

Neural Machine Translation with Adequacy-Oriented Learning

Xiang Kong
Carnegie Mellon University
xiangk@andrew.cmu.edu

Zhaopeng Tu*
Tencent AI Lab
zptu@tencent.com

Shuming Shi
Tencent AI Lab
shumingshi@tencent.com

Eduard Hovy
Carnegie Mellon University
hovy@cs.cmu.edu

Tong Zhang
Tencent AI Lab
bradymzhang@tencent.com

Abstract

Although Neural Machine Translation (NMT) models have advanced state-of-the-art performance in machine translation, they face problems like the inadequate translation. We attribute this to that the standard Maximum Likelihood Estimation (MLE) cannot judge the real translation quality due to its several limitations. In this work, we propose an adequacy-oriented learning mechanism for NMT by casting translation as a stochastic policy in Reinforcement Learning (RL), where the reward is estimated by explicitly measuring translation adequacy. Benefiting from the sequence-level training of RL strategy and a more accurate reward designed specifically for translation, our model outperforms multiple strong baselines, including (1) standard and coverage-augmented attention models with MLE-based training, and (2) advanced reinforcement and adversarial training strategies with rewards based on both word-level BLEU and character-level CHRF3. Quantitative and qualitative analyses on different language pairs and NMT architectures demonstrate the effectiveness and universality of the proposed approach.

Introduction

During the past several years, rapid progress has been made in the field of Neural Machine Translation (NMT) (Kalchbrenner and Blunsom 2013; Sutskever, Vinyals, and Le 2014; Bahdanau, Cho, and Bengio 2015; Gehring et al. 2017; Wu et al. 2016; Vaswani et al. 2017).

Although NMT models have advanced the community, they still face inadequate translation problems: one or multiple parts of the input sentence are not translated (Tu et al. 2016). We attribute this problem to the lack of the mechanism to guarantee the generated translation being as sufficient as human translation. NMT models are generally trained in an end-to-end manner to maximize the likelihood of the output sentence. Maximum Likelihood Estimation (MLE), however, could not judge the real quality of generated translation due to its several limitations

1. *Exposure bias* (Ranzato et al. 2016): The models are trained on the groundtruth data distribution, but at test time are used to generate target words based on previous model predictions, which can be erroneous;

* Zhaopeng Tu is the corresponding author. Work was mainly done when Xiang Kong was interning at Tencent AI Lab.
Copyright © 2019, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

2. *Word-level loss* (Shen et al. 2016): Likelihood is defined at word-level, which might hardly correlate well with sequence-level evaluation metrics like BLEU.
3. *Focusing more on fluency than adequacy* (Tu et al. 2017): Likelihood does not measure how well the complete source information is transformed to the target side, thus does not correlate well with translation adequacy. Adequacy metric is regularly employed to assess the translation quality in practice.

Some recent work partially alleviates one or two of the above problems with advanced training strategies. For example, the first two problems are tackled by sequence level training using the REINFORCE algorithm (Ranzato et al. 2016; Bahdanau et al. 2017), minimum risk training (Shen et al. 2016), beam search optimization (Wiseman and Rush 2016) or adversarial learning (Wu et al. 2017; Yang et al. 2018). The last problem can be alleviated by introducing an auxiliary reconstruction-based training objective to measure translation adequacy (Tu et al. 2017).

In this work, we aim to fully solve all the three problems in a unified framework. Specifically, we model the translation as a stochastic policy in Reinforcement Learning (RL) and directly perform gradient policy update. The RL reward is estimated on a complete sequence produced by the NMT model, which is able to correlate well with a sequence-level task-specific metric. To explicitly measure translation adequacy, we propose a novel metric called *Coverage Difference Ratio* (CDR) which is calculated by counting how many source words are under-translated via directly comparing generated translation with human translation. Benefiting from the sequence-level training of RL strategy and a more accurate reward designed specifically for translation, the proposed approach is able to alleviate all the aforementioned limitations of MLE-based training.

We conduct experiments on Chinese \Rightarrow English and German \Leftrightarrow English translation tasks, using both the RNN-based NMT model (Bahdanau, Cho, and Bengio 2015) and the recently proposed TRANSFORMER (Vaswani et al. 2017). The consistent improvements across language pairs and NMT architectures demonstrate the effectiveness and universality of the proposed approach. The proposed adequacy-oriented learning improves translation performance not only over a standard attention model, but

also over a coverage-augmented attention model (Tu et al. 2016) that alleviates the inadequate translation problem at the word-level. In addition, the proposed metric – CDR score, consistently outperforms the commonly-used word-level BLEU (Papineni et al. 2002) and character-level CHR3 (Popović 2015) scores in both the reinforcement learning and adversarial learning frameworks, indicating the superiority and necessity of an adequacy-oriented metric in training effective NMT models.

