

Informative Language Representation Learning for Massively Multilingual Neural Machine Translation

Renren Jin and Deyi Xiong*

College of Intelligence and Computing, Tianjin University, Tianjin, China

{rrjin, dyxiong}@tju.edu.cn

Abstract

In a multilingual neural machine translation model that fully shares parameters across all languages, an artificial language token is usually used to guide translation into the desired target language. However, recent studies show that prepending language tokens sometimes fails to navigate the multilingual neural machine translation models into right translation directions, especially on zero-shot translation. To mitigate this issue, we propose two methods, language embedding embodiment and language-aware multi-head attention, to learn informative language representations to channel translation into right directions. The former embodies language embeddings into different critical switching points along the information flow from the source to the target, aiming at amplifying translation direction guiding signals. The latter exploits a matrix, instead of a vector, to represent a language in the continuous space. The matrix is chunked into multiple heads so as to learn language representations in multiple subspaces. Experiment results on two datasets for massively multilingual neural machine translation demonstrate that language-aware multi-head attention benefits both supervised and zero-shot translation and significantly alleviates the off-target translation issue. Further linguistic typology prediction experiments show that matrix-based language representations learned by our methods are capable of capturing rich linguistic typology features.¹

1 Introduction

Multilingual neural machine translation (MNMT) (Johnson et al., 2017; Ha et al., 2016; Aharoni et al., 2019; Arivazhagan et al., 2019b), unlike bilingual machine translation with task-specific engineering (language-specific features), uses a single learning system for multiple language pairs, which is jointly

trained in a mult-task learning formalism (Collobert et al., 2011). Parameter sharing across different languages in MNMT models enables transferring of intermediate representations and knowledge among languages, which makes it beneficial to machine translation of low-resource and even zero-resource languages (Firat et al., 2016b; Gu et al., 2018; Neubig and Hu, 2018; Gu et al., 2019; Zhang et al., 2020).

According to the degree of parameter sharing across languages, multilingual neural machine translation (MNMT) approaches can be categorized into two strands: full parameter sharing and partial parameter sharing (Sachan and Neubig, 2018). The former uses a unified model where both the encoder and decoder are shared for all languages. To guide the translation direction, a special token is usually prepended to the beginning of the source or target sentence to indicate the target language (Johnson et al., 2017; Fan et al., 2021). Instead of sharing all model parameters for all translation directions, the latter uses language-specific components, e.g., separate encoders (in many-to-one translation), separate decoders (in one-to-many translation), or separate cross-attention networks (Zoph and Knight, 2016; Firat et al., 2016a; Blackwood et al., 2018; Vázquez et al., 2020; Kong et al., 2021).

The key for full sharing models to distinguish target languages lies in the prepended tokens.² Prior studies find that the embeddings of the prepended tokens encode typological properties of languages (Östling and Tiedemann, 2017; Malaviya et al., 2017; Bjerva et al., 2019; Oncevay et al., 2020). Hence the token embeddings are also referred to as language representations or language embeddings, guiding translation into different target languages whose typological features varies a lot. While the

*Corresponding author.

¹The source code is publicly available at <https://github.com/cordercorder/nmt-multi>.

²When there are multiple languages on the target side, full sharing models cannot even converge during training if no target language information is provided by the prepended tokens (Wang et al., 2019).

prepended tokens play a significant role in language-specific knowledge learning, they are not usually working appropriately, making MNMT models generate translations in wrong languages, especially in zero-shot translation. Such off-target translation issue (Zhang et al., 2020; Yang et al., 2021) implies that prepending special tokens to sentence pairs in full-sharing models is not adequate to learn sufficient language-specific information to guide the MNMT models to translate into right directions.

Partial sharing may be suitable for learning language-specific properties. Nevertheless, it suffers from a rapid growth in the number of parameters with the increase in the number of languages. Additionally, which modules should be language-specific still remains to be further explored.

