

ConsistTL: Modeling Consistency in Transfer Learning for Low-Resource Neural Machine Translation

Zhaocong Li¹ Xuebo Liu^{2*} Derek F. Wong^{1*} Lidia S. Chao¹ Min Zhang²

¹NLP²CT Lab, Department of Computer and Information Science, University of Macau
nlp2ct.zhaocong@gmail.com, {derekfw, lidiasc}@um.edu.mo

²Institute of Computing and Intelligence, Harbin Institute of Technology, Shenzhen, China
{liuxuebo, zhangmin2021}@hit.edu.cn

Abstract

Transfer learning is a simple and powerful method that can be used to boost model performance of low-resource neural machine translation (NMT). Existing transfer learning methods for NMT are static, which simply transfer knowledge from a parent model to a child model once via parameter initialization. In this paper, we propose a novel transfer learning method for NMT, namely ConsistTL, which can continuously transfer knowledge from the parent model during the training of the child model. Specifically, for each training instance of the child model, ConsistTL constructs the semantically-equivalent instance for the parent model and encourages prediction consistency between the parent and child for this instance, which is equivalent to the child model learning each instance under the guidance of the parent model. Experimental results on five low-resource NMT tasks demonstrate that ConsistTL results in significant improvements over strong transfer learning baselines, with a gain up to 1.7 BLEU over the existing back-translation model on the widely-used WMT17 Turkish-English benchmark. Further analysis reveals that ConsistTL can improve the inference calibration of the child model. Code and scripts are freely available at <https://github.com/NLP2CT/ConsistTL>.

1 Introduction

Neural machine translation (NMT) has achieved success on the high-resource language pairs (Vaswani et al., 2017). However, it achieves an inadequate performance on the low-resource language pairs with merely bilingual data due to the data sparsity (Koehn and Knowles, 2017; Sennrich and Zhang, 2019). Transfer learning is a simple and powerful method that can be used to boost the model performance of low-resource NMT, which can transfer knowledge from an off-the-shelf high-

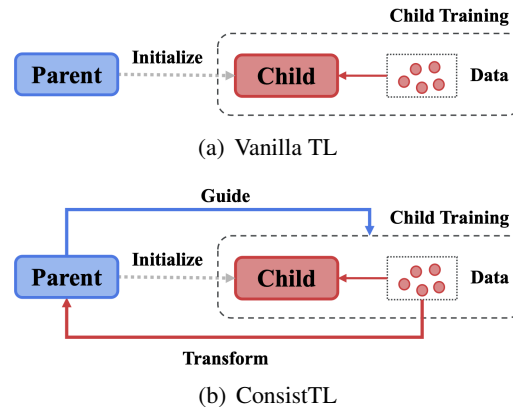


Figure 1: Comparison between vanilla TL and ConsistTL. The dashed arrow represents the initialization of the parameters from the parent model. Apart from disposable parameter transfer, ConsistTL continuously transfers knowledge from the parent model during the child training.

resource parent model to the low-resource child model (Zoph et al., 2016).

The goal of transfer learning for low-resource NMT is to transfer knowledge sufficiently from the parent model. Prior works attempt to transfer more information from the embedding layer of the parent model, using methods such as building a joint dictionary (Kocmi and Bojar, 2018; Gheini and May, 2019), creating cross-lingual token mapping (Kim et al., 2019), and transferring the word of same form (Aji et al., 2020). These strategies achieve success in parameter transfer. The methods used in these works is summarized in Figure 1(a): the parent model transfers its parameters to the child model once and lacks continual guidance on the child training. This method can not utilize the parent model sufficiently, which potentially limits the learning of the child model.

This paper attempts to leverage a continual knowledge transfer from the parent model for the child training. The prediction distribution of the parent model is usually more informative since

*Co-corresponding author

it has learned to predict various translations during training (Dreyer and Marcu, 2012; Khayrallah et al., 2020). Therefore, utilizing the prediction of the parent model to continuously guide the child training might be a nice plus.

To this end, we propose **Consistency-based Transfer Learning (ConsistTL)**, a novel framework that can be used to model the consistency between the parent model and the child model during the child training. For each instance from the child data, ConsistTL constructs a semantically-equivalent pseudo source sentence for the parent by translating the child target sentence. Then ConsistTL encourages the child prediction to be consistent with the parent prediction since these two predictions are conditioned on the same target sentence and semantically-equivalent source sentences. ConsistTL is summarized in Figure 1(b): unlike in vanilla TL, the parent model releases continual guidance for the child model during training by receiving the transformation of the training data of the child model. With ConsistTL, the child model can acquire more knowledge from the parent predictions beyond the simple parameter transfer.

We conduct experiments on five low-resource NMT benchmarks. Experimental results show that ConsistTL achieves significant improvements over strong baselines of transfer learning. Combined with the representative data augmentation method of back-translation (Sennrich et al., 2016a), our method can outperform the existing method up to 1.7 BLEU on the widely-used WMT17 Turkish-English benchmark. An extensive analysis reveals that our method, as an approach to model external consistency, is complementary to the method of modeling inner consistency.

Our contributions are as follows:

- We propose ConsistTL to model the consistency between the parent and child models in transfer learning for low-resource NMT. The cross-model consistency extends the knowledge transfer process to child training.
- ConsistTL outperforms strong baselines significantly on five low-resource benchmarks, and is complementary to the other methods for low-resource NMT (e.g., back-translation and inner-model consistency learning).
- We find that ConsistTL can improve the calibration of the child model by increasing translation accuracy and reducing overconfidence.