Background

Neural Machine Translation (NMT) is an end-to-end structure which could directly model the translation probability between a source sentence $\mathbf{x} = x_1, x_2, \dots, x_J$ and a target sentence $\mathbf{y} = y_1, y_2, \dots, y_I$ word by word:

$$P(\mathbf{y}|\mathbf{x}) = \prod_{i=1}^I P(y_i | \mathbf{y}_{<i}, \mathbf{x}; \theta) \quad (1)$$

where $\mathbf{y}_{<i}$ is the partial translation before decoding step i and θ is parameters of the NMT. The probability of generating the i -th word $P(y_i | \mathbf{y}_{<i}, \mathbf{x}; \theta)$ is calculated by

$$P(y_i | \mathbf{y}_{<i}, \mathbf{x}; \theta) \propto \exp \{f(y_{i-1}, \mathbf{s}_i, \mathbf{c}_i; \theta)\} \quad (2)$$

where \mathbf{s}_i is the i -th hidden state of the decoder and $f(\cdot)$ is a non-linear activation function of the decoder state. \mathbf{c}_i is a distinct source representation for time i , calculated as a weighted sum of the source annotations: $\mathbf{c}_i = \sum_{j=1}^J \alpha_{i,j} \cdot \mathbf{h}_j$, where \mathbf{h}_j is the annotation of x_j from an encoder, and its weight $\alpha_{i,j}$ is computed by

$$\alpha_{i,j} = \frac{\exp(e_{i,j})}{\sum_{j'=1}^J \exp(e_{i,j'})} \text{ with } e_{i,j} = a(\mathbf{s}_{i-1}, \mathbf{h}_j) \quad (3)$$

where $a(\cdot)$ is an *attention model* that scores how well y_i and \mathbf{h}_j (i.e., x_j) match. The encoder and decoder can be implemented as Recurrent Neural Network (RNN) (Bahdanau, Cho, and Bengio 2015), Convolutional Neural Network (CNN) (Gehring et al. 2017), or Self-Attention Network (SAN) (Vaswani et al. 2017).

The parameters of the NMT θ are trained to maximize the likelihood of training instances $\{[x^n, y^n]\}_{n=1}^N$:

$$L(\theta) = \arg \max_{\theta} \sum_{n=1}^N \log P(y^n | x^n; \theta) \quad (4)$$

Although likelihood is a widely-used training objective for its simpleness and effectiveness, it has several aforementioned limitations including *exposure bias* (Ranzato et al. 2016; Wiseman and Rush 2016), *word-level estimation* (Shen et al. 2016), and *focusing more on fluency than adequacy* (Tu et al. 2017).

Approach

Intuition

In this work, we try to solve the three problems mentioned above in a unified framework. Our objective is three-fold:

Input	二十六岁男司机被困车内，由到场的消防员救出。
REF	The 26-year-old driver was trapped inside the car after the impact and was rescued by the firemen.
NMT	The 26-year-old driver was trapped inside the vehicle.

(a) Example of human (REF) and generated (NMT) translations.

	二十六岁	男	司机	被困	车内	，	由	到场的	消防员	救出。
REF										
NMT										

(b) Examples of covered source words (i.e. shadow boxes).

Figure 1: An example to illustrate coverage difference ratio.

1. We solve the exposure bias problem by modeling the translation as a stochastic policy in reinforcement learning (RL) and directly performing policy gradient update.
2. The RL reward is estimated on a complete sequence, which correlates well with either sequence-level BLEU or a more adequacy-oriented metric, as described below.
3. We design a sequence-level metric – *Coverage Difference Ratio* (CDR) – to explicitly measure translation adequacy which focuses on the commonly-cited weaknesses of NMT models: producing fluent yet inadequate translations. We expect that the model can benefit from linguistic insights that correlate well with human intuitions.

Coverage Difference Ratio (CDR) We measure translation adequacy by the number of under-translated words via comparing generated translation with human translation. We take an example to illustrate how to measure translation adequacy in terms of coverage difference ratio. Figure 1(a) shows one inadequate translation. Following (Luong, Pham, and Manning 2015; Tu et al. 2016), we extract only one-to-one alignments (hard alignments) by selecting the source word with the highest alignment for each target word from the word alignments produced by NMT models.¹ A source word is considered to be translated when it is covered by the hard alignments, as shown in Figure 1(b). Comparing source words covered by generated translation with those covered by human translation, we can find that the two sets are very different for inadequate translation. Specifically, the difference generally lies in the untranslated source words that cause inadequate translation problem, indicating that coverage difference ratio is a good way to measure the adequacy of generated translation.

¹For generated translations, we directly use the attention probability distributions from decoding procedure; for human translations, we obtain attention distributions by force decoding the target sentences with the same NMT model.