In this work, we attempt to learn more informative language representations (beyond language embeddings) to improve translation quality of the MNMT models with the full parameter sharing strategy. We argue that the prepended token cannot provide sufficient target language information to guide the translation direction, which is detrimental to both supervised and zero-shot translation. We conjecture that the direction control supervision provided by the prepended token is becoming weaker as the translation information flows to deeper layers. Therefore, instead of prepending special tokens to sentences, we propose a **L**anguage **E**mbedding **E**mbodiment (LEE) strategy that embodies language embeddings at critical switching points (e.g., in-between self-attention and FFN, or self-attention and cross-attention) across layers in both the encoder and decoder along the information flow from the source to the target, so as to amplify the translation direction guiding signal, shown in Figure 1. Experiment results show a boosted translation performance compared with the standard full sharing MNMT model (Johnson et al., 2017).

Inspired by the performance improvement, we further propose **L**anguage-**A**ware Multi-Head **A**ttention (LAA) to model typological features of languages. LAA estimates a matrix, instead of a vector, as the language representation for each language, which is able to encode more language information than a single fixed-length vector that is commonly adopted in previous studies. Motivated by multi-head attention (Vaswani et al., 2017), we split the matrix along the column to learn information from different representation subspaces. In order to probe the typological properties encoded in

language representations, we extract the language representations from the trained MNMT model and use them for linguistic typology prediction.

The main contributions of our work can be summarized as follows:

- We empirically show that both the supervised and zero-shot translation of the standard full-sharing Transformer model are sensitive to the ways of indicating the desired target language.
- We propose the language embedding embodiment and language-aware multi-head attention to learn informative language representations for MNMT. We verify our proposal on two public datasets for massively MNMT in the many-to-many setting. Experimental results indicate that both the supervised and zero-shot translation can significantly benefit from LAA.
- We show that the language representations learned by LAA are generalizable and informative, which suggests that a proper language modeling strategy is crucial for MNMT.

2 Related Work

Multilingual Neural Machine Translation Pioneering studies on multilingual neural machine translation mostly favor to extend the standard bilingual model to MNMT by designing language specific components (Dong et al., 2015; Luong et al., 2016; Zoph and Knight, 2016; Firat et al., 2016a), where the number of parameters grows rapidly with the number of languages. Alternatively, Johnson et al. (2017); Ha et al. (2016) propose to prepend an artificial target language token to source sentences without modifying the model architecture, which is parameter efficient and eases model design. It hence becomes the dominant approach to massively MNMT due to its simplicity and effectiveness (Aharoni et al., 2019; Arivazhagan et al., 2019b; Freitag and Firat, 2020; Rios et al., 2020; Wu et al., 2021). Despite that, it usually lags behind the bilingual counterpart on high-resource language pairs and suffers from off-target translation issue (Zhang et al., 2020) on zero-shot translation. To alleviate these issues, subsequent studies propose approaches such as increasing the model capacity by deepening the model (Zhang et al., 2020), increasing the model cardinality (Xu et al., 2021), inserting mixture-of-experts (MoE) layers (Fedus

et al., 2021), designing lightweight language specific modules (Wang et al., 2018, 2019; Blackwood et al., 2018; Bapna and Firat, 2019; Philip et al., 2020; Zhang et al., 2021; Zhu et al., 2021), reducing negative interference by clustering similar languages (Tan et al., 2019), resolving the gradient conflicts (Wang et al., 2021), dividing the model parameters into a shared part and language-specific part (Lin et al., 2021; Xie et al., 2021; Gong et al., 2021; Wang and Zhang, 2022), bridging the representation gap by data argumentation (Lin et al., 2020) and contrastive learning (Pan et al., 2021), introducing language-agnostic regularization (Ariavzhagan et al., 2019a; Pham et al., 2019) and explicit word alignment supervision (Raganato et al., 2021), to name a few. However, these studies neglect the capacity bottleneck in language representations as they all resort to a language embedding with the same dimension as word embeddings to encode the information of languages whose typological features diverse a lot. In contrast, LAA adopts the matrix as the language representation to alleviate the bottleneck. Moreover, the number of additional parameters brought by LAA is manageable thus it can adapt to massive MNMT settings.