2 Related Work

Transfer Learning for NMT Zoph et al. (2016) propose a parent-child framework to transfer the parameters from the parent model trained on a high-resource corpus to a low-resource child model with the same target language. The follow-up works aim to transfer the knowledge from the parent model parameters more sufficiently since Zoph et al. (2016) ignore the embedding layer mismatch caused by the vocabulary mismatch. These works include building the shared joint dictionary (Kocmi and Bojar, 2018; Gheini and May, 2019; Liu et al., 2019a,b), using an extra trained transformation to connect the embedding layers of two models (Kim et al., 2019; Sato et al., 2020; Liu et al., 2021a), or transferring partial word embedding from the parent model (Aji et al., 2020; Xu and Hong, 2022). In comparison, this presenting paper tries to transfer the knowledge contained in the parent prediction beyond parameters.

There are also some works that consider multilingualism to enhance transfer learning. Neubig and Hu (2018); Tan et al. (2019) propose to transfer knowledge from a multilingual parent model. Gu et al. (2018a) develop a universal architecture with shared lexical and sentence level representation to implement transfer learning. Gu et al. (2018b) propose a meta-learning approach to learn the initialization process using several high-resource language pairs. In this paper, we mainly focus on the bilingual parent model, since off-the-shelf bilingual models are easier to obtain.

Consistency Learning in NMT Consistency learning aims to model consistency across different model predictions. The core idea is to craft the cross-view supervision and enhance cross-view consistency, which enriches the supervision signals. In the early studies, consistency learning methods focus on semi-supervised learning since consistency signals enable models to learn from data without manual labels (Xie et al., 2020; Miyato et al., 2019; Lowell et al., 2021a). The semi-supervised consistency learning also enhances NMT models with the combination of forward-translation (Clark et al., 2018). Recent studies of consistency learning gradually consider supervised learning areas (Chen et al., 2021; Lowell et al., 2021b). R-Drop is one of the simple and generalized methods (Liang et al., 2021), which enhances the consistency between multiple model structures produced by dropout,




Figure 2: Overview of ConsistTL. ConsistTL first constructs semantically-equivalent pseudo parent source sentences from child target sentences using the reversed parent model. Then, given the sentence group with the same semantic meaning, ConsistTL encourages the child prediction to be consistent with the parent prediction during child training.

and improves NMT performance.

There are some consistency learning methods specially designed for supervised NMT based on data augmentation. Shen et al. (2020); Guo et al. (2022); Gao et al. (2022) leverage token-level perturbations to craft different predictions. Kambhatla et al. (2022) explore the insight of cryptology to create a ciphertext with the same semantic meaning for each training instance. Xie et al. (2021) develop a scheme to augment target side input by the output prediction from the first forward. Wang et al. (2022) enhance the Chinese representation by modeling the consistency between stroke representation and its corresponding encryption. These consistency modeling works in NMT mainly exploit the inner model consistency, i.e., modeling consistency within the same model. Compared with these works, we exploit transferable information from the cross-model consistency in transfer learning for low-resource NMT. Chen et al. (2017) also exploit the consistency between two models and the differences from us are: 1) we construct the parent data instead of the child data; 2) we enhance the transfer learning framework for low-resource NMT while their method relies on a pivot language to improve zero-resource NMT.

3 ConsistTL

3.1 Motivation

Previous works concentrate on knowledge transfer from the parent model parameters (Kocmi and Bojar, 2018; Kim et al., 2019; Aji et al., 2020). This paper argues that continual knowledge transfer from the predictions of the parent model is also beneficial to transfer learning. Since the parent

model is trained on a large number of diverse translations (Dreyer and Marcu, 2012; Khayrallah et al., 2020), the prediction of the parent model might be informative for child training.

In this paper, we propose ConsistTL, a framework that can be used to jointly transfer the knowledge from the parent parameters and the prediction distribution of the parent model. ConsistTL can enhance the consistency between the distribution from the parent model (e.g., German-English) and the child model (e.g., Turkish-English) conditioned on the semantically-equivalent sentences. ConsistTL is composed of two parts: semantically-equivalent parent data construction and parent-child cross-model consistency learning. The overall framework is illustrated in Figure 2.

3.2 Semantically-Equivalent Parent Data Construction

To implement the proposed ConsistTL, we need to model the cross-model consistency between the parent-side source sentence and child-side source sentence. As the parent model and child model are trained for different language pairs, for each instance in the child data, we need a semantically-equivalent sentence written in the source language of the parent language pair.

However, during the child training, we may have only the small-scale bilingual child data $\mathcal{D}_c = \{\langle \mathbf{x}_c^n, \mathbf{y}^n \rangle\}_{n=1}^N$. Acquiring a semantically-equivalent source sentence for each instance of the parent in \mathcal{D}_c is labor-consuming. Inspired by back-translation (Sennrich et al., 2016a), ConsistTL creates synthetic parent data by generating pseudo source sentences $\tilde{\mathbf{x}}_p$ from \mathbf{x}_c , or \mathbf{y} .

