

Towards Reliable Neural Machine Translation with Consistency-Aware Meta-Learning

Rongxiang Weng^{1,2}, Qiang Wang^{3,4}, Wensen Cheng², Changfeng Zhu², Min Zhang¹

¹Soochow University, Suzhou, China

²miHoYo AI, Shanghai, China

³Zhejiang University, Hangzhou, China

⁴RoyalFlush AI Research Institute, Hangzhou, China

wengrongxiang@gmail.com, wangqiangneu@gmail.com, vinson7973@gmail.com, bajiuchangfeng@126.com, minzhang@suda.edu.cn

Abstract

Neural machine translation (NMT) has achieved remarkable success in producing high-quality translations. However, current NMT systems suffer from a lack of reliability, as their outputs that are often affected by lexical or syntactic changes in inputs, resulting in large variations in quality. This limitation hinders the practicality and trustworthiness of NMT. A contributing factor to this problem is that NMT models trained with the *one-to-one* paradigm struggle to handle the *source diversity* phenomenon, where inputs with the same meaning can be expressed differently. In this work, we treat this problem as a bilevel optimization problem and present a consistency-aware meta-learning (CAML) framework derived from the model-agnostic meta-learning (MAML) algorithm to address it. Specifically, the NMT model with CAML (named CoNMT) first learns a consistent meta representation of semantically equivalent sentences in the outer loop. Subsequently, a mapping from the meta representation to the output sentence is learned in the inner loop, allowing the NMT model to translate semantically equivalent sentences to the same target sentence. We conduct experiments on the NIST Chinese to English task, three WMT translation tasks, and the TED M2O task. The results demonstrate that CoNMT effectively improves overall translation quality and reliably handles diverse inputs.

Introduction

Neural machine translation (NMT) based on the *encoder-decoder structure* (Bahdanau, Cho, and Bengio 2014; Sutskever, Vinyals, and Le 2014) and *attention mechanism* (Luong, Pham, and Manning 2015), has achieved prominent success in the vast majority of translation tasks (Cho et al. 2014; Vaswani et al. 2017) in recent years. Most NMT models are trained by parallel data in an end-to-end manner. In a nutshell, the encoder encodes an input sentence into sequential hidden states as contextual representation. Then, the decoder fetches the contextual representation through cross attention and generates the corresponding target sentence.

However, many previous studies have shown that NMT models are highly susceptible to lexical or syntactic changes in inputs (Ott et al. 2018; He, Meister, and Su 2020). This means that altering the structure of the input sentence or even

Source: 在没有新来的驾驶员的情况下，你让我校准了操纵杆。

Trans: You make me calibrate joystick without a new driver.

Source: 在没有新来的驾驶员的情况下，你让我校准了控制杆。

Trans: Under the case of no new *pilot*, you ask me to *control the control stick*.

Source: 你让我校准了操纵杆，在新来的那个驾驶员还没来的情况下。

Trans: You let me calibration the joystick *when you haven't come yet*.

Figure 1: Translations (Trans) from semantically equivalent inputs (Source). Compared to the first one, Some *semantically irrelevant* modifications in the following inputs bring apparent changes in the outputs.

just a few words, without changing the semantics, can result in a significant change in the output, frequently leading to inaccurate translations (See Figure 1). One reason for this is that the *one-to-one* training paradigm conflicts with the *many-to-many* characteristic of natural language. We define the aforementioned phenomenon as the *source diversity* problem and focus on the *many-to-one* situation in this work. In reality, the *many-to-one* mapping is commonplace in NMT. For example, individuals may express the same concept in different ways, and even two descriptions of the same concept from a single individual may vary significantly. When NMT correctly handles the source diversity, it will become more reliable in diversified application scenarios.

Some recent works have noticed this problem and sought to address it. Nguyen et al. (2020) proposed a data augmentation method, which back-translates the original sentence by several machine translation systems, and then directly adds them into the model training process. However, it does not constrain the NMT model, leaving them to deal only with synthetic data. Wei et al. (2020) proposed a variational method to model uncertainty, so as to map semantically equivalent sen-

tences into a sentence representation. However, this method is tendentious to get a trivial solution, where all sentences are embedded in a meaningless vector. Moreover, studies on the robustness of NMT can also mitigate this problem. Liu et al. (2019) incorporated noisy input with similar pronunciations into NMT training to address *homophone noise*. Cheng, Jiang, and Macherey (2019); Cheng et al. (2020) proposed to use adversarial examples to address *synonym problem*. However, they only focus on the noisy input, part of the source diversity phenomenon.

In this work, we divide this source diversity problem into two subproblems: The first one is *how to learn a consistent context representation from semantically equivalent sentences*; The second one is *how to map representation to the ground-truth outputs*. Following this intuition, we present a consistency-aware meta-learning (CAML) framework based on the model-agnostic meta-learning (MAML) algorithm (Finn, Abbeel, and Levine 2017) to address it. Unlike the traditional MAML, our CAML contains two task-specific training objectives to constrain the meta representation. Concretely, we design 1). a sentence-level objective to force semantically equivalent sentences to generate each other; 2). a word-level objective to constrain the output distributions from these sentences to be similar. CAML can promote NMT to learn a consistent meta representation from semantically equivalent samples in the outer loop. After this, NMT will focus on learning a mapping from the meta representation to the output sentence in the inner loop.

To illustrate the effectiveness of the proposed CAML, we conduct experiments on widely used machine translation tasks with different settings. Specifically, we first evaluate our approach on the WMT14 English to German (En→De), WMT17 Chinese to English (Zh→En), and WMT16 Romanian to English (Ro→En) tasks, and gets more than 1 BLEU score gain in all of them. Further, we apply our approach to three varieties: non-autoregressive generation, multilingual translation, and larger model structure. The empirical results are indicative of a good versatility of our method. Finally, we make a simulation experiment, which demonstrating that our approach effectively alleviates the influence of the different forms of inputs.