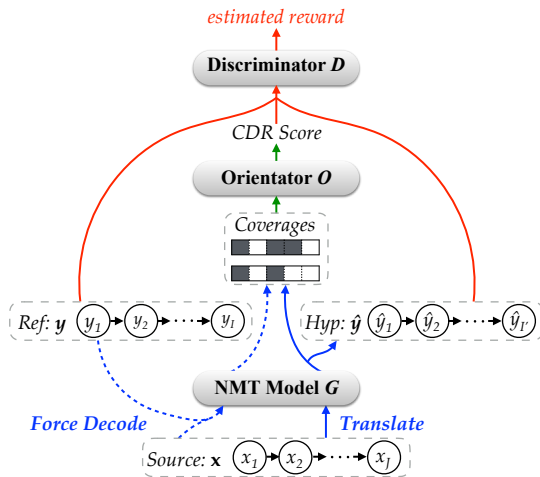


Figure 2: Architecture of adequacy-oriented NMT. The newly added orientator O reads coverages of generated and human translations to generate a CDR score for each generated translation, which guides the discriminator D to differentiate good generated translations from bad ones.

Formally, we calculate the CDR score of a given generated translation \hat{y} by

$$\text{CDR}(\hat{y}|\mathbf{y}, \mathbf{x}) = 1 - \frac{|\mathbf{C}_{ref} \setminus \mathbf{C}_{gen}|}{|\mathbf{C}_{ref}|} \quad (5)$$

where \mathbf{C}_{ref} and \mathbf{C}_{gen} is the set of source words covered by human translation and generated translation, respectively. $\mathbf{C}_{ref} \setminus \mathbf{C}_{gen}$ denotes the covered source words in \mathbf{C}_{ref} but not in \mathbf{C}_{gen} . We use \mathbf{C}_{ref} as the reference coverage to eliminate the effect of null-aligned source words which are not aligned to any target word. As seen, $\text{CDR}(\hat{y}|\mathbf{y}, \mathbf{x})$ is a number between 0 and 1, where 1 means “completely adequate translation” and 0 means “completely inadequate translation”. Taking Figure 1(b) as an example, the CDR score is $1 - (7 - 4)/7 = 0.57$.

Architecture

As shown in Figure 2, the proposed model consists of a *generator*, a *discriminator*, and an *orientator*.

Generator The generator G generates the translation \hat{y} conditioned on the input sentence \mathbf{x} . Because we need word alignments to calculate adequacy scores in terms of CDR, an attention-based NMT model is employed as the generator.

Orientator The orientator O reads the word alignments produced by NMT attention model when generating (or force decoding) the two translations and outputs an adequacy score for the generated translation in terms of the aforementioned CDR score. Then, the orientator is used to guide the discriminator to distinguish adequate translation from inadequate ones. Accordingly, adequate translations with higher CDR scores would contribute more to parameter tuning, as described in the following section.

Discriminator We employ a RNN-based discriminator to differentiate generated translation from human translation, given the input sentence. The discriminator reads the input sentence \mathbf{x} and its translation (either \mathbf{y} or $\hat{\mathbf{y}}$), and use two RNNs to summarize the two sentences individually. The concatenation of the two summarized representation vectors is fed into a fully-connected neural network.

Adequacy-Oriented Training

In order to train the system efficiently and effectively, we employ a periodical training strategy, which is commonly used in adversarial training (Goodfellow et al. 2014; Wu et al. 2017). Specifically, we optimize two networks with two objective functions and periodically freeze the parameters of each network during training.

Train Generator and Freeze Discriminator Following Wu et al. (2017), we use the REINFORCE algorithm (Williams 1992) to back-propagate the error signals from D to G , given the discretely generated $\hat{\mathbf{y}}$ from G . The objective of the generator is to maximize the expected reward:

$$L = \mathbb{E}_{(\mathbf{x}, \hat{\mathbf{y}}) \in G_\theta} [D(\mathbf{x}, \hat{\mathbf{y}})] \quad (6)$$

whose gradient is

$$\nabla_\theta = E_{(\mathbf{x}, \hat{\mathbf{y}}) \in G_\theta} [D(\mathbf{x}, \hat{\mathbf{y}}) \nabla_\theta \log G_\theta(\hat{\mathbf{y}}|\mathbf{x})] \quad (7)$$

The gradient is approximated by a sample from G using the REINFORCE algorithm (Williams 1992):

$$\nabla_\theta \approx \hat{\nabla}_\theta = D(\mathbf{x}, \hat{\mathbf{y}}) \nabla_\theta \log G_\theta(\hat{\mathbf{y}}|\mathbf{x}) \quad (8)$$

where $\nabla_\theta \log G(\hat{\mathbf{y}}|\mathbf{x})$ is the standard NMT gradient which is calculated by the maximum likelihood estimation. Therefore, the final update function for the generator is:

$$\theta = \theta - \eta \hat{\nabla}_\theta \quad (9)$$

where the η is the learning rate. Based on the update function, when the $D(\mathbf{x}, \hat{\mathbf{y}})$ is large (*i.e.*, ideally, the generated translation $\hat{\mathbf{y}}$ has a high adequacy score), the larger reward the NMT model will get, and thus parameters are updated more based on the adequate training instance $(\mathbf{x}, \hat{\mathbf{y}})$.