Linguistic Typology Prediction Linguistic typology mainly studies the classification of languages based on their structural properties. Our work is closely related to inferring typological features with language representations from the trained MNMT models. Previous studies demonstrate that language representations in multilingual neural models can capture cross-lingual similarities between languages (Östling and Tiedemann, 2017; Malaviya et al., 2017; Bjerva and Augenstein, 2018; Bjerva et al., 2019; Yu et al., 2021), which can potentially recover missing features in typological databases. Oncevay et al. (2020) fuse language representations in trained MNMT models with features from typological databases to enhance knowledge transfer in MNMT models. They perform typological feature prediction to analyze typological knowledge in the learned language representations. Similarly, we conduct typological feature prediction to evaluate the quality of the language representations learned in LEE and LAA.

3 Language Embedding Embodiment

The central idea for the language embedding embodiment strategy is shortening distance between the language embedding and target translation in



Figure 1: Illustration of the language embedding embodiment strategy. Six switching points in the standard Transformer are particularly marked to insert the language embedding.

order to enable language-embodied translation generation. In Transformer-based MNMT models, instead of only feeding the language embedding into the first layer (Conneau and Lample, 2019), we embody the language embedding into many different layers along the translation information flow from the source to the target, as illustrated in Figure 1. Since the prepended artificial token has been already in vocabulary, feeding the language embedding into other layers will not result in extra parameters. As shown in Figure 1, six candidate positions in the standard Transformer can be used to embody the language embedding. Consider that the language embedding is to be embodied in position 5. Let $\mathbf{y}^{lang} = (\mathbf{y}_1^{lang}, \mathbf{y}_2^{lang}, \dots, \mathbf{y}_n^{lang})$ denote the input of the target language $lang$ into position 5. \mathbf{W}^Q , \mathbf{W}^K and \mathbf{W}^V are parameter matrices of self-attention. \mathbf{E}^{lang} is the language embedding of language $lang$. We compute the output $\mathbf{z}^{lang} = (\mathbf{z}_1^{lang}, \mathbf{z}_2^{lang}, \dots, \mathbf{z}_n^{lang})$ after the language embedding is embodied as follows:

$$e_{ij} = \frac{((\mathbf{y}_i^{lang} + \mathbf{E}^{lang})\mathbf{W}^Q)((\mathbf{y}_j^{lang} + \mathbf{E}^{lang})\mathbf{W}^K)^T}{\sqrt{d}} \quad (1)$$

$$\mathbf{z}_i^{lang} = \sum_{j=1}^n \frac{\exp(e_{ij})}{\sum_{k=1}^n \exp(e_{ik})} (\mathbf{y}_j^{lang} + \mathbf{E}^{lang})\mathbf{W}^V \quad (2)$$

The language embedding embodiment for other positions can be done in a similar way: the language

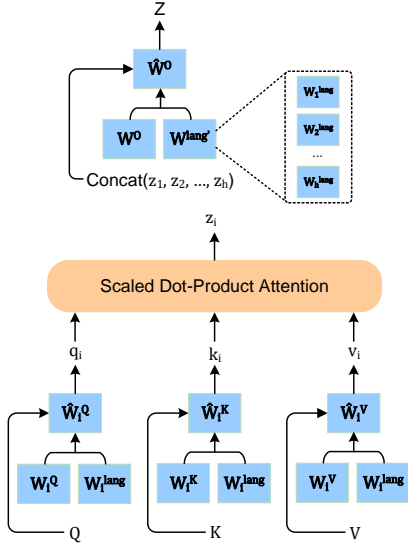


Figure 2: The architecture of the language-aware multi-head attention. W_i^{lang} is the trainable language-specific matrix of the i -th head.

embedding is added to either word embeddings or hidden states fed into the corresponding position so as to make them language-specific. The embodied language embedding is helpful for the MNMT model to distinguish words/subwords that occur in different languages with different meanings.