The Proposed Approach

In this section, we will briefly introduce NMT based on the Transformer (Vaswani et al. 2017) structure and the application of MAML (Finn, Abbeel, and Levine 2017) to NMT. Then, we will describe the proposed CAML based on the Transformer backbone in detail.

Background and Notations

Neural machine translation. Given a *source-target* sentence pair $\{x, y\}$ from the parallel data-set \mathcal{B} , where $|\mathcal{B}|$ is the number of sentence pairs. The paradigm of NMT is:

$$P(y|x) \propto \text{DEC}(\text{ENC}(x; \theta_e); \theta_d), \quad (1)$$

where the $\text{ENC}(\cdot)$ and $\text{DEC}(\cdot)$ are encoder and decoder networks respectively, and θ_e and θ_d are corresponding parameters. In Transformer, they are composed of multiple self-

attention layers. Then the NMT model is optimized by minimizing the negative log-likelihood, denoted by:

$$\mathcal{L}_N(y, f_\theta(x)) = -\mathbb{E}_{\{x, y\} \sim \mathcal{B}} [\log P(y|x; \theta)], \quad (2)$$

where $\theta = \{\theta_e, \theta_d\}$, $f_\theta(x)$ denotes the model prediction of x with model parameters θ .

MAML for NMT. Model-agnostic meta-learning (MAML) is an widely used meta-learning (or *learning to learn*) algorithm, which divides the training process into two stages (Finn, Abbeel, and Levine 2017), named *meta-train* and *meta-test*. In *meta-train* step, a general model is learned from different data (called supported set), while a task-specific data (called target set) is used in the *meta-test* step.

In our setup, we take the original parallel data $\{x, y\}$ as the target set, and construct the support set by introducing a set \mathbf{x}^s which is semantically equivalent to x . $\mathbf{x}^s = \{x_1^s, \dots, x_i^s, \dots, x_I^s\}$, where I is the number of sentences. The \mathbf{x}^s is generated from a sampling function $\varphi(x, y)$, which will be described in the Section . Following Gu et al. (2018); Li et al. (2018a), we can readily employ MAML on NMT to utilize \mathbf{x}^s . In this way, we first update the current model parameters from θ to θ' by the gradients in the meta-train step:

$$\theta' \leftarrow \theta - \alpha \nabla_\theta \{\mathcal{L}_{\text{Train}}(y_i, f_\theta(x_i^s))\}, \quad (3)$$

where α is the learning rate and $\mathcal{L}_{\text{Train}}$ is the loss function for the meta-train step. We can calculate $\mathcal{L}_{\text{Train}}$ through the support set as:

$$\mathcal{L}_{\text{Train}}(y_i, f_\theta(x_i^s)) = \frac{1}{I} \sum_i \{\mathcal{L}_N(y_i, f_\theta(x_i^s))\}. \quad (4)$$

Given θ' , we further leverage the target set to update the model parameters through the meta-test step:

$$\mathcal{L}_{\text{Test}}(y, f_\theta(x)) = \mathcal{L}_N(y, f_{\theta'}(x)). \quad (5)$$

Consistency-aware Meta-learning for NMT

The ideal case for increasing the reliability of NMT is that semantically equivalent sentences have the same representation in the semantic space, which allows the model to generate stable outputs from various inputs. Following this motivation, the proposed CAML concentrates on promoting NMT to learn a consistent meta representation from semantically equivalent samples and to translate from the meta representation to the ground truth. In this section, we will first introduce two simple training objectives in the CAML, which constrain the contextual representation in different aspects. Then, we will expound on how to incorporate the CAML into the training process to build a reliable NMT model.

Modeling consistency through multi-level constraints.

In Transformer, the output is decided by the contextual representation encoded from the encoder. Namely, when the representations from sentences in \mathbf{x}^s are equivalent:

$$\|\text{ENC}(x_i^s; \theta_e) - \text{ENC}(x_j^s; \theta_e)\|_2^2 = 0, \{x_i^s, x_j^s\} \subset \mathbf{x}^s, \quad (6)$$

where the NMT model can generate the same output by the contextual representation. However, it is hard to optimize this

Algorithm 1: The training process of NMT with CAML

Input: Parallel data set \mathcal{B} ; Sampling function: $\varphi(\cdot)$; Epoch E
Output: The NMT model $f_\theta(\cdot)$

```

1: for  $e = 0$  to  $E$  do
2:   for  $\{x, y\}$  in  $\mathcal{B}$  do
3:     Optimizing  $f_{\hat{\theta}}(\cdot)$  by the Equation 2
4:   end for
5: end for
6: for  $e = 0$  to  $E$  do
7:   for  $\{x, y\}$  in  $\mathcal{B}$  do
8:     Sampling  $\mathbf{x}^s$  by the  $\varphi(x, y)$ 
9:     for  $_$  in  $\mathbf{T}$  do
10:      Optimizing  $f_{\theta'}(\cdot)$  by the Equation 9
11:    end for
12:    Optimizing  $f_\theta(\cdot)$  by the Equation 10
13:  end for
14: end for
15: return  $f_\theta(\cdot)$ 

```

objective due to the variable sentence structure. From another perspective, we consider two training objectives of different granularity to optimize the NMT model.