Train Discriminator Oriented by Adequacy and Freeze Generator

Ideally, a good translation $\hat{\mathbf{y}}$ should be assigned a high adequacy score $D(\mathbf{x}, \hat{\mathbf{y}})$ and thus contribute more to updating the generator. Therefore, we expect the discriminator to not only differentiate generated translations from human translations but also distinguish bad generated translations from good ones. Therefore, a new objective of discriminator is to assign a precise score for each generated translation, which is consistent with their adequacy score:

$$\min_D |\text{CDR}(\hat{\mathbf{y}}|\mathbf{x}, \mathbf{y}) - D(\mathbf{x}, \hat{\mathbf{y}})|^2 \quad (10)$$

where $\text{CDR}(\hat{\mathbf{y}}|\mathbf{x}, \mathbf{y})$ is the coverage difference ratio of $\hat{\mathbf{y}}$. As seen, a well trained discriminator would assign a distinct score to each generated translation, which can better measure its adequacy.

Related Work

This work is related to modeling translation as policy gradient and adequacy modeling. For the former, we take minimum risk training, reinforcement learning and adversarial learning as representative strategies.

Minimum Risk Training In response to the exposure bias and word-level loss problems of MLE training, Shen et al. (2016) minimize the expected loss in terms of evaluation metrics on the training data. Our simplified model is analogous to their MRT model, if we directly use CDR as the reward to update parameters:

$$\hat{\nabla}_{\theta} = \text{CDR}(\hat{\mathbf{y}}|\mathbf{x}, \mathbf{y}) \nabla_{\theta} \log G_{\theta}(\hat{\mathbf{y}}|\mathbf{x}) \quad (11)$$

The simplified model differs in that (1) we use adequacy-oriented metric (*i.e.*, CDR) while they use sequence-level BLEU, and (2) we only need to sample one candidate to calculate reinforcement reward while they generate multiple samples to calculate the expected risk. In addition, our discriminator gives a smoother and dynamically-updated objective compared with directly using the adequacy-oriented metric, because the latter is highly sensitive to the slight coverage difference (Koehn and Knowles 2017).

Reinforcement Learning Recent work shows that maximum likelihood training could be sub-optimal due to the different conditions between training and test modes (Bengio et al. 2015; Ranzato et al. 2016). In order to address the exposure bias and the loss which does not operate at the sequence level, Ranzato et al. (2016) employ the REINFORCE algorithm (Williams 1992) to decide whether or not tokens from a sampled prediction could contribute to a high task-specific score (*e.g.*, BLEU). Bahdanau et al. (2017) use the actor-critic method from reinforcement learning to directly optimize a task-specific score.

Adversarial Learning Recently, adversarial learning (Goodfellow et al. 2014) has been successfully applied to neural machine translation (Wu et al. 2017; Yang et al. 2018; Cheng et al. 2018). In the adversarial framework, NMT models generally serve as the generator which defines the policy to generate the target sentence \mathbf{y} given the source sentence \mathbf{x} . A discriminator tries to distinguish the translation result $\hat{\mathbf{y}} = G(\mathbf{x})$ from the human-generated one \mathbf{y} , given the source sentence \mathbf{x} .

If we remove the orientator \mathbf{O} , our model is roll-backed to the adversarial NMT, and the training objective of the discriminator D is rewritten as

$$\max_D \{\log D(\mathbf{x}, \mathbf{y}) + \log(1 - D(\mathbf{x}, \hat{\mathbf{y}}))\} \quad (12)$$

The goal of the discriminator is try to maximize the likelihood of human translation $D(\mathbf{x}, \mathbf{y})$ to 1 and minimize that of generated translation $D(\mathbf{x}, \hat{\mathbf{y}})$ to 0.

As seen, the discriminator uses a binary classification by uniformly treating all generated translations as negative examples (*i.e.*, labeling “0”) and all human translations as positive examples (*i.e.*, labeling “1”), regardless of the quality of

the generated translations. However, intuitively, high-quality translations and low-quality translations should be treated differently by the discriminator, otherwise, inaccurate reward signals would be propagated back to the generator. In our proposed architecture, this problem can be alleviated by replacing the simple binary outputs with the more informative adequacy-oriented metric CDR, which is calculated by directly comparing generated and human translations.

Adequacy Modeling Inadequate translation problem is a commonly-cited weakness of NMT models (Tu et al. 2016). A number of recent efforts have explored ways to alleviate this problem. For example, Tu et al. (2016) and Mi et al. (2016) employ coverage vector as a lexical-level indicator to indicate whether a source word is translated or not. Zheng et al. (2018) and Meng et al. (2018) move one step further and directly model translated and untranslated source contents by operating on the attention context vector. He et al. (2017) use a prediction network to estimate the future cost of translating the uncovered source words. Our approach is complementary to theirs since they model the adequacy learning at the word-level inside the generator (*i.e.*, NMT models), while we model it at the sequence-level outside the generator. We take the representative coverage mechanism (Tu et al. 2016) as another stronger baseline model for its simplicity and efficiency, and experimental results show that our model can further improve performance.

In the context of adequacy-oriented training, Tu et al. (2017) introduce an auxiliary objective to measure the adequacy of translation candidates, which is calculated by reconstructing generated translations back to the original inputs. Benefiting from the flexible framework of reinforcement training, we are able to directly compare generated translations with human translations and define a more straightforward metric, *i.e.*, CDR to measure adequacy of generated sentences.