4 Language-Aware Multi-Head Attention

LEE adopts a fixed-length language embedding to guide the MNMT model to generate translations of the target language. However, the fixed-length language vector may not be sufficient to capture all key linguistic features and diversities in languages, which are important for translation, due to its limited representational capacity. Experiment results of LEE on both zero-shot translation and linguistic typology prediction in Section 5 have also empirically verified the capacity bottleneck in the fixed-length language embedding. To alleviate the capacity bottleneck, we further propose the language aware multi-head attention (shown in Figure 2), representing a language with a matrix, rather than a vector. We incorporate the language-specific matrix into the attention modules (self-attention in the encoder/decoder or cross-attention) since they are the essential components in Transformer. Following the strategy of multi-head attention (Vaswani et al., 2017), we split the matrix along its column to learn information from different representation subspaces of languages.

In the standard multi-head attention, three parameter matrices W_i^Q , W_i^K , W_i^V are used in the i -th attention head to project Q (queries), K (keys) and V (values). After projection, the scaled dot-product attention is applied on each head. The outputs from all heads are then concatenated and multiplied with W^O . $W_i^Q \in \mathbb{R}^{d_{model} \times d_k}$, $W_i^K \in \mathbb{R}^{d_{model} \times d_k}$, $W_i^V \in \mathbb{R}^{d_{model} \times d_v}$, $W^O \in \mathbb{R}^{hd_v \times d_{model}}$ and h is the number of heads. For simplicity, we set $d_k = d_v = d$ and $d_{model} = hd$. In practice, the parameter matrices of all heads from Q , K , V are concatenated into W^Q , W^K , W^V respectively and the projected results are split into h heads, thus benefiting parallel computation and simplifying implementation.

In LAA, we use the matrix for language representation to relax the capacity constraint. Let matrix W^{lang} be the representation of language $lang$, where $W^{lang} \in \mathbb{R}^{d_{model} \times d_{model}}$. W^{lang} is split into h heads along the column, which produces $W_i^{lang} \in \mathbb{R}^{d_{model} \times d}$ for the i -th head. We add W_i^{lang} to W_i^Q , W_i^K , W_i^V to inject the target language information into MNMT models, which can be formulated as follows:

$$\begin{aligned} q_i &= Q(W_i^Q + W_i^{lang}) \\ k_i &= K(W_i^K + W_i^{lang}) \\ v_i &= V(W_i^V + W_i^{lang}) \end{aligned} \quad (3)$$

Let z_i be the scaled dot-product attention output of the i -th head. z_i is computed as:

$$z_i = \text{softmax}\left(\frac{q_i k_i^T}{\sqrt{d}}\right) v_i \quad (4)$$

We split W^O into h heads along the row and pack the outputs from different heads together as follows:

$$Z = \sum_{i=1}^h z_i (W_i^O + (W_i^{lang})^T) \quad (5)$$

where $W_i^O \in \mathbb{R}^{d \times d_{model}}$. As Eq. (3) and (5) show, the language representation matrix can be incorporated into the self-attention layer of the encoder/decoder or the cross-attention layer. It is optimized with other parameters to capture the typological properties of a given language. The MNMT models with the matrix-based W^{lang} learned in this way, have sufficient representation capacity to model both language-specific features (so as to alleviate the negative interference (Wang et al., 2021) between dissimilar languages) and universal features shared by all languages. Additionally, LAA

Dataset	Type	#Languages	#Supervised directions	#Zero-shot directions	#Training sentences
TED-59	English-centric	59	116	3306	5M
OPUS-100	English-centric	100	198	30	55M

Table 1: Statistics of the two datasets.

is able to scale to massively MNMT as there is only one additional matrix for each language and it introduces a very small extra computational cost in training and inference. We show several tips to implement LAA efficiently in Appendix A.

5 Experiments

We conducted two types of experiments, many-to-many translation and linguistic typology prediction, aiming at: (1) systematically evaluating the traditional language token prepending method and (2) examining the effectiveness of the proposed language embedding embodiment and language-aware multi-head attention.