The first is the consistency of the sentence representation, which we consider to be optimized by a partial reconstruction manner. We randomly sample a sentence pair $\{x_i^s, x_j^s\}$ from \mathbf{x}^s and mask the same part of them, which are denoted as $\{x_i^m, x_j^m\}$. Then, the NMT model is asked to generate one when given another. Specifically, the sentence level training loss is formalized as:

$$\mathcal{L}_S(x_j^s, x_i^s) = \mathcal{L}_N(x_j^s, f_\theta(x_i^s)) + \mathcal{L}_N(x_i^s, f_\theta(x_j^s)) = \quad (7)$$

$$- \mathbb{E}_{\{x_i^s, x_j^s\} \sim \varphi(\cdot)} \mathbb{E}_{\{x, y\} \sim \mathcal{B}} [\log P(x_j^m | x_i^s; \theta) P(x_i^m | x_j^s; \theta)].$$

Here, \mathcal{L}_S encourages the contextual representation to recover sentences that have similar semantics as itself. This is almost equal to enabling NMT to generate the same sentence representations between semantically equivalent sentences.

Besides modeling the sentence-level semantics embedded in the sequential representations, we need to treat the difference in syntactic structure, determining how the decoder gets the correct contextual information. In contrast, instead of exploiting the syntactic information, which needs human-annotated data, we adopt a word-level training objective as an alternative to making the output distribution generated by one sentence governed by the others.

Specifically, we constrain the output distributions from semantically equivalent sentences to be similar at each decoding step. Formally, the word-level loss function is summarized as:

$$\mathcal{L}_W(f_{\tilde{\theta}}(x), f_\theta(x_i^s)) = \quad (8)$$

$$- \mathbb{E}_{x_i^s \sim \varphi(\cdot)} \mathbb{E}_{\{x, y\} \sim \mathcal{B}} [P(y|x; \tilde{\theta}) \log P(y|x_i^s; \theta)],$$

where $\tilde{\theta}$ is the copy of the NMT parameters, which will not be updated by the gradient. For different forms of input, the NMT decoder trained by this objective can get the similar contextual information at each step, avoiding the the influence of the structural differences.

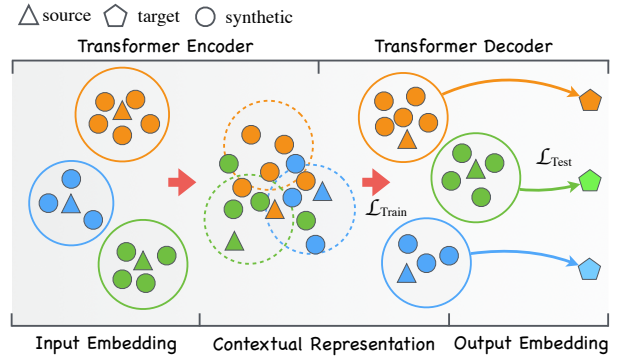


Figure 2: An intuitive illustration of the impact of the proposed CAML on the contextual representation of NMT.

Meta-learning with \mathcal{L}_S and \mathcal{L}_W . Following the MAML framework (Finn, Abbeel, and Levine 2017), we sample T pairs from \mathbf{x}^s as the *support set*, and then use the $\{x_i^s, x, y\}$ as the *target set*. The training function of the meta-train step can be summarized as follows:

$$\mathcal{L}_{\text{Train}}(x_{\setminus i}^s, f_\theta(x_i^s)) = \quad (9)$$

$$\beta \sum \left\{ \mathcal{L}_S(x_{\setminus i}^s, f_\theta(x_i^s)) + \mathcal{L}_W(f_\theta(x_{\setminus i}^s), f_\theta(x_i^s)) \right\},$$

where $x_{\setminus i}^s$ is a sentence from \mathbf{x}^s which not equal to x_i^s , β is the learning rate for the meta-train step. Compared to the Eq. 4, we use the aforementioned self-supervised training objectives instead of the NMT training loss (*i.e.*, \mathcal{L}_T) in this step. Our goal is to facilitate the model to pay attention to learning the contextual representation in the outer loop rather than directly learning the translation from noisy input data.

The next step of the CAML is to learn to translate based on contextual representation. An intuitive approach is to adopt Eq. 9 to train the model. However, it may cause the *forgetting* problem since the objective functions of the two steps are different. To avoid this problem, we integrate \mathcal{L}_S and \mathcal{L}_W into the meta-test step as a regularization term to maintain the representation consistently. Specifically, the meta-test step in the CAML is reorganized as:

$$\mathcal{L}_{\text{Test}}(y, f_\theta(x)) = \gamma \mathcal{L}_N(y, f_\theta(x)) \quad (10)$$

$$+ \epsilon \underbrace{\left\{ \mathcal{L}_S(x, f_{\theta'}(x_i^s)) + \mathcal{L}_W(f_\theta(x), f_{\theta'}(x_i^s)) \right\}}_{\text{maintaining consistency}}.$$

The coefficients γ and ϵ are used to balance the preference among these losses, which are empirically set to 1 and 0.5, respectively. Here, the updated parameters θ' are computed by T gradient descent updates on the support task, which equals $2I$. The one gradient update in the meta-train step is:

$$\theta' \leftarrow \theta - \gamma \nabla_\theta \left\{ \mathcal{L}_S(x_{\setminus i}^s, f_\theta(x_i^s)) + \mathcal{L}_W(f_\theta(x_{\setminus i}^s), f_\theta(x_i^s)) \right\},$$

where γ is the learning rate of the meta-train step. The overall training algorithm of our approach is shown in Algorithm 1. An intuitive illustration is shown in Figure 2.

Semantically Equivalent Sentences

A key factor of our approach is how to build the \mathbf{x}^s by the $\varphi(\cdot)$. It is difficult to obtain perfect semantically equivalent sentences by data augmentation methods. Here, we use two steps to construct semantically similar sentences as approximations. Coincidentally, previous studies show that noisy samples enhance the performance of meta-learning (Wang, Hu, and Hu 2020; Yao et al. 2021). It is worthy to note that our work focuses on the training strategy of NMT, so we only use two common data augmentation approaches. Any advanced approaches can be easily applied here (Wang et al. 2018; Hoang et al. 2018; Nguyen et al. 2020).