Experiments

Setup

We conduct experiments on the widely-used Chinese (Zh) \Rightarrow English (En) and German (De) \Leftrightarrow English (En) translation tasks. For Zh \Rightarrow En translation, the training corpus contains 1.25M sentence pairs extracted from LDC corpora. NIST 2002 (MT02) dataset is the validation set and the test data consists of NIST 2003 (MT03), NIST2004 (MT04), NIST 2005 (MT05) and NIST 2006(MT06). For De \Leftrightarrow En translation, to compare with the results reported by previous work (Shen et al. 2016; Bahdanau et al. 2017; Wu et al. 2017; Vaswani et al. 2017), we use both the IWSLT 2014 and WMT 2014 data. The former contains 153K sentence pairs and the latter consists of 4.56M sentence pairs. The 4-gram NIST BLEU score (Papineni et al. 2002) is used as the evaluation metric and sign-test (Collins, Koehn, and Kučerová 2005) is employed to test statistical significance.

For training all neural models, we set the vocabulary size to 30K for Zh \Rightarrow En, for IWSLT 2014 De \Rightarrow En, we follow the preprocessing procedure as used in Ranzato et al. (2016)

#	Model	# Para.	Speed	MT03	MT04	MT05	MT06	All
1	RNNSEARCH	86.7M	1.4K	36.00	38.26	35.88	35.98	36.76
2	+ MRT _{BLEU}	+0M	0.3K	37.32	39.41	36.78	37.22	37.92 [‡]
3	+ D _{CNN}	+0.23M	1.0K	37.11	38.84	35.97	37.36	37.54 [‡]
4	+ D _{RNN}	+0.17M	1.2K	36.96	38.92	36.13	37.54	37.59 [‡]
5	+ O _{BLEU}	+0M	1.1K	37.02	38.49	36.62	36.21	37.44 [†]
6	+ O _{CHRF3}	+0M	1.1K	37.91	39.80	36.57	35.95	37.75 [†]
7	+ O _{CDR}	+0M	1.0K	37.42	39.52	36.86	37.34	38.02 [‡]
8	+ D _{RNN} + O _{CDR}	+0.17M	0.8K	37.61	40.05	37.58	36.87	38.42 [‡]
9	RNNSEARCH-COVERAGE	+1.03M	1.2K	38.04	41.09	38.73	36.52	39.13
10	+ D _{RNN} + O _{CDR}	+1.20M	0.7K	38.62	41.98	39.39	37.42	39.81 [†]

Table 1: Evaluation of translation performance on Zh⇒En translation. “D” denotes discriminator and “O” denotes orientator. “MRT” indicates minimum risk training (Shen et al. 2016), and “D_{CNN}” indicates adversarial training with a CNN-based discriminator (Wu et al. 2017). “# Para.” denotes the number of parameters, and “Speed” denotes the training speed (words/second). “†” and “‡” indicate statistically significant difference ($p < 0.05$ and $p < 0.01$ respectively) from the corresponding baseline.

and for WMT 2014 En⇒De, preprocessing method described in Vaswani et al. (2017) is borrowed. We pre-train the discriminator on translation samples produced by the pre-trained generator. After that, the discriminator and the generator are trained together, and the generator is updated by the REINFORCE algorithm mentioned above. We also follow the training tips mentioned in Shen et al. (2016) and Wu et al. (2017). The hyper-parameter α which could control the sharpness of the generator distribution in our system is $1e-4$, which could also be regarded as a baseline to reduce the variance of the REINFORCE algorithm. We also randomly choose 50% minibatches trained with our objective function and the other with the MLE principle. In MRT training strategy (Shen et al. 2016), the sample size is 25, the hyper-parameter α is $5e-3$ and the loss function is negative smoothed sentence-level BLEU.

We validate our models on two representative model architectures, namely RNNSEARCH and TRANSFORMER. For the RNNSEARCH model, mini-batch size is 80, the word-embedding dimension is 620, and the hidden layer size is 1000. We use a neural coverage model for RNNSEARCH-COVERAGE and the dimensionality of coverage vector is 100. The baseline models are trained for 15 epochs, which are used as the initial generator in the proposed framework. For the TRANSFORMER model, we implement our proposed approach on top of an open source toolkit THMUT (Zhang et al. 2017). Configurations in Vaswani et al. (2017) are used to train the baseline models.

Chinese-English Translation Task

Table 1 lists the results of various translation models on Zh⇒En corpus. As seen, all advanced systems significantly outperform the baseline system (*i.e.*, RNNSEARCH), although there are still considerable differences among different variants.

Architectures of Discriminator (Rows 3-4) We evaluate two architectures for the discriminator. The CNN-based discriminator is composed of two convolution layers with 3×3 window, two max-pooling layers with 2×2 window and

one softmax layer. The feature map size is 10 and the feed-forward hidden size is 20. The RNN-based discriminator consists of two two-layer RNN encoders with 32 LSTM units and a fully-connected neural network with 32 units. We find that the RNN discriminator achieves similar performance with its CNN counterpart (37.59 vs. 37.54), while has a faster training speed (1.2K vs. 1.0K words/second). The main reason is that the CNN-based discriminator requires high computation and space cost to utilize multiple layers with convolution and pooling from a large input matrix.