The translation quality is crucial for generating pseudo sentences since low-quality translation would cause poor semantic equivalence between the generated sentence \tilde{x}_p and the existing sentence pair $\langle x_c, y \rangle$. Obviously, the reversed parent model $\mathcal{M}_{y \rightarrow x_p}$ is easier to obtain and tends to produce higher quality translation than the low-resource translation model $\mathcal{M}_{x_c \rightarrow x_p}$. Therefore, we choose to back-translate the target sentence y into \tilde{x}_p using the reversed parent model $\mathcal{M}_{y \rightarrow x_p}$ in ConsistTL. Then, we can concatenate the pseudo source data $\{\tilde{x}_p^n\}_{n=1}^N$ with the target side data $\{y^n\}_{n=1}^N$ from the child language pair resulting in the synthetic parent data $\tilde{D}_p = \{\langle \tilde{x}_p^n, y^n \rangle\}_{n=1}^N$. After constructing the synthetic parent data, ConsistTL steps into the next stage to take advantage of this synthetic parent data.

3.3 Parent-Child Cross-Model Consistency Learning

In this part, we encourage the child model to learn the cross-model consistency information from the child data and synthetic parent data. Every instance $\langle x_c, y \rangle$ from the child data \mathcal{D}_c and its semantically-equivalent counterpart $\langle \tilde{x}_p, y \rangle$ from the synthetic parent data \tilde{D}_p are fed into the child and parent models, respectively, at the same time. The corresponding prediction distributions can be obtained from their forwards. Our goal is to make these two distributions as consistent as possible. We propose a cross-model consistency loss to make the parent model distribution and the child model distribution closer during child training.

Cross-Model Consistency Loss To make the distribution of the parent model and child model consistent, we need to measure and minimize the gap between these distributions. The proposed framework minimizes the distribution gap with the following cross-model loss:

$$\mathcal{L}_d(\theta_c) = \sum_{t=1}^T F[P(*|\mathbf{y}_{<t}, \mathbf{x}_c, \theta_c, \tau), P(*|\mathbf{y}_{<t}, \tilde{\mathbf{x}}_p, \theta_p^*, \tau)] \quad (1)$$

where F measures the distribution gap between parent model and child model. We choose Jensen–Shannon divergence (Lin, 1991) as F . θ_p^* and θ_c represent parameters of parent model and child model respectively. $P(*|\mathbf{y}_{<t}, \tilde{\mathbf{x}}_p, \theta_p^*, \tau)$ and $P(*|\mathbf{y}_{<t}, \mathbf{x}_c, \theta_c, \tau)$ denote the prediction distributions of parent model and child model at time step

t respectively. The parent parameters are fixed during child training as it is well trained on the parent language pair. We use τ to control the smoothness of the prediction distribution:

$$P(y_t = k|\mathbf{y}_{<t}, \mathbf{x}, \theta) = \frac{\exp(z_k/\tau)}{\sum_{i=1}^{|V|} \exp(z_i/\tau)} \quad (2)$$

where z_i denotes the logits output before softmax is computed, and $|V|$ represents the size of the target-side vocabulary. We can smooth the distribution of the acquired prediction with $\tau > 1$.

Final Child Training Objective Function We interpolate the weighted cross-model consistency loss and negative log-likelihood loss to create the final child training objective function. The final child training objective is formulated as:

$$\mathcal{L}(\theta_c) = \mathcal{L}_{null}(\theta_c) + \alpha \mathcal{L}_d(\theta_c) \quad (3)$$

In this manner, the child model can learn each instance under the guidance of the parent model.

4 Experiments

4.1 Datasets

Parent Language Pair This paper follows Aji et al. (2020) to transfer knowledge from the German-English (De-En) task. Our parent model is trained on the WMT17 De-En training set and validated on *newstest2013*. The training set consists of 5.8M sentences. The vocabulary is implemented using joint source-target BPE with 40k merge operations (Sennrich et al., 2016b).

Child Language Pairs We conduct transfer learning experiments on five low-resource translation benchmarks. There are four translation benchmarks from Global Voices (Tiedemann, 2012; Khayrallah et al., 2020): Polish (Pl), Hungarian (Hu), Indonesian (Id), and Catalan (Ca). The sentences of corresponding languages are paired with English sentences. The subset splits follow Khayrallah et al. (2020). For the fifth language pair, we adopt WMT17 Turkish (Tr)-English benchmark and use *newstest2016* as the validation set. Before subword segmentation, we apply normalization and tokenization to all datasets using the Moses (Koehn et al., 2007). The sentences with more than 60 words and a length ratio of over 1.5 in training data would be filtered for the Tr-En language pair. After data filtering, we then apply joint source-target BPE to the child language pairs with fewer merge

Task	Train	Valid	Test	BPE
Tr-En	192,614	3,000	3,007	16,000
Pl-En	39,943	2,000	2,000	8,000
Ca-En	15,176	2,000	2,000	8,000
Id-En	8,448	2,000	2,000	8,000
Hu-En	7,712	2,000	2,000	8,000

Table 1: Numbers of sentences and BPE merge operations of preprocessed datasets.

operations than were used for the parent language pair. We display detailed statistics of the preprocessed datasets in Table 1.

4.2 Model Configuration

We use the Transformer-base architecture (Vaswani et al., 2017) to implement NMT models in our experiments from the toolkit FAIRSEQ (Ott et al., 2019). For all the low-resource translation models, we tie the input embedding layers of the decoder and the output projections (Press and Wolf, 2017; Inan et al., 2017). For the high-resource parent model, we tie all embedding layers.

Vanilla NMT We implement the transformer models training from scratch as baselines. These models share the vocabularies of source and target languages.