Step 1: replacing words with alignment relationships.

Given the sentence x , we sample and replace 20% words with similar words. More specifically, for a selected word in x , we choose 100 words from the vocabulary \mathcal{V} ordered by the frequency of a word alignment model, which is trained on the same parallel data.¹ Then, we form these words as a subset and randomly choose one to replace the original word.

Step 2: noisy round-trip translation.

Then, we use noisy round-trip translation to process the data from Step 1 further. Given sentences from Step 1, we add random vectors sampled from the standard normal distribution to them and translate them to target language sentences $\mathbf{y}' = \{y'_1, \cdot, y'_j, \cdot, y'_J\}$, where J is the number of sentences. We select several sentences from the \mathbf{y}' according to the weight $\text{EDITDIST}(y'_j, y) * \text{OVERLAP}(y'_j, y)$. The $\text{EDITDIST}(\cdot)$ is edit distance and $\text{OVERLAP}(\cdot)$ is word overlap ratio. Then, these sentences are back-translated to the source language to compose the \mathbf{x}^s .

Experiments

Implementation Detail

Data-set. We first conduct experiments on the NIST Chinese to English (Zh→En) task and three widely-used WMT translation tasks: WMT14 English to German (En→De), WMT16 Romanian to English (Ro→En), and WMT18 Chinese to English (Zh→En). On the NIST Zh→En task, we use NIST 2006 (MT06) as the dev set and NIST 2002 (MT02), 2003 (MT03), 2004 (MT04), 2005 (MT05), 2008 (MT08) as the test sets. On the En→De task, we use `newstest2013` as the dev set and `newstest2014` as the test set. On the Ro→En task, we use `newstest2015` as the dev set and `newstest2016` as the test set. On the Zh→Ee task, we use `newsdev2017` as the dev set, and `newstest2017` as the test set. Moreover, we evaluate our method on the TED multilingual machine translation task. The overall data-set is divided into two training sets according to language diversity, which is named *related* and *diverse*. The detail is shown in Wang, Tsvetkov, and Neubig (2020). We only conduct the many-to-one (M2O), which translates eight languages to English.

Setting. We adopt the Transformer-*base* for the NIST Zh→En and WMT Ro→En, Transformer-*big* for the WMT Zh→En, and both *base* and *big* for the WMT En→De. We

¹https://github.com/clab/fast_align

apply the byte pair encoding (BPE) (Sennrich, Haddow, and Birch 2016b) to all language pairs and limit vocabulary size to 32K. All out-of-vocabulary words were mapped to the *UNK*. In order to avoid the problem of vocabulary mismatch, we adopt an additional four-layer parallel decoder (Ghazvininejad et al. 2019) for the \mathcal{L}_S . We set label smoothing as 0.1 and dropout rate as 0.1. The Adam is adopted as the optimizer, and the β_1/β_2 is set as 0.9/0.98 for the *base* setting and 0.9/0.998 for the *big* setting. The initial learning rate is 0.001. We adopt the warm-up strategy with 4000 steps. The learning rate γ of *meta-train* remains one-tenth of NMT. Other settings are followed as Vaswani et al. (2017).

We implement our approach on *fairseq*². We use 4 V100 GPUs and accumulate the gradient four iterations on the WMT En→De and Zh→En. Other tasks run on 2 V100 GPUs. We use beam search as the decoding algorithm and set the beam size as 5 and the length penalty as 0.6. For a fair comparison, we calculate the *case-sensitive tokenized* BLEU with the *multi-bleu.perl* script for the NIST Zh→En, and use the *sacreBLEU*^{3,4} to calculate *case-sensitive* BLEU (Papineni et al. 2002) for all WMT and TED tasks.

Results on NIST Zh→En Task

In order to compare with the related work, we first carried out experiments on the NIST Zh→En task with the same setting (Wei et al. 2020; Cheng et al. 2020). Following Cheng et al. (2020), we sample 1.25M English sentences from the Xinhua portion of the gigaword corpus as the extra corpus to enhance the model.⁵

The overall results are shown in Table 1. We divide the experiment into three parts. Firstly, we compare our model to the Transformer baselines. We implement the vanilla Transformer-base trained by the parallel data (Transformer, Row 10) and both the parallel and synthetic data (Transformer_{sync}, Row 11), respectively. The synthetic data brings a 1.36 BLEU score gain. Transformer with the proposed CAML (CONMT, Row 12) attains 47.64 BLEU. When incorporating the synthetic data, our approach improves 1.24 to 48.89 BLEU (CONMT_{sync}, Row 13). Summarily, our approach gets 3.23/3.12 absolute improvements compared to the corresponding baselines on the NIST Zh→En task.

Then, we compare our approach to several related work including the robust NMT (Ebrahimi, Lowd, and Dou 2018; Cheng, Jiang, and Macherey 2019; Cheng et al. 2020), data-augmentation (Wei et al. 2020), and enhancing the contextual representation (Hao et al. 2019). In the same way, we also report the results from the model trained with or without synthetic data. Our approach achieves the best performance on both settings. Furthermore, our method outperforms the AdvAug by 0.57/0.50 BLEU and achieves state-of-the-art in four test sets.

