolian-Chinese task according to the sentence length. Figure 4 shows the BLEU scores for different sentence lengths. No matter on LSTM, Transformer or compared to joining RL with the best baseline performance, our work have outstanding behaviors continuously. It is due to our method trains machine translation models based on ADP strategies and increases the ability to analyze low-resource sentence structures through parameter optimization, which makes our method perform better on both long and short sentences.

## 5 Conclusion

Since different reinforcement learning sampling methods affect the performance of the model and in order to alleviate the quality problem of generated translations in low-resource machine translation tasks, this paper proposes an Asynchronous Dynamic sampling strategy based on reinforcement learning. This strategy selectively samples and stores each time according to the experience priority situation, thereby reducing the consumption of computing resources and improving the value of the experience state. Experimental results show that this method can effectively improve the performance of low-resource machine translation. In the next step, we will combine more low-resource language features to explore more suitable for low-resource machine translation, so as to further improve the performance of low-resource machine translation systems.

## References

- Benyamin Ahmadnia and Bonnie J. Dorr. 2019. Enhancing phrase-based statistical machine translation by learning phrase representations using long short-term memory network. In Ruslan Mitkov and Galia Angelova, editors, *Proceedings of the International Conference on Recent Advances in Natural Language Processing, RANLP 2019, Varna, Bulgaria, September 2-4, 2019*, pages 25–32. INCOMA Ltd.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.



- Yiming Cui, Shijin Wang, and Jianfeng Li. 2016. LSTM neural reordering feature for statistical machine translation. In Kevin Knight, Ani Nenkova, and Owen Rambow, editors, *NAACL HLT 2016, The 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, San Diego California, USA, June 12-17, 2016*, pages 977–982. The Association for Computational Linguistics.
- Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N. Dauphin. 2017. Convolutional sequence to sequence learning. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, volume 70 of *Proceedings of Machine Learning Research*, pages 1243–1252. PMLR.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- 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 of the Association for Computational Linguistics, July 6-12, 2002, Philadelphia, PA, USA*, pages 311–318. ACL.
- Marc’Aurelio Ranzato, Sumit Chopra, Michael Auli, and Wojciech Zaremba. 2016. Sequence level training with recurrent neural networks. In Yoshua Bengio and Yann LeCun, editors, *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. 2016. Neural machine translation of rare words with subword units. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers*. The Association for Computer Linguistics.
- 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, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers*. The Association for Computer Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008.
- Lijun Wu, Xu Tan, Di He, Fei Tian, Tao Qin, Jianhuang Lai, and Tie-Yan Liu. 2018a. Beyond error propagation in neural machine translation: Characteristics of language also matter. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 3602–3611.
- Lijun Wu, Fei Tian, Tao Qin, Jianhuang Lai, and Tie-Yan Liu. 2018b. A study of reinforcement learning for neural machine translation. In Ellen Riloff, David Chiang, Julia Hockenmaier, and Jun’ichi Tsujii, editors, *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 3612–3621. Association for Computational Linguistics.
- Zhen Yang, Wei Chen, Feng Wang, and Bo Xu. 2018. Improving neural machine translation with conditional sequence generative adversarial nets. In Marilyn A. Walker, Heng Ji, and Amanda Stent, editors, *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2018, New Orleans, Louisiana, USA, June 1-6, 2018, Volume 1 (Long Papers)*, pages 1346–1355. Association for Computational Linguistics.
- Hyeon-gu Yun, Yongkeun Hwang, and Kyomin Jung. 2020. Improving context-aware neural machine translation using self-attentive sentence embedding. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 9498–9506. AAAI Press.