TL Zoph et al. (2016) propose a transfer learning method to randomly assign a parent word embedding to each child word. The rest of the parameters are copied from the parent model. We implement this method as a baseline in our experiment. This vanilla transfer learning method is marked as “TL”.

TM-TL Aji et al. (2020) propose a method named “token matching” (TM), which is simple and effective. We implement this transfer learning method as a baseline in our experiments marked as “TM-TL”. Some tokens are common to both parent model and child model vocabularies for child language. TM-TL assigns the embedding parameters of such common tokens in the source embedding layer of the child model. For the other parameters of the source embedding layer, TM-TL initializes them randomly following the vanilla NMT. The rest settings of TM-TL follow TL.

ConsistTL Since our method, ConsistTL, transfers knowledge from the prediction distribution of the parent model, it is orthogonal to the parame-

ter transfer methods. We adopt TM to implement ConsistTL.

4.3 Settings

Training We train all the NMT models using the Adam optimizer (Kingma and Ba, 2015) with $\beta = (0.9, 0.98)$. We also use the inverse square root schedule to control the variation of the learning rate during training. For the parent model, the value of linear warmup steps is set to 10,000 and the peak learning rate is set to 0.001. The parent model is trained for 200 epochs with 460K (3584×128) tokens per batch and a low dropout rate of 0.1. We use 4,000/2,000 max tokens per batch for Tr/Pl-En tasks, and 1,000 for the other tasks. For the the training of the low-resource translation model from scratch, we set the number of linear warm-up steps to 8,000 and the peak learning rate is set to 0.0005. For the low-resource translation models with transfer learning, as the inner layers are adequately trained on the parent language pair, we set narrow warm-up steps to 1,000. We set a lower peak learning rate to 0.0002 to prevent overfitting. All the low-resource translation models are trained for 200 epochs. To prevent the models from overfitting on low-resource language pairs, we set the dropout rate to 0.3, both the attention dropout rate and activation dropout rate to 0.1. For checkpoint selection, the checkpoints with the best validation BLEU would be selected as the final model checkpoints. We train an En-De model on the same parent data as the reversed parent model. We set τ to 1.0 and tune α on $\{3, 4, 5, 6, 7\}$.

Evaluation We used beam search with a beam width of 5 and a length penalty of 1 to evaluate all the NMT models. And we also use such decoding setting to generate pseudo parent source sentences. The translation quality is evaluated using SacreBLEU (Post, 2018)¹ and the BERTScore (Zhang et al., 2020)², which evaluate the surface-form correctness and semantic correctness, respectively.

4.4 Main Results

Table 2 displays the results of the five aforementioned translation tasks in terms of the BLEU score and the BERTScore. We report the existing results from Khayrallah et al. (2020) and Baziotis et al.

¹Signature: nrefs:1 + case:mixed + eff:no + tok:13a + smooth:exp + version:2.0.0

²https://github.com/Tiiiger/bert_score

Model	Id-En		Ca-En		Hu-En		Pl-En		Tr-En	
	BLEU	BS	BLEU	BS	BLEU	BS	BLEU	BS	BLEU	BS
Existing	12.3	-	20.0	-	5.4	-	18.0	-	18.5	-
Vanilla	1.1	13.2	1.1	15.5	0.9	0.9	1.5	18.9	17.8	51.8
+ TL	13.5	37.7	21.6	51.8	5.9	27.4	19.9	55.3	17.6	51.9
+ TM-TL	18.6	49.9	25.3	58.9	10.6	41.2	21.4	58.2	18.6	53.9
+ ConsistTL	19.7	52.2	26.6	60.0	11.9	43.9	22.4	59.9	19.3	55.9

Table 2: Main results of the five translation tasks. BS denotes the BERTScore. “Existing” are reported from Khayrallah et al. (2020); Baziotis et al. (2020). Both BLEU score and BERTScore reflect that ConsistTL achieves significant improvements compared with the strong baseline TM-TL on all the tasks ($p < 0.01$) (Koehn, 2004).

(2020). ConsistTL outperforms the existing works on the same test sets in terms of the BLEU score.

Global Voices {Id,Ca,Hu,Pl}-En The data scale of the four language pairs (Id-En, Ca-En, Hu-En, Pl-En) collected from Global Voices is extremely lower than that of WMT17 Tr-En. The parameter transfer methods TL and TM-TL are strong since they outperform vanilla NMT in terms of both BLEU and BERTScore. Our method ConsistTL achieves a significant average improvement of 1.1 BLEU and 2.0 BERTScore over the TM-TL.

WMT17 Tr-En Different from the four aforementioned language pairs, the data scale of Tr-En is larger. Compared with the aforementioned extremely low-resource language pairs, the transfer learning baseline methods have lower performance gains (TM-TL) or even no performance gains (TL) over Vanilla NMT on this benchmark with larger data. Although baseline methods are less informative on the Tr-En, our method still achieves a significant improvement over the strong baseline TM-TL with 0.7 BLEU and 2.0 BERTScore. The results of BERTScore on five translation benchmarks indicate that ConsistTL provides consistent and significant improvements of the semantic correctness. Our follow-up analysis is based on the WMT17 Tr-En benchmark.