5.1 Experiment Settings for Many-to-Many Translation

Dataset We used two publicly available MNMT datasets TED-59 (Qi et al., 2018) and OPUS-100 (Zhang et al., 2020). Table 1 shows the statistics of the two datasets. Due to the multi-way parallel nature of the TED-59 dataset,³ we were able to construct 3306 (58×57) test sets for zero-shot translation directions by pairing sentences of any two languages via aligned English sentences in the original test sets. To balance training data for various language pairs, we employed the temperature-based sampling strategy (Arivazhagan et al., 2019b) with $T = 5$. More details are in Appendix B.

Model & Training We built our MNMT models on the Transformer base. More detailed settings are shown in Appendix C.

Evaluation We adopted case-sensitive de-tokenized sacreBLEU⁴ (Post, 2018) as the evaluation metric. BLEU scores were averaged over test sets. Following Zhang et al. (2020), we

³Each English sentence may have multiple corresponding translations in different languages.

⁴For TED-59, the signature is: BLEU+case.mixed+numrefs.1+smooth.exp+tok.{13a,zh,ja-mecab-0.996}+version.1.5-.1, tok.zh and tok.ja-mecab-0.996 are only for Chinese and Japanese respectively. For OPUS-100, We adopted the same signature as previous work to be comparable with them: BLEU+case.mixed+numrefs.1+smooth.exp+tok.13a+version.1.4.1

used LangAcc as a complementary evaluation metric for zero-shot translation, which calculates the proportion of the translations in the desired target language. More details about the evaluation settings are in Appendix D.

5.2 Experiment Settings for Linguistic Typology Prediction

Dataset We employed typological features from URIEL typological database (Littell et al., 2017)⁵ for experiments. We mainly used syntax, phonology and phonetic inventory typological features in our work. More details are in Appendix E.

Prediction Methods We adopted k -nearest neighbors approach (k-NN) for linguistic typology prediction. More details about the prediction methods are in Appendix F.

5.3 Many-to-Many Translation Results

We conducted a series of experiments to evaluate both LEE and LAA in many-to-many translation.

5.3.1 Language Embedding Embodiment

We mainly carried out experiments on the TED-59 dataset to examine the effectiveness of LEE. Table 2 shows the experimental results. We further evaluated LEE_{4,5} (the best system in terms of average BLEU on all supervised translation directions on the TED-59 dataset) on the OPUS-100 dataset. As shown in Table 2, poor LangAccs on zero-shot translation directions suggest that the majority of models suffer from the severe off-target translation issue (Zhang et al., 2020). Besides, LEE_{4,5} can slightly boost the average BLEU on supervised translation directions on both datasets. Findings from Table 2 can be summarized as follows:

- The ways to indicate the desired target language to MNMT models can significantly affect the translation performance, especially on zero-shot translation directions.
- The performance gap between $\text{Token}_{\text{src}}$ and $\text{Token}_{\text{tgt}}$ on zero-shot translation directions