²<https://github.com/pytorch/fairseq>

³<https://github.com/mjpost/sacreBLEU>

⁴BLEU+case.mixed+lang.\${Src}-\${Trg}+num-refs.1+smooth.exp+test.\${Task}+tok.13a+version.1.5.1

⁵<https://catalog.ldc.upenn.edu/LDC2003T05>

#	Model	MT06	MT02	MT03	MT04	MT05	MT08	Avg
1	Vaswani et al. (2017)	44.57	45.49	44.55	46.20	44.96	35.11	43.26
2	Ebrahimi, Lowd, and Dou (2018)	45.28	45.95	44.68	45.99	45.32	35.84	43.56
3	Cheng, Jiang, and Macherey (2019)	46.95	47.06	46.48	47.39	46.58	37.38	44.98
4	Hao et al. (2019)	44.00	N/A	43.98	45.60	44.28	N/A	N/A
5	Cheng et al. (2020)	49.26	49.03	47.96	48.86	49.88	39.63	47.07
6	Sennrich, Haddow, and Birch (2016a)*	46.20	47.78	46.93	47.80	46.81	36.79	45.22
7	Cheng, Jiang, and Macherey (2019)*	47.74	48.13	47.83	49.13	49.04	38.61	46.55
8	Wei et al. (2020)*	48.88	N/A	N/A	49.15	49.21	40.94	N/A
9	Cheng et al. (2020)*	49.98	50.34	49.81	50.61	50.72	40.45	48.39
10	Transformer	45.24	46.89	46.32	46.83	45.72	36.31	44.41
11	Transformer* _{sync}	46.57	47.13	47.69	47.51	46.88	39.67	45.77
12	CoNMT	49.41	49.83	48.83	49.51	50.03	39.99	47.64
13	CoNMT* _{sync}	49.74	50.79	50.53	50.96	51.27	40.88	48.89

Table 1: The effectiveness of our approach (CoNMT) on the NIST Zh→En task. “*” indicates the model uses extra monolingual data-set. “Avg” is the average score of all test sets.

#	Model	Trans. Base	Trans. Big
1	Vaswani et al. (2017)	27.30	28.40
2	Li et al. (2018b)	28.51	29.28
3	Cheng et al. (2018)*	28.09	N/A
4	Gao et al. (2019)*	N/A	29.70
5	Jiao et al. (2020)*	28.30	29.20
6	Transformer	27.53	28.73
7	Transformer* _{sync}	28.82	29.71
8	CoNMT	28.64 [‡]	29.58 [‡]
9	CoNMT* _{sync}	30.22[‡]	30.87[‡]

Table 2: The effectiveness of our approach (CoNMT) on the En→De task. “†/‡” indicate the model is significantly better than the baseline ($p < 0.05/0.01$) (Neubig et al. 2019). “*” indicates using extra monolingual data. “Trans. Base” and “Trans. Big” are the *base* and *big* setting, respectively.

Results on WMT Tasks

We further evaluate our approach on three widely-used WMT benchmarks, including WMT14 En→De, WMT17 Zh→En, and WMT16 En→Ro. Meanwhile, we report several previous work as a comparison.⁶ The results on WMT14 En→De are summarized in the Table 2. Our model gets 28.64/29.58 BLEU in the *base/big* setting (Row 8). Compared to the baseline (Row 7), our model achieves 1.11/0.85 gains. We use 8M German monolingual sentences sampled from Caswell, Chelba, and Grangier (2019) to enhance the model. When combined with the synthetic data, our model gets 30.22/30.87 BLEU (Row 9), which gets 1.4/1.06 improvements and outperforms all previous work.

The results on the WMT17 Zh→En and WMT16 Ro→En are shown in Table 3. We use 10M monolingual sentences sampled from WMT News Craw and 2M monolingual sen-

⁶We consider the work which measure the results by the *score-bleu* toolkit.

#	Model	Zh→En	Ro→En
1	Vaswani et al. (2017)	25.21	31.90
2	Wei et al. (2020)	26.48	N/A
3	Yang et al. (2019)	25.15	N/A
4	Zhou et al. (2020)	25.04	N/A
5	Gu et al. (2018)	N/A	31.76
6	Transformer	25.37	31.93
7	Transformer* _{sync}	26.63	36.94
8	CoNMT	26.38	33.72
9	CoNMT* _{sync}	27.49	38.44

Table 3: The effectiveness of our approach (CoNMT) on the WMT Zh→En and WMT16 Ro→En. “*” indicates that the model uses extra monolingual data. We use the Transformer-big on the Zh→En and Transformer-base on the Ro→En.

tences from Sennrich, Haddow, and Birch (2016a) for the Zh→En and Ro→En, respectively. Compared to baselines, our work gets 1.01/0.86 gains on the Zh→En, and 1.79/1.50 gains on the Ro→En. The extensive experiments on the WMT benchmark demonstrate more powerfully the effect of the proposed CAML algorithm.

Effectiveness on Different Varieties

The proposed CAML framework only changes the training schema of NMT, which means it can be employed on any Transformer-based NMT tasks. In this section, we conduct condensed experiments on three popular NMT directions: non-autoregressive translation, multi-lingual translation, and larger model size.

Non-autoregressive NMT. Our approach is tangential to the decoding mode. Thus, whether our approach works on NAT (Gu et al. 2017; Ghazvininejad et al. 2019) is valuable to investigate. Here, we employ the CAML based on the

Model	Speedup	En→De
Vaswani et al. (2017)	1.0x	27.30
Gu et al. (2017)	2.36x	19.17
Ghazvininejad et al. (2019)	1.7x	27.03
Transformer	1.0x	27.53
Mask-Predict	1.63x	27.14
CoNMT	1.63x	28.21

Table 4: The results of the proposed CAML (CoNMT) on the NAT structure.