5 Analysis

5.1 Effect of Loss Type

There are many ways to measure the difference between two distributions. Different measure functions would result in different consistency loss types. This part compares two typical choices Jensen–Shannon divergence (Lin, 1991) and Kullback–Leibler divergence (Kullback and Leibler, 1951) to study the effect of the chosen loss type

Tr-En	Consistency Loss	Valid	Test
TL	-	17.8	17.6
TM-TL	-	18.8	18.6
ConsistTL	Kullback-Leibler	19.2	18.9
	Jensen–Shannon	19.8	19.3

Table 3: Comparison of different consistency loss. ConsistTL works better with JS loss than KL loss.

on ConsistTL. To make the most of their capabilities, we tune the hyperparameters (α, τ) for each loss type. The final hyperparameters of the JS loss and the KL loss are (7.0, 1.0) and (4.0, 2.0) respectively. As seen in Table 3, the model with JS achieves the best improvements, while the model with KL loss realizes limited improvements. Perhaps minimizing the KL divergence is more difficult than minimizing the JS divergence since we observe that lowering the KL loss through smoothing ($\tau > 1$) is helpful.

5.2 Comparison of Learning Curves

The learning curve is one way to describe the temporal effect of the proposed method during training, which has also been adopted in previous consistency modeling works (Liang et al., 2021; Kambhatla et al., 2022; Xie et al., 2021). This part compares the proposed method, ConsistTL, with the baseline methods TL and TM-TL in terms of their learning curves described by validation set BLEU. Figure 3 displays the learning curves of these three transfer learning methods. During the child training, the ConsistTL and TM-TL grow into larger validation BLEU values with higher convergence speeds. The difference between ConsistTL and TM-TL is invisible in the early learning phase and gradually changes into a stable and signifi-




Figure 3: Learning curves of different TL methods. ConsistTL achieves the largest upper bound among them.

Tr-En	Parent Source	Valid	Test
TL	-	17.8	17.6
TM-TL	-	18.8	18.6
	Sampling	19.3	18.8
ConsistTL	Greedy Search	19.8	19.3
	Beam Search	19.8	19.3

Table 4: Effect of pseudo parent source generation. Both greedy search and beam search can ensure the semantic equivalence of pseudo parent sources.

cant value. Finally, ConsistTL achieves the largest upper bound among these three learning curves, which indicates that ConsistTL is informative to child training.

5.3 Effect of Pseudo Source Generation

Although beam search is a de facto decoding method that is used to generate sentences in the NMT task, there are also other decoding methods that can be used to generate synthetic data (Ott et al., 2018; Edunov et al., 2018), including greedy search, and greedy sampling. To study the effect of decoding methods in parent-side synthetic data generation, we compare the results produced by beam search, greedy search and greedy sampling. As shown in Table 4, we find that greedy sampling produces lower improvement in contrast to the beam search. This can be interpreted to mean that greedy sampling introduces low-quality translations due to the randomness during decoding (Ippolito et al., 2019). Though greedy search performs worse than beam search in usual model testing, the final results in Table 4 empirically shows that there is no significant difference in final results between beam search and greedy search. That means the translation quality of greedy search meets the ConsistTL requirement. Therefore, we can adopt the greedy

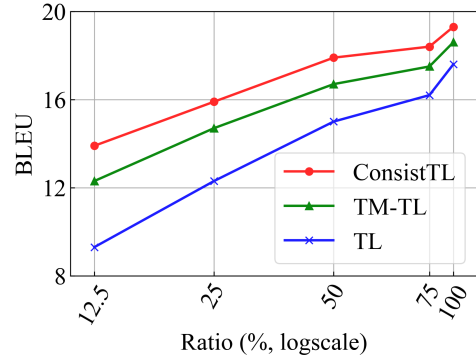


Figure 4: BLEU score with varying data scale.

Tr-En	w/ BT	Valid	Test
TM-TL (Aji et al., 2020)	Y	-	20.6
TM-TL (Ours)	N	18.8	18.6
	Y	22.3	21.6
ConsistTL	N	19.8	19.3
	Y	22.9	22.3

Table 5: Effect of combining with back-translation. ConsistTL still outperforms TM-TL+BT significantly.

search to generate pseudo parent source sentences for future usage.

5.4 Effect of Data Scale

To verify the robustness of the improvements brought by ConsistTL across data scales, we compare our method and the baselines methods, TM-TL and TL, on the subsets with different scales. The training subsets are randomly sampled from the full Tr-En training data according to the specific ratios. Figure 4 demonstrates the result of data ablation. The performances of ConsistTL and TM-TL are more stable than TL on different data scales, since TL drops more performance on smaller-scale data. ConsistTL achieves significant improvements consistently in contrast to TM-TL across different data scales. This result indicates the robust effectiveness of ConsistTL with data scales varying.

5.5 Effect of Back-Translation

Back-translation (Sennrich et al., 2016a; Edunov et al., 2018) is an essential method for NMT that is used to augment training corpus by generating synthetic sentences from target-side monolingual data. In this part, to study the complementarity between this work and back-translation, we compare the proposed method and the baseline TM-TL on the training data augmented by back-translation.

Tr-En	w/ Inner Consist.	Valid	Test
TM-TL	N	18.8	18.6
	Y	20.1	20.0
ConsistTL	N	19.8	19.3
	Y	20.8	20.6

Table 6: Effect of combining with inner consistency modeling.