⁵https://www.cs.cmu.edu/~dmortens/projects/07_project

ID	Dataset	Model	#Param	En \rightarrow XX		XX \rightarrow En		All		Zero-shot		
				BLEU	WR	BLEU	WR	BLEU	WR	BLEU	LangAcc	WR
①	TED-59	Token _{tgt}	77M	20.74	<i>ref</i>	24.08	<i>ref</i>	22.41	<i>ref</i>	2.42	37.62	<i>ref</i>
②	TED-59	Token _{src}	77M	21.24	91.38	23.77	15.52	22.50	53.45	10.50	71.82	98.91
③	TED-59	LEE ₂	77M	21.26	89.66	23.42	5.17	22.34	47.41	8.87	73.97	97.70
④	TED-59	LEE _{1,2}	77M	20.92	65.52	23.75	32.76	22.33	49.14	8.99	73.97	97.58
⑤	TED-59	LEE ₅	77M	20.70	48.28	24.32	77.59	22.51	62.93	4.10	58.63	83.42
⑥	TED-59	LEE _{4,5}	77M	20.83	65.52	24.49	84.48	22.66	75.00	3.94	59.60	75.05
⑦	TED-59	LEE _{3,4,5}	77M	20.65	39.66	24.56	93.10	22.61	66.38	4.98	61.76	87.69
⑧	TED-59	LEE _{4,5,6}	77M	20.71	51.72	24.39	89.66	22.55	70.69	6.17	64.83	94.65
⑨	TED-59	LEE _{2,5}	77M	21.26	86.21	23.79	27.59	22.53	56.90	9.82	74.44	98.55
⑩	OPUS-100	Token _{tgt}	77M	23.46	<i>ref</i>	29.49	<i>ref</i>	26.47	<i>ref</i>	5.81	55.92	<i>ref</i>
⑪	OPUS-100	Token _{src}	77M	24.04	81.91	28.74	8.51	26.39	45.21	4.06	34.53	23.33
⑫	OPUS-100	LEE _{4,5}	77M	23.38	50.00	29.64	47.87	26.51	48.94	7.37	65.84	56.67

Table 2: Experimental results of LEE on the two datasets. En \rightarrow XX and XX \rightarrow En: translation directions from and to English respectively. All: all supervised translation directions. Zero-shot: zero-shot translation directions. #Param: the number of trainable parameters in the model. WR (Win Rate): the proportion of translation directions that outperform the *ref* system in terms of BLEU. LangAcc: the proportion of translations in the correct target language among all translations. Token_{tgt}: prepending a target language tag to the target side. Token_{src}: prepending a target language tag to the source side. The subscript of LEE denotes the positions to embody the language embedding as shown in Figure 1. It’s worth noting that for supervised translation directions, the test sets in OPUS-100 only cover 94 language pairs and BLEU for "All" is hence averaged on 188 translation directions.

varies substantially across different datasets. We perform analysis to find the main reasons behind the performance discrepancy and the details are shown in Appendix G.

- Prepending a target language token to the source side (Token_{src}) and to the target side (Token_{tgt}) benefit En \rightarrow XX and XX \rightarrow En translation directions respectively.
- Embodying the target language embedding in the decoder is preferable for supervised translation directions.

Please refer to Appendix G for an in-depth analysis that supports these findings.

5.3.2 Language-Aware Multi-Head Attention

Baselines Since LAA enlarges the language representation with a language-specific matrix, introducing additional parameters than LEE, we compared it against two state-of-the-art MNMT baselines. The first is the monolingual adapter (Philip et al., 2020). We omit the bottleneck dimension and activation function in the original adapter and adopt a matrix as the adapter. The adapter is the same size as the language representation in LAA. We share the monolingual adapter across all layers to keep the number of parameters identical to LAA and jointly train it with the other components of the MNMT model. In addition to the monolingual

adapter, we also compared against the combination of LALN (Language-Aware Layer Normalization) and LALT (Language-Aware Linear Transformation) proposed by Zhang et al. (2020).

Table 3 shows the detailed results. We observe that incorporating LAA into the decoder benefits both supervised and zero-shot translation (① vs. ③⑤)(② vs. ④)(⑫ vs. ⑬⑭), while incorporating LAA into the self-attention layer of the encoder deteriorates translation performance (① vs. ②)(③ vs. ④), which is consistent with the finding in LEE (i.e., the incorporation into the the decoder is better than the encoder). Similarly, incorporating the monolingual adapter into the decoder is more effective than into the encoder (⑧ vs. ⑨). **Moreover, LAA outperforms all strong baselines under the condition of using equal number of parameters (⑧⑨⑩⑪ vs. ③⑤),** which demonstrates that LAA can capture rich language information to improve translation performance of massively MNMT model. To examine whether LAA and LEE are complementary to each other, we combine LAA_{dec.self} and LEE_{4,5} and the results show a slightly boosted translation performance (③ vs. ⑦) (⑬ vs. ⑮). However, adopting LAA while prepending the language tag to the target side leads to a slightly performance drop on supervised translation directions (③ vs. ⑥).