Model	Related	Diverse
Uni. ($\tau = \infty$)	22.63	24.81
Temp. ($\tau = 5$)	24.00	26.01
Prop. ($\tau = 1$)	24.88	26.68
Wang, Tsvetkov, and Neubig (2020)	25.26	27.00
Wu et al. (2021)	26.28	27.82
CoNMT	26.77	28.09

Table 5: The results of our method (CoNMT) on the M2O task. *Uni.*: sampling data with equal frequency. *Temp.*: sampling data according to size exponentiated by τ . *Prop.*: sampling data in portions equivalent to the size of each language.

Mask-Predict (Ghazvininejad et al. 2019) framework⁷ and evaluate on the En→De task. The results are shown in Table 4. Our work largely outperforms the Mask-Predict baseline (+1.07 BLEU) without dropping the decoding speed.

Multilingual NMT. We believe that sentences from different languages with the same semantics can be regarded as semantically equivalent sentences. To verify that, we employ our approach on the many-to-one (M2O) translation task following the setting from Wang, Tsvetkov, and Neubig (2020). We sample I sentences and use a one-to-many NMT model to translate to other languages. Then these synthetic sentences are used as training data for the meta-train step.

The results are shown in Table 5. Compared to Wang, Tsvetkov, and Neubig (2020), our approach obtains 1.51/1.09 improvements in the related/diverse language pairs. Furthermore, our approach gets better results than Wu et al. (2021), especially in related language pairs. Empirically, our approach works better for similar language pairs than diverse ones, which indicate that it is difficult for NMT to learn consistent representations for diverse languages.

Larger model size. A spontaneous confusion is whether a larger model can eliminate the source diversity problem. To verify the availability in this situation, we employ the CAML on the pre-trained mBART (Liu et al. 2020) and Deep Transformer (Wang et al. 2019) on the En→De task. The results are shown in Table 6. Compared to the fine-tuning, our approach gets a 1.41 improvement. Then, when combined with Deep Transformer, our approach gets a 1.44 gain, outperforming

⁷The detail of this model is shown in Ghazvininejad et al. (2019).

Model	#Param	En→De
Vaswani et al. (2017)	213M	28.40
Wei et al. (2022)	265M	29.80
Wang et al. (2019)	137M	29.30
Zhu et al. (2020)	500M	30.75
Raffel et al. (2020)	700M	30.90
Deep Transformer	159M	29.82
mBART Fine-tuning	610M	28.82
CoNMT _{sync}	207M	30.87
mBART w/ CAML	610M	30.23
Deep Transformer w/ CAML	159M	31.26

Table 6: The results of the proposed method with larger model size on the En→De task.




Figure 3: The BLEU curve under the different ratio of diverse input. The dotted line is the BLEU from the original inputs.

all previous work. The evaluation results demonstrate that a larger model can not escape the source diversity problem, and our CAML still works in this situation.

Analysis

The influence of the diverse input. We sample 300 samples from MT02-05, each one has one source sentence and four references, to appraise the ability of our approach to handle the source diversity problem. We ask two professional translators to translate source sentences according to the first reference and the remaining three, respectively. Then, we use the first reference and the corresponding translation as the standard test set, and replace input sentences from another translator in different proportions. The results are presented in Figure 3. As the replacement proportion increases, the performance of all models decreases. However, our model is more reliable than the baseline, which performs similarly to the original input even with a large replacement proportion.

Distribution across translation quality. Then, we sample one thousand translations on the En→De task and calculate the sentence-BLEU of each one.⁸ We divide these sentences into several groups according to the BLEU score and count

⁸We use the script from THUMT: <https://github.com/THUNLP-MT/THUMT/blob/master/thumt/utis/bleu.py>




Figure 4: The number of sentences divided according to the sentence BLEU.

Model	En→De	Δ
Transformer (Vaswani et al. 2017)	27.53	–
+ Data Augmentation	27.87	+0.34
+ MAML (Eq.5-4)	27.41	-0.12
+ MTL with \mathcal{L}_W and \mathcal{L}_S	28.01	+0.48
+ CAML (CoNMT)	28.64	+1.11
+ \mathcal{L}_T in Eq. 9	28.37	+0.84
– \mathcal{L}_S in Eq. 9 and Eq. 10	28.27	+0.74
– \mathcal{L}_W in Eq. 9 and Eq. 10	28.12	+0.59

Table 7: Ablation study on the WMT14 En→De task.

the number of sentences at each group. The results are shown in Figure 4. Compared to the baseline, the translations of our model are more concentrated into interval [60-80) and have less number in interval [0-40). The evaluation results show that the proposed CAML can accelerate the model to generate more stable results, which effectively improves the reliability of NMT.

Ablation study. To further show the effect of each module in our approach, we make an ablation study which is shown in Table 7. On the one hand, we implement three methods to analyze effect of \mathbf{x}^s , \mathcal{L}_S and \mathcal{L}_W . The first one is to use $\{\mathbf{x}^s, y\}$ with the original training set to train the model. The second one is to apply MAML as described in Eq.4-5. The third one is to use \mathcal{L}_S and \mathcal{L}_W in a multi-task learning (MTL) manner. The first and third methods can get 0.34 and 0.48 gains, respectively. It is unexpected that the MAML does not work. We believe one of the reasons is that tasks in two stages of MAML are too similar.

On the other hand, when ablating \mathcal{L}_S and \mathcal{L}_W , the BLEU score drops 0.37 and 0.52, respectively. We further attempts to add the \mathcal{L}_T to meta-train step, which leads to a 0.27 drop. We think the reason for this is the same as MAML. In summary, the CAML algorithm is conducive to training, making the model achieve better results. \mathcal{L}_S and \mathcal{L}_W play different roles and brings certain improvements to the model.