We sample 200k English monolingual data from News Crawl 2015 for back-translation, following the approximate 1:1 ratio according to Aji et al. (2020). In Aji et al. (2020), the parent model for Tr-En is also a Transformer-base model trained on WMT17 De-En data. Table 5 displays the results reported from Aji et al. (2020) and the results of our implementation. Our proposed method ConsistTL outperforms the existing baselines with 1.7 BLEU. In combination with back-translation, our method yields a performance gain with the significance of 95% ($p < 0.01$) on the test set, which means ConsistTL is complementary to back-translation. This result confirms the generality of ConsistTL since ConsistTL can be generalized to a scenario of in which training data are augmented by back-translation.

5.6 Combination with Inner Consistency

As ConsistTL learns the cross-model consistency, here we investigate the effectiveness of such cross-model consistency under inner consistency modeling. R-Drop is a typical method to model inner consistency (Liang et al., 2021). We compare the transfer learning methods combined with inner consistency implemented by R-Drop. Table 6 shows that our method achieves the best performance and consistent improvements across the validation set and test set ($p < 0.01$). This result validates that the cross-model consistency is indeed effective since ConsistTL still improves performance when combined with R-Drop.

5.7 Model Calibration

As this paper proposes leveraging the prediction distribution of parent model, the following questions arises: *how does the parent prediction help shape the child prediction?* The inference calibration can be used to describe the prediction distribution for NMT models, which requires mitigating the gap between the prediction confidence and translation accuracy during inference (Wang

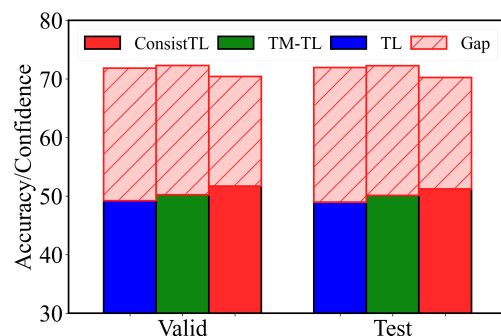


Figure 5: Visualization of prediction confidence and accuracy for the Tr-En model. The complete column represents the confidence while the pure color column represents the accuracy. The striped part denotes the gap between confidence and accuracy. ConsistTL narrows the gap and achieves the best prediction accuracy, resulting in better model calibration.

et al., 2020). We exploit the gap between the averaged confidence and translation accuracy to evaluate inference calibration³. The average accuracy illustrates the correctness of generated tokens, which is annotated by the translation error rate (TER) (Snover et al., 2006). The averaged confidence is computed on the estimated probabilities of generated tokens. More details can refer to Wang et al. (2020). Figure 5 shows the averaged accuracy and the averaged confidence of child models implemented by ConsistTL, TM-TL, and TL. The visualization results show that the averaged confidences are higher than the corresponding averaged accuracy, namely overconfidence of the models. Compared with other methods, ConsistTL improves the translation accuracy and reduces overconfidence. In this way, ConsistTL narrows the gap between the confidence and accuracy, which indicates that ConsistTL can be beneficial to child model inference calibration by introducing the prediction distribution of the parent model. This provides a reasonable explanation of why ConsistTL can give a performance boost to the child model.

6 Conclusion and Future Work

We introduce ConsistTL for transfer learning in low-resource NMT, which can continuously transfer the knowledge from parent prediction during child training. In order to transfer knowledge from parent prediction, we propose to model cross-model consistency between parent model and child model with the child source sentences and pseudo parent source sentences. Experimental results on

³<https://github.com/shuo-git/InfECE>

five translation benchmarks verify the effectiveness of ConsistTL, which can significantly improve the translation performance. ConsistTL is complementary to other powerful methods for low-resource NMT (e.g., back-translation). Further analysis reveals that introducing parent prediction is helpful to shape child model prediction distribution, resulting in better inference calibration.

In the future, we would like to apply curriculum learning (Liu et al., 2020; Zhan et al., 2021) to better organize the learning of the child model. It is also worthwhile to enhance the parent and child models by utilizing pre-trained knowledge learned from unlabeled data (Liu et al., 2021b,c).

Limitations

Our proposed framework would occupy extra computation resources compared with baseline methods. In the part of semantically-equivalent parent data construction, we need to train an additional reversed parent model to back-translate the target sentences of child data. In the part of parent-child cross-model consistency learning, at each training step of child model, the parent model needs an additional forward pass to generate prediction distribution for guidance.

Acknowledgments

This work was supported in part by the National Natural Science Foundation of China (Grant No. 62206076), the Science and Technology Development Fund, Macau SAR (Grant No. 0101/2019/A2), Shenzhen College Stability Support Plan (Grant No. GXWD20220811173340003 and GXWD20220817123150002) and the Multi-year Research Grant from the University of Macau (Grant No. MYRG2020-00054-FST). We would like to thank the anonymous reviewers and meta-reviewer for their insightful suggestions.

References

Alham Fikri Aji, Nikolay Bogoychev, Kenneth Heafield, and Rico Sennrich. 2020. [In neural machine translation, what does transfer learning transfer?](#) In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7701–7710, Online. Association for Computational Linguistics.

Christos Baziotis, Barry Haddow, and Alexandra Birch. 2020. [Language model prior for low-resource neural](#)

[machine translation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7622–7634, Online. Association for Computational Linguistics.

Jiaao Chen, Dinghan Shen, Weizhu Chen, and Diyi Yang. 2021. [HiddenCut: Simple data augmentation for natural language understanding with better generalizability](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4380–4390, Online. Association for Computational Linguistics.

Yun Chen, Yang Liu, Yong Cheng, and Victor O.K. Li. 2017. [A teacher-student framework for zero-resource neural machine translation](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1925–1935, Vancouver, Canada. Association for Computational Linguistics.