The experiments in Table 2 and 3 are conducted

ID	Dataset	Model	#Param	En → XX		XX → En		All		Zero-shot		
				BLEU	WR	BLEU	WR	BLEU	WR	BLEU	LangAcc	WR
①	TED-59	Token _{tgt}	77M	20.74	<i>ref</i>	24.08	<i>ref</i>	22.41	<i>ref</i>	2.42	37.62	<i>ref</i>
②	TED-59	LAA _{enc,self}	92M	20.11	17.24	21.13	1.72	20.62	9.48	6.30	71.92	92.77
③	TED-59	LAA _{dec,self}	92M	21.29	84.48	25.14	96.55	23.21	90.52	8.94	74.68	98.58
④	TED-59	LAA _{enc,self,dec,self}	92M	21.04	68.97	23.10	1.72	22.07	35.34	8.89	74.03	97.58
⑤	TED-59	LAA _{dec,self,cros}	92M	21.47	89.66	24.79	94.83	23.13	92.24	12.69	74.57	99.15
⑥	TED-59	LAA _{dec,self} + Token _{tgt}	92M	21.19	79.31	25.11	96.55	23.15	87.93	9.07	74.65	98.40
⑦	TED-59	LAA _{dec,self} + LEE _{4,5}	92M	21.28	87.93	25.29	100.00	23.28	93.97	9.48	74.69	99.09
⑧	TED-59	Adapter _{enc}	92M	21.04	68.97	22.91	0.00	21.97	34.48	8.20	74.44	96.79
⑨	TED-59	Adapter _{dec}	92M	21.43	91.38	24.14	62.07	22.79	76.72	10.61	75.36	98.97
⑩	TED-59	Adapter _{enc,dec}	92M	21.10	79.31	22.95	1.72	22.03	40.52	9.17	75.10	98.09
⑪	TED-59	LALN + LALT	94M	21.37	75.86	22.85	1.72	22.11	38.79	7.69	73.18	95.58
⑫	OPUS-100	Token _{tgt}	77M	23.46	<i>ref</i>	29.49	<i>ref</i>	26.47	<i>ref</i>	5.81	55.92	<i>ref</i>
⑬	OPUS-100	LAA _{dec,self}	103M	24.42	90.43	29.78	79.79	27.10	85.11	11.23	81.53	100.00
⑭	OPUS-100	LAA _{dec,self,cros}	103M	24.42	89.36	29.58	68.09	27.00	78.72	12.13	87.65	100.00
⑮	OPUS-100	LAA _{dec,self} + LEE _{4,5}	103M	24.41	88.30	29.83	76.60	27.12	82.45	11.31	81.70	100.00

Table 3: Experiment results of LAA on the two datasets. The results of ① and ⑫ are from Table 2. The subscript of LAA denotes the layers where the language representation is incorporated. The subscript of the adapter denotes the modules where the monolingual adapter is incorporated. enc/dec denote the encoder/decoder. self/cros denote self-attention/cross-attention layer. Considering that both LAA_{dec,self} and LAA_{dec,self} + LEE_{4,5} achieve superior average BLEU scores on supervised translation directions and LAA_{dec,self,cros} obtains the best BLEU on zero-shot translation, we verify them again on the OPUS-100 dataset.

under a temperature-based sampling setting of $T = 5$, which oversamples training data of low-resource language pairs. Oversampling usually leads to improvements in average BLEU on supervised translation directions (Table 3 vs. Table 4). To eliminate such effect and compare with previous work (Zhang et al., 2020; Xu et al., 2021),⁶ we carried out experiments without oversampling. Table 4 summarizes our experimental results without oversampling. The results suggest that LAA yields larger BLEU improvements on supervised translation directions when training on the raw data distribution. e.g., (⑬) - (⑫) = +0.63 BLEU, Table 3) vs. (⑫) - (⑨) = +1 BLEU, Table 4). Additionally, the performance gain on zero-shot translation is also retained mostly ((⑬) - (⑫) = +5.42 BLEU, Table 3) vs. (⑫) - (⑨) = +5.3 BLEU, Table 4). We enlarge the MNMT model by increasing the number of layers in both the encoder and decoder to 24 so that the number of parameters in our models is approximately the same as the largest model used by Zhang et al. (2020) and Xu et al. (2021).⁷ As the model deepens, we adjusted the dropout rate to 0.2. Our methods outperform previous results (Zhang

⁶As the two studies are carried out on the OPUS-100 dataset without oversampling, we remove oversampling to make a fair comparison with them.