Related Work

Reliability vs. Robustness. There is some overlap between robustness and reliability in NMT (Sperber, Niehues, and Waibel 2017). Thus, some studies on the robustness of NMT can partially improve reliability. For example, Liu et al. (2019) proposed to incorporate noisy sentences into NMT training to address *homophone noise*, where words are replaced by others with similar pronunciations. Ebrahimi, Lowd, and Dou (2018); Cheng et al. (2020) proposed to use adversarial examples to address the *synonym problem*. However, robustness focuses on noisy inputs, which can deal with the errors in the sentence. Some pre-processing methods, e.g., grammar correction models, are another way to solve the robustness problem. Our work improves the reliability of NMT by addressing the source diversity phenomenon. Various expressions, all of which are proper, are more common and intractable than incorrect inputs.

Meta-learning in NLP. Meta-learning offers a strong generalization ability with limited data. Previous work has successfully employed model-agnostic meta-learning (MAML) (Finn, Abbeel, and Levine 2017) to various tasks, e.g., text classification (Bansal et al. 2020) and generation (Qian and Yu 2019). In NMT, Gu et al. (2018) proposed to improve the low-resource machine translation with MAML. Li et al. (2018a) proposed to improve the domain generalization of NMT with MAML. The proposed CAML is a derivative of the standard MAML, which focuses on a more fundamental issue of how to solve the problem of source diversity to improve the reliability of NMT.

Self-supervised training. Recently, self-supervised training has been widely used in NLP (Devlin et al. 2019; Brown et al. 2020; Lample and Conneau 2019). In NMT, Sennrich, Haddow, and Birch (2016a) firstly proposed back-translation to translate target monolingual data to synthetic parallel data. Caswell, Chelba, and Grangier (2019) proposed to add a tag at the source side of the synthetic data. These data-augmentation methods objectively enhance the ability of NMT to deal with source diversity. Wei et al. (2020, 2022) proposed to model continuous semantics towards reducing the source uncertainty problem. Following these studies, our work adopts the self-supervised method to learn the consistent representation from semantically equivalent sentences.

Conclusion

In this work, we clarify the shortcoming of current NMT models, specifically models trained by the one-to-one paradigm are hard to deal with the multi-to-one scenario. To address this problem, we present a consistency-aware meta-learning (CAML) framework derived from the model-agnostic meta-learning algorithm. In the proposed CAML, two customized training objectives are used with the MAML manner to model the consistent representation, thereby enhancing the reliability of the NMT model. Extensive experiments show that CAML effectively improves translation quality by producing stable outputs for diverse inputs.

Acknowledgments

We would like to thank the reviewers for their insightful comments. This work is supported by the National Natural Science Foundation of China (Grant No. 62036004).

References

- Bahdanau, D.; Cho, K.; and Bengio, Y. 2014. Neural Machine Translation by Jointly Learning to Align and Translate. *arXiv:1409.0473*.
- Bansal, T.; Jha, R.; Munkhdalai, T.; and McCallum, A. 2020. Self-Supervised Meta-Learning for Few-Shot Natural Language Classification Tasks. In *EMNLP*, 522–534.
- Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J. D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. 2020. Language models are few-shot learners. *Neurips*, 1877–1901.
- Caswell, I.; Chelba, C.; and Grangier, D. 2019. Tagged back-translation. *arXiv:1906.06442*.
- Cheng, Y.; Jiang, L.; and Macherey, W. 2019. Robust Neural Machine Translation with Doubly Adversarial Inputs. In *ACL*, 3044–3049.
- Cheng, Y.; Jiang, L.; Macherey, W.; and Eisenstein, J. 2020. AdvAug: Robust Adversarial Augmentation for Neural Machine Translation. In *ACL*, 5961–5970.
- Cheng, Y.; Tu, Z.; Meng, F.; Zhai, J.; and Liu, Y. 2018. Towards Robust Neural Machine Translation. In *ACL*, 1756–1766.
- Cho, K.; van Merriënboer, B.; Gulcehre, C.; Bahdanau, D.; Bougares, F.; Schwenk, H.; and Bengio, Y. 2014. Learning Phrase Representations using RNN Encoder–Decoder for Statistical Machine Translation. In *EMNLP*, 1724–1734.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *NAACL*, 4171–4186.
- Ebrahimi, J.; Lowd, D.; and Dou, D. 2018. On Adversarial Examples for Character-Level Neural Machine Translation. In *Coling*, 653–663.
- Finn, C.; Abbeel, P.; and Levine, S. 2017. Model-agnostic meta-learning for fast adaptation of deep networks. In *ICML*, 1126–1135.
- Gao, F.; Zhu, J.; Wu, L.; Xia, Y.; Qin, T.; Cheng, X.; Zhou, W.; and Liu, T.-Y. 2019. Soft Contextual Data Augmentation for Neural Machine Translation. In *ACL*, 5539–5544.
- Ghazvininejad, M.; Levy, O.; Liu, Y.; and Zettlemoyer, L. 2019. Mask-Predict: Parallel Decoding of Conditional Masked Language Models. In *EMNLP-IJCNLP*, 6112–6121.
- Gu, J.; Bradbury, J.; Xiong, C.; Li, V. O.; and Socher, R. 2017. Non-autoregressive neural machine translation. *arXiv preprint arXiv:1711.02281*.
- Gu, J.; Wang, Y.; Chen, Y.; Li, V. O.; and Cho, K. 2018. Meta-Learning for Low-Resource Neural Machine Translation. In *EMNLP*, 3622–3631.
- Hao, J.; Wang, X.; Shi, S.; Zhang, J.; and Tu, Z. 2019. Multi-Granularity Self-Attention for Neural Machine Translation. In *EMNLP-IJCNLP*, 887–897.
- He, P.; Meister, C.; and Su, Z. 2020. Structure-invariant testing for machine translation. In *ICSE*, 961–973. IEEE.
- Hoang, V. C. D.; Koehn, P.; Haffari, G.; and Cohn, T. 2018. Iterative Back-Translation for Neural Machine Translation. In *NMTG*, 18–24.
- Jiao, W.; Wang, X.; He, S.; King, I.; Lyu, M.; and Tu, Z. 2020. Data Rejuvenation: Exploiting Inactive Training Examples for Neural Machine Translation. In *EMNLP*, 2255–2266.
- Lample, G.; and Conneau, A. 2019. Cross-lingual Language Model Pretraining. *arXiv preprint arXiv:1901.07291*.
- Li, D.; Yang, Y.; Song, Y.-Z.; and Hospedales, T. 2018a. Learning to generalize: Meta-learning for domain generalization. In *AAAI*, volume 32.
- Li, J.; Tu, Z.; Yang, B.; Lyu, M. R.; and Zhang, T. 2018b. Multi-Head Attention with Disagreement Regularization. In *EMNLP*, 2897–2903.
- Liu, H.; Ma, M.; Huang, L.; Xiong, H.; and He, Z. 2019. Robust Neural Machine Translation with Joint Textual and Phonetic Embedding. In *ACL*, 3044–3049.
- Liu, Y.; Gu, J.; Goyal, N.; Li, X.; Edunov, S.; Ghazvininejad, M.; Lewis, M.; and Zettlemoyer, L. 2020. Multilingual Denoising Pre-training for Neural Machine Translation. *TACL*, 726–742.
- Luong, M.; Pham, H.; and Manning, C. D. 2015. Effective Approaches to Attention-based Neural Machine Translation. In *EMNLP*, 1412–1421.
- Neubig, G.; Dou, Z.; Hu, J.; Michel, P.; Pruthi, D.; Wang, X.; and Wieting, J. 2019. compare-mt: A Tool for Holistic Comparison of Language Generation Systems. *CoRR*, abs/1903.07926.
- Nguyen, X.-P.; Joty, S.; Kui, W.; and Aw, A. T. 2020. Data Diversification: A Simple Strategy For Neural Machine Translation. In *NIPS*, 10018–10029.
- Ott, M.; Auli, M.; Grangier, D.; and Ranzato, M. 2018. Analyzing uncertainty in neural machine translation. In *ICML*, 3956–3965.
- Papineni, K.; Roukos, S.; Ward, T.; and Zhu, W.-J. 2002. BLEU: A Method for Automatic Evaluation of Machine Translation. In *ACL*, 311–318.
- Qian, K.; and Yu, Z. 2019. Domain Adaptive Dialog Generation via Meta Learning. In *ACL*, 2639–2649.
- Raffel, C.; Shazeer, N.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; Liu, P. J.; et al. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *JMLR*, 21(140): 1–67.
- Sennrich, R.; Haddow, B.; and Birch, A. 2016a. Improving Neural Machine Translation Models with Monolingual Data. In *ACL*, 86–96.
- Sennrich, R.; Haddow, B.; and Birch, A. 2016b. Neural Machine Translation of Rare Words with Subword Units. In *ACL*, 1715–1725.
- Sperber, M.; Niehues, J.; and Waibel, A. 2017. Toward robust neural machine translation for noisy input sequences. In *ICSLT*, 90–96.