Adequacy Metrics for Orientator (Rows 5-7) As aforementioned, the CDR score can be directly used as a reward to update the parameters, which is in analogy to the MRT (Shen et al. 2016) except that we use 1-best sample while they use n-best samples. For comparison, we also used the word-level BLEU score (Row 5) and character-level CHRF3 score (Popović 2015) (Row 6) as the rewards.

As seen, this strategy consistently improves translation performance, without introducing any new parameters. The extra computation cost is mainly from generating translation sentence and force decoding the human translation with the NMT model. We find that CDR not only outperforms its 1-best counterpart “O_{BLEU}” and “O_{CHRF3}”, but also surpasses “MRT_{BLEU}” using 25 samples. We attribute this to the fact that CDR can better estimate the adequacy of the translation, which is the key problem of NMT models, and go beyond the the simple low-level n-gram matching measured by BLEU and CHRF3.

Combining Them Together (Row 8) By combining advantages of both reinforcement learning and adequacy-oriented objective, our model achieves the best performance, which is 1.66 BLEU points better than the baseline “RNNSEARCH”, up to 0.98 BLEU points better than using single component and significantly improve the performance of “MRT_{BLEU}” model. One more observation can be made. “+D+O” outperforms its “+O” counterpart (*e.g.*, 8 vs. 7), which confirms our claim that the discriminator gives a

System	Model	De⇒En
<i>Existing end-to-end NMT systems</i>		
(Ranzato et al. 2016)	CNN encoder + Sequence level objective	20.73
(Bahdanau et al. 2017)	CNN encoder + Actor-critic	22.45
(Wiseman and Rush 2016)	RNNSEARCH + Beam search optimization	25.48
(Wu et al. 2017)	RNNSEARCH + Adversarial objective	26.98
<i>Our end-to-end NMT systems</i>		
<i>this work</i>	RNNSEARCH	26.51
	+ MRT _{BLEU}	27.29 [†]
	+ D_{CNN}	27.24 [†]
	+ O_{CDR}	27.31 [†]
	+ D_{RNN} + O_{CDR}	27.79 [‡]

Table 2: Comparing with previous works of applying reinforcement learning for NMT on IWSLT 2014 De⇒En translation task. “[†]” and “[‡]” indicate statistically significant difference ($p < 0.05$ and $p < 0.01$ respectively) from the RNNSEARCH model.

Model	BLEU
GNMT + RL (Wu et al. 2016)	26.30
ConvS2S (Gehring et al. 2017)	26.43
Transformer (Base) (Vaswani et al. 2017)	27.3
Transformer (Big) (Vaswani et al. 2017)	28.4
TRANSFORMER-BASE	27.30
+ O_{CDR}	27.80
+ D_{RNN} + O_{CDR}	28.01 [†]
TRANSFORMER-BIG	28.35
+ O_{CDR}	28.63
+ D_{RNN} + O_{CDR}	28.99 [†]

Table 3: BLEU scores on WMT 2014 En⇒De testset using the state-of-the-art TRANSFORMER model. “[†]” indicates statistically significant difference ($p < 0.05$) from the corresponding TRANSFORMER baseline model.

smoother and dynamically-updated score than directly using the calculated one.

Working with Coverage Model (Rows 11-12) Tu et al. (2016) propose a coverage model to indicate whether a source word is translated or not, which alleviates the inadequate translation problem of NMT models. We argue that our model is complementary to theirs, because we model the adequacy learning outside the generator by using an additional adequacy-oriented discriminator, while they model it inside the generator. Experimental results validate our hypothesis: the proposed approach further improves performance by 0.58 BLEU points over the coverage-augmented model RNNSEARCH-COVERAGE.

English-German Translation Tasks

To compare with previous work of applying reinforcement learning for NMT (Ranzato et al. 2016; Bahdanau et al. 2017; Wiseman and Rush 2016; Wu et al. 2017), we first conduct experiments on IWSLT 2014 De⇒En translation task. As listed in Table 2, we reproduce the results of adversarial training reported by Wu et al. (2017) (27.24 vs. 26.98). Furthermore, the proposed approach consistently outper-

Model	MAN	Δ	CDR	Δ
RNNSEARCH	3.31 ± 0.70	-	0.68	-
+ D	3.57 ± 0.61	7.9%	0.71	4.4%
+ O	3.69 ± 0.48	11.5%	0.75	10.3%
+ $D + O$	3.79 ± 0.47	14.5%	0.80	17.6%

Table 4: Adequacy scores on randomly selected 100 sentences on Zh⇒En task, which are measured by CDR and human evaluation (“MAN”).

forms previous works, demonstrating the effectiveness of our models.