Kevin Clark, Minh-Thang Luong, Christopher D. Manning, and Quoc Le. 2018. [Semi-supervised sequence modeling with cross-view training](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1914–1925, Brussels, Belgium. Association for Computational Linguistics.

Markus Dreyer and Daniel Marcu. 2012. [HyTER: Meaning-equivalent semantics for translation evaluation](#). In *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 162–171, Montréal, Canada. Association for Computational Linguistics.

Sergey Edunov, Myle Ott, Michael Auli, and David Grangier. 2018. [Understanding back-translation at scale](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 489–500, Brussels, Belgium. Association for Computational Linguistics.

Pengzhi Gao, Zhongjun He, Hua Wu, and Haifeng Wang. 2022. [Bi-SimCut: A simple strategy for boosting neural machine translation](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3938–3948, Seattle, United States. Association for Computational Linguistics.

Mozhdeh Gheini and Jonathan May. 2019. [A universal parent model for low-resource neural machine translation transfer](#). *ArXiv preprint*, abs/1909.06516.

Jiatao Gu, Hany Hassan, Jacob Devlin, and Victor O.K. Li. 2018a. [Universal neural machine translation for extremely low resource languages](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics:*

- Human Language Technologies, Volume 1 (Long Papers)*, pages 344–354, New Orleans, Louisiana. Association for Computational Linguistics.
- Jiatao Gu, Yong Wang, Yun Chen, Victor O. K. Li, and Kyunghyun Cho. 2018b. [Meta-learning for low-resource neural machine translation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3622–3631, Brussels, Belgium. Association for Computational Linguistics.
- Dengji Guo, Zhengrui Ma, Min Zhang, and Yang Feng. 2022. [Prediction difference regularization against perturbation for neural machine translation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, *ACL 2022, Dublin, Ireland, May 22-27, 2022*, pages 7665–7675. Association for Computational Linguistics.
- Hakan Inan, Khashayar Khosravi, and Richard Socher. 2017. [Tying word vectors and word classifiers: A loss framework for language modeling](#). In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.
- Daphne Ippolito, Reno Kriz, João Sedoc, Maria Kustikova, and Chris Callison-Burch. 2019. [Comparison of diverse decoding methods from conditional language models](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3752–3762, Florence, Italy. Association for Computational Linguistics.
- Nishant Kambhatla, Logan Born, and Anoop Sarkar. 2022. [Cipherdaug: Ciphertext based data augmentation for neural machine translation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, *ACL 2022, Dublin, Ireland, May 22-27, 2022*, pages 201–218. Association for Computational Linguistics.
- Huda Khayrallah, Brian Thompson, Matt Post, and Philipp Koehn. 2020. [Simulated multiple reference training improves low-resource machine translation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 82–89, Online. Association for Computational Linguistics.
- Yunsu Kim, Yingbo Gao, and Hermann Ney. 2019. [Effective cross-lingual transfer of neural machine translation models without shared vocabularies](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1246–1257, Florence, Italy. Association for Computational Linguistics.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Tom Kocmi and Ondřej Bojar. 2018. [Trivial transfer learning for low-resource neural machine translation](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 244–252, Brussels, Belgium. Association for Computational Linguistics.
- Philipp Koehn. 2004. [Statistical significance tests for machine translation evaluation](#). In *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing*, pages 388–395, Barcelona, Spain. Association for Computational Linguistics.
- Philipp Koehn, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran, Richard Zens, Chris Dyer, Ondřej Bojar, Alexandra Constantin, and Evan Herbst. 2007. [Moses: Open source toolkit for statistical machine translation](#). In *Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics Companion Volume Proceedings of the Demo and Poster Sessions*, pages 177–180, Prague, Czech Republic. Association for Computational Linguistics.
- Philipp Koehn and Rebecca Knowles. 2017. [Six challenges for neural machine translation](#). In *Proceedings of the First Workshop on Neural Machine Translation*, pages 28–39, Vancouver. Association for Computational Linguistics.
- Solomon Kullback and Richard A Leibler. 1951. [On information and sufficiency](#). *The annals of mathematical statistics*, 22(1):79–86.
- Xiaobo Liang, Lijun Wu, Juntao Li, Yue Wang, Qi Meng, Tao Qin, Wei Chen, Min Zhang, and Tiejun Liu. 2021. [R-drop: Regularized dropout for neural networks](#). In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 10890–10905.
- Jianhua Lin. 1991. [Divergence measures based on the shannon entropy](#). *IEEE Trans. Inf. Theory*, 37(1):145–151.
- Xin Liu, Baosong Yang, Dayiheng Liu, Haibo Zhang, Weihua Luo, Min Zhang, Haiying Zhang, and Jinsong Su. 2021a. [Bridging subword gaps in pretrain-finetune paradigm for natural language generation](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 6001–6011, Online. Association for Computational Linguistics.
- Xuebo Liu, Houtim Lai, Derek F. Wong, and Lidia S. Chao. 2020. [Norm-based curriculum learning for neural machine translation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 427–436, Online. Association for Computational Linguistics.