Sutskever, I.; Vinyals, O.; and Le, Q. V. 2014. Sequence to Sequence Learning with Neural Networks. In *NIPS*, volume 27.

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*, volume 30.

Wang, Q.; Li, B.; Xiao, T.; Zhu, J.; Li, C.; Wong, D. F.; and Chao, L. S. 2019. Learning Deep Transformer Models for Machine Translation. In *ACL*, 1810–1822.

Wang, X.; Pham, H.; Dai, Z.; and Neubig, G. 2018. SwitchOut: an Efficient Data Augmentation Algorithm for Neural Machine Translation. In *EMNLP*, 856–861.

Wang, X.; Tsvetkov, Y.; and Neubig, G. 2020. Balancing Training for Multilingual Neural Machine Translation. In *ACL*, 8526–8537.

Wang, Z.; Hu, G.; and Hu, Q. 2020. Training noise-robust deep neural networks via meta-learning. In *CVPR*, 4524–4533.

Wei, X.; Yu, H.; Hu, Y.; Weng, R.; Luo, W.; and Jin, R. 2022. Learning to Generalize to More: Continuous Semantic Augmentation for Neural Machine Translation. In *ACL*, 7930–7944.

Wei, X.; Yu, H.; Hu, Y.; Weng, R.; Xing, L.; and Luo, W. 2020. Uncertainty-Aware Semantic Augmentation for Neural Machine Translation. In *EMNLP*, 2724–2735.

Wu, M.; Li, Y.; Zhang, M.; Li, L.; Haffari, G.; and Liu, Q. 2021. Uncertainty-Aware Balancing for Multilingual and Multi-Domain Neural Machine Translation Training. In *EMNLP*, 7291–7305.

Yang, B.; Li, J.; Wong, D. F.; Chao, L. S.; Wang, X.; and Tu, Z. 2019. Context-aware self-attention networks. In *AAAI*, 387–394.

Yao, H.; Wang, Y.; Wei, Y.; Zhao, P.; Mahdavi, M.; Lian, D.; and Finn, C. 2021. Meta-learning with an Adaptive Task Scheduler. In *Neurips*, 7497–7509.

Zhou, Y.; Yang, B.; Wong, D. F.; Wan, Y.; and Chao, L. S. 2020. Uncertainty-Aware Curriculum Learning for Neural Machine Translation. In *ACL*, 6934–6944.

Zhu, J.; Xia, Y.; Wu, L.; He, D.; Qin, T.; Zhou, W.; Li, H.; and Liu, T.-Y. 2020. Incorporating BERT into Neural Machine Translation. In *arXiv:2002.06823*.