We also evaluate our model on the recently proposed TRANSFORMER model (Vaswani et al. 2017) on WMT 2014 En⇒De corpus. As shown in Table 3, our models significantly improve performances in all cases. Combining with previous results, our model consistently improve translation performance across various language pairs and NMT architectures, demonstrating the effectiveness and universality of the proposed approach.

Analysis

To better understand our models, we conduct extensive analyses on the Zh⇒En translation task.

Adequacy Evaluation To better evaluate the adequacy, we randomly choose 100 sentences from the test set, and ask two human evaluators to judge the quality of generated translations. Five scales have been set up, i.e., $\{1, 2, 3, 4, 5\}$, where “1” means that it is irrelevant between the source sentence and the translation sentence, and “5” means that from semantic and syntactic aspect, the translation sentence and the source sentence is completely equivalent.

Table 4 lists the results of human evaluation and the proposed CDR score. First, our models consistently improve the translation adequacy under both human evaluation and the CDR score, indicating that the proposed approaches indeed alleviate the inadequate translation problem. Second, the relative improvement on CDR is consistent with that on subjective evaluation. The Pearson Correlation Coefficient




Figure 3: BLEU and CDR scores of the translations with respect to the input lengths.

between CDR and manual evaluation score is 0.64, indicating that the proposed CDR is a reasonable metric to measure translation adequacy.

Length Analysis We group sentences of similar lengths and compute both the BLEU score and CDR score for each group, as shown in Figure 3. The four length spans contain 1386, 2284, 1285, and 498 sentences, respectively. From the perspective of the BLEU score, the proposed model (*i.e.*, “+D+O”) outperforms RNNSEARCH in all length segments. In contrast, using discriminator only (*i.e.*, “+D”) outperforms RNNSEARCH in most cases, except long sentences (*i.e.*, > 45). One possible reason is that it is difficult for the discriminator to differentiate generated translations from human translations for long source sentences, thus the generator cannot learn well about these instances due to the “mistaken” rewards from the discriminator. Accordingly, using the CDR score (*i.e.*, “+O”) alleviates this problem by providing a sequence-level score, which better estimates the adequacy of the translations. The final model combines the advantages of both a smoother and dynamically-updated objective from the discriminator (“+D”), and a more accurate objective specifically designed for the translation task from the orientator (“+O”).

The CDR scores for all models degrade when the length of source sentence increases. This is mainly due to that inadequate translation problem is more serious on longer sentences for NMT models (Tu et al. 2016). The adversarial model (*i.e.*, “+D”) improves CDR scores while the improvement degrades faster with the increase of sentence length. However, our proposed approach consistently improves CDR performance in all length segments.

Input	美国线上时代华纳董事长凯斯今晚宣布他将在五月辞职,他说,此举「对公司而言是最好的」。
REF	Steve case chairman of AOL time warner announced tonight that he will resign in May, he said that his move is the best interest of the company.
BASE	The chairman of the board of time warner of the united states announced tonight that he will resign in May. (<i>CDR: 0.56; BLEU: 36.18</i>)
OURS	Chairman of AOL time warner announced tonight that he will resign in May, he said this is the best thing for the company. (<i>CDR: 0.88; BLEU: 58.82</i>)

Figure 4: Example translations on Zh \Rightarrow En test set.

Effect of the Discriminator Koehn and Knowles (2017) point out that the attention model does not always correspond to word alignment and may considerably diverge. Accordingly, the attention matrix-based CDR score may not always correctly reflect the adequacy of generation sentences. However, our discriminator is able to give a smoother and dynamically-updated objective, and thus could provide more accurate adequacy scores of generation sentences. From the above quantitative and qualitative results, the discriminator indeed leads to better performance (*i.e.*, “+D+O” vs. “+O”).

Case Study To better understand the advantage of our proposed model, we show a translation case in Figure 4. Specially, we provide a Zh \Rightarrow En example with two translation results from the RNNSearch and Adequacy-NMT models respectively, as well as the corresponding CDR and BLEU scores. We emphasize on their different parts with bold fonts which lead to different translation quality. As seen, the latter part of the source sentence is not translated by the RNNSearch model while our proposed model correct this mistake. Accordingly, our model improves both CDR and BLEU scores.

Conclusion

In this work, we propose a novel learning approach for RL-based NMT models, which integrates into the policy gradient with an adequacy-oriented reward designed specifically for translation. The proposed approach combines the advantages of both sequence-level training of reinforcement learning, as well as a more accurately estimated reward by considering the translation adequacy in terms of coverage difference ratio (CDR). Experimental results on different language pairs show that our proposed approach not only significantly outperforms standard NMT models, but also further improves performance over those using the policy gradient and the adequacy-oriented reward individually. In addition, the proposed approach is also complementary to the coverage models (Tu et al. 2016), because the two models aim to alleviate the inadequate translation problem from two different perspectives (*i.e.*, sequence-level vs. word-level).

Future directions include validating our approach on other architectures such as CNN-based NMT models (Gehring et al. 2017) and improved TRANSFORMER models (Shaw, Uszkoreit, and Vaswani 2018; Shen et al. 2018), as well as combining with other advanced techniques in reinforcement

learning and adversarial learning (Bahdanau et al. 2017; Yu et al. 2017; Yang et al. 2018).