- Xuebo Liu, Longyue Wang, Derek F. Wong, Liang Ding, Lidia S. Chao, Shuming Shi, and Zhaopeng Tu. 2021b. [On the complementarity between pre-training and back-translation for neural machine translation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 2900–2907, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Xuebo Liu, Longyue Wang, Derek F. Wong, Liang Ding, Lidia S. Chao, Shuming Shi, and Zhaopeng Tu. 2021c. [On the copying behaviors of pre-training for neural machine translation](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 4265–4275, Online. Association for Computational Linguistics.
- Xuebo Liu, Derek F. Wong, Lidia S. Chao, and Yang Liu. 2019a. [Latent attribute based hierarchical decoder for neural machine translation](#). *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 27(12):2103–2112.
- Xuebo Liu, Derek F. Wong, Yang Liu, Lidia S. Chao, Tong Xiao, and Jingbo Zhu. 2019b. [Shared-private bilingual word embeddings for neural machine translation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3613–3622, Florence, Italy. Association for Computational Linguistics.
- David Lowell, Brian Howard, Zachary C. Lipton, and Byron Wallace. 2021a. [Unsupervised data augmentation with naive augmentation and without unlabeled data](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 4992–5001, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- David Lowell, Brian Howard, Zachary C. Lipton, and Byron Wallace. 2021b. [Unsupervised data augmentation with naive augmentation and without unlabeled data](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 4992–5001, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Takeru Miyato, Shin-ichi Maeda, Masanori Koyama, and Shin Ishii. 2019. [Virtual adversarial training: A regularization method for supervised and semi-supervised learning](#). *IEEE Trans. Pattern Anal. Mach. Intell.*, 41(8):1979–1993.
- Graham Neubig and Junjie Hu. 2018. [Rapid adaptation of neural machine translation to new languages](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 875–880, Brussels, Belgium. Association for Computational Linguistics.
- Myle Ott, Michael Auli, David Grangier, and Marc’Aurelio Ranzato. 2018. [Analyzing uncertainty in neural machine translation](#). In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pages 3953–3962. PMLR.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. [fairseq: A fast, extensible toolkit for sequence modeling](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 48–53, Minneapolis, Minnesota. Association for Computational Linguistics.
- Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Brussels, Belgium. Association for Computational Linguistics.
- Ofir Press and Lior Wolf. 2017. [Using the output embedding to improve language models](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 157–163, Valencia, Spain. Association for Computational Linguistics.
- Shoetsu Sato, Jin Sakuma, Naoki Yoshinaga, Masashi Toyoda, and Masaru Kitsuregawa. 2020. [Vocabulary adaptation for distant domain adaptation in neural machine translation](#). *ArXiv preprint*, abs/2004.14821.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016a. [Improving neural machine translation models with monolingual data](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 86–96, Berlin, Germany. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016b. [Neural machine translation of rare words with subword units](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725, Berlin, Germany. Association for Computational Linguistics.
- Rico Sennrich and Biao Zhang. 2019. [Revisiting low-resource neural machine translation: A case study](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 211–221, Florence, Italy. Association for Computational Linguistics.
- Dinghan Shen, Mingzhi Zheng, Yelong Shen, Yanru Qu, and Weizhu Chen. 2020. [A simple but tough-to-beat data augmentation approach for natural language understanding and generation](#). *ArXiv preprint*, abs/2009.13818.
- Matthew Snover, Bonnie Dorr, Rich Schwartz, Linnea Micciulla, and John Makhoul. 2006. [A study of translation edit rate with targeted human annotation](#). In *Proceedings of the 7th Conference of the Association*

- for *Machine Translation in the Americas: Technical Papers*, pages 223–231, Cambridge, Massachusetts, USA. Association for Machine Translation in the Americas.
- Xu Tan, Yichong Leng, Jiale Chen, Yi Ren, Tao Qin, and Tie-Yan Liu. 2019. [A study of multilingual neural machine translation](#). *ArXiv preprint*, abs/1912.11625.
- Jörg Tiedemann. 2012. [Parallel data, tools and interfaces in OPUS](#). In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC'12)*, pages 2214–2218, Istanbul, Turkey. European Language Resources Association (ELRA).
- 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 *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.
- Shuo Wang, Zhaopeng Tu, Shuming Shi, and Yang Liu. 2020. [On the inference calibration of neural machine translation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3070–3079, Online. Association for Computational Linguistics.
- Zhijun Wang, Xuebo Liu, and Min Zhang. 2022. [Breaking the representation bottleneck of chinese characters: Neural machine translation with stroke sequence modeling](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Qizhe Xie, Zihang Dai, Eduard H. Hovy, Thang Luong, and Quoc Le. 2020. [Unsupervised data augmentation for consistency training](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Shufang Xie, Ang Lv, Yingce Xia, Lijun Wu, Tao Qin, Tie-Yan Liu, and Rui Yan. 2021. [Target-side input augmentation for sequence to sequence generation](#). In *International Conference on Learning Representations*.
- Minhan Xu and Yu Hong. 2022. [Sub-word alignment is still useful: A vest-pocket method for enhancing low-resource machine translation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers), ACL 2022, Dublin, Ireland, May 22-27, 2022*, pages 613–619. Association for Computational Linguistics.
- Runzhe Zhan, Xuebo Liu, Derek F. Wong, and Lidia S. Chao. 2021. [Meta-curriculum learning for domain adaptation in neural machine translation](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35(16), pages 14310–14318.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. [Bertscore: Evaluating text generation with BERT](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Barret Zoph, Deniz Yuret, Jonathan May, and Kevin Knight. 2016. [Transfer learning for low-resource neural machine translation](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1568–1575, Austin, Texas. Association for Computational Linguistics.