References

- Bahdanau, D.; Brakel, P.; Xu, K.; Goyal, A.; Lowe, R.; Pineau, J.; Courville, A.; and Bengio, Y. 2017. An actor-critic algorithm for sequence prediction. In *ICLR*.
- Bahdanau, D.; Cho, K.; and Bengio, Y. 2015. Neural machine translation by jointly learning to align and translate. In *ICLR*.
- Bengio, S.; Vinyals, O.; Jaitly, N.; and Shazeer, N. 2015. Scheduled sampling for sequence prediction with recurrent neural networks. In *NIPS*.
- Cheng, Y.; Tu, Z.; Meng, F.; Zhai, J.; and Liu, Y. 2018. Towards robust neural machine translation. In *ACL*.
- Collins, M.; Koehn, P.; and Kučerová, I. 2005. Clause restructuring for statistical machine translation. In *ACL*.
- Gehring, J.; Auli, M.; Grangier, D.; Yarats, D.; and Dauphin, Y. N. 2017. Convolutional sequence to sequence learning. In *ICML*.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. In *NIPS*.
- He, D.; Lu, H.; Xia, Y.; Qin, T.; Wang, L.; and Liu, T. 2017. Decoding with value networks for neural machine translation. In *NIPS*.
- Kalchbrenner, N., and Blunsom, P. 2013. Recurrent continuous translation models. In *EMNLP*.
- Koehn, P., and Knowles, R. 2017. Six challenges for neural machine translation. In *Proceedings of the First Workshop on Neural Machine Translation*, 28–39.
- Luong, M.-T.; Pham, H.; and Manning, C. D. 2015. Effective approaches to attention-based neural machine translation. In *EMNLP*.
- Meng, F.; Tu, Z.; Cheng, Y.; Wu, H.; Zhai, J.; Yang, Y.; and Wang, D. 2018. Neural machine translation with key-value memory-augmented attention. In *IJCAI*.
- Mi, H.; Sankaran, B.; Wang, Z.; and Ittycheriah, A. 2016. Coverage embedding models for neural machine translation. In *EMNLP*.
- Papineni, K.; Roukos, S.; Ward, T.; and Zhu, W.-J. 2002. BLEU: a method for automatic evaluation of machine translation. In *ACL*.
- Popović, M. 2015. chrF: character n-gram f-score for automatic mt evaluation. In *Proceedings of the Tenth Workshop on Statistical Machine Translation*, 392–395.
- Ranzato, M.; Chopra, S.; Auli, M.; and Zaremba, W. 2016. Sequence level training with recurrent neural networks. In *ICLR*.
- Shaw, P.; Uszkoreit, J.; and Vaswani, A. 2018. Self-Attention with Relative Position Representations. In *NAACL*.
- Shen, S.; Cheng, Y.; He, Z.; He, W.; Wu, H.; Sun, M.; and Liu, Y. 2016. Minimum risk training for neural machine translation. In *ACL*.
- Shen, T.; Zhou, T.; Long, G.; Jiang, J.; Pan, S.; and Zhang, C. 2018. DiSAN: directional self-attention network for RNN/CNN-free language understanding. In *AAAI*.
- Sutskever, I.; Vinyals, O.; and Le, Q. V. 2014. Sequence to sequence learning with neural networks. In *NIPS*.
- Tu, Z.; Lu, Z.; Liu, Y.; Liu, X.; and Li, H. 2016. Modeling coverage for neural machine translation. In *ACL*.
- Tu, Z.; Liu, Y.; Shang, L.; Liu, X.; and Li, H. 2017. Neural machine translation with reconstruction. In *AAAI*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *NIPS*.
- Williams, R. J. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning* 8(3-4):229–256.
- Wiseman, S., and Rush, A. M. 2016. Sequence-to-sequence learning as beam-search optimization. In *EMNLP*.
- Wu, Y.; Schuster, M.; Che, Z.; and Le, Q. V. e. a. 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv*.
- Wu, L.; Xia, Y.; Zhao, L.; Tian, F.; Qin, T.; Lai, J.; and Liu, T.-Y. 2017. Adversarial neural machine translation. *arXiv*.
- Yang, Z.; Chen, W.; Wang, F.; and Xu, B. 2018. Improving neural machine translation with conditional sequence generative adversarial nets. In *NAACL*.
- Yu, L.; Zhang, W.; Wang, J.; and Yu, Y. 2017. SeqGAN: sequence generative adversarial nets with policy gradient. In *AAAI*.
- Zhang, J.; Ding, Y.; Shen, S.; Cheng, Y.; Sun, M.; Luan, H.; and Liu, Y. 2017. THUMT: An open source toolkit for neural machine translation. *arXiv preprint arXiv:1706.06415*.
- Zheng, Z.; Zhou, H.; Huang, S.; Mou, L.; Xinyu, D.; Chen, J.; and Tu, Z. 2018. Modeling past and future for neural machine translation. *TACL*.