⁷Although Xu et al. (2021) don't report the exact number of parameters, the number of parameters in their model is approximately equal to that of Zhang et al. (2020) as stated in paper.

et al., 2020; Xu et al., 2021) on both supervised and zero-shot translation directions ((⑥⑦⑧) vs. (⑭)).⁸

5.4 Linguistic Typology Prediction Results

Similarly, we conducted experiments on the linguistic typology prediction for both LEE and LAA which are trained on the two datasets.

5.4.1 Language Embedding Embodiment

We selected language embeddings learned by Token_{src}, Token_{tgt} and LEE_{4,5} to perform linguistic typology prediction given that Token_{src} and Token_{tgt} are the dominant practice in MNMT and LEE_{4,5} achieves fairly good performance over all supervised translation directions across the two datasets. We present the experimental results in Table 6 (Appendix H). Optimal k s for Token_{src}, Token_{tgt} and LEE_{4,5} are not always consistent and it is difficult to conclude the most superior one. We therefore adopt the maximum accuracy under different settings of k as the evaluation metric. We observe that there is no significant performance

⁸Yang et al. (2021) and Wang et al. (2022) report better results than us. However, Yang et al. (2021) oversampled the training data for low-resource language pairs and removed five language pairs without test data. We removed oversampling and used all training data of the OPUS-100 dataset, hence the model in our work needs to schedule its capacity over the extra ten translation directions. Wang et al. (2022) conducted experiments with the MNMT model of up to 3.8B parameters, which is several times larger than ours. To the best of our knowledge, there are no other results superior to ours except theirs.

ID	Dataset	Model	#Param	En \rightarrow XX		XX \rightarrow En		All		Zero-shot		
				BLEU	WR	BLEU	WR	BLEU	WR	BLEU	LangAcc	WR
①	TED-59	Token _{tgt} ⁻	77M	19.54	<i>ref</i>	24.23	<i>ref</i>	21.89	<i>ref</i>	2.84	40.94	<i>ref</i>
②	TED-59	Token _{src} ⁻	77M	20.25	96.55	23.65	8.62	21.95	52.59	9.65	65.45	96.77
③	TED-59	LEE _{4,5} ⁻	77M	19.56	51.72	24.42	86.21	21.99	68.97	6.43	66.25	87.57
④	TED-59	LAA _{dec,self} ⁻	92M	20.64	96.55	25.16	98.28	22.90	97.41	9.93	75.04	97.67
⑤	TED-59	LAA _{dec,self} ⁻ + LEE _{4,5} ⁻	92M	20.67	94.83	25.05	93.10	22.86	93.97	10.02	75.26	97.82
⑥	OPUS-100	Zhang et al. (2020)	254M	23.96	-	31.36	-	27.66	-	5.24	47.91	-
⑦	OPUS-100	Zhang et al. (2020)	254M	23.36	-	30.98	-	27.17	-	14.08	87.68	-
⑧	OPUS-100	Xu et al. (2021)	-	24.17	-	32.19	-	28.18	-	14.71	-	-
⑨	OPUS-100	Token _{tgt} ⁻	77M	21.82	<i>ref</i>	28.45	<i>ref</i>	25.14	<i>ref</i>	6.63	58.76	<i>ref</i>
⑩	OPUS-100	Token _{src} ⁻	77M	22.15	74.47	27.68	10.64	24.91	42.55	4.91	37.80	30.00
⑪	OPUS-100	LEE _{4,5} ⁻	77M	21.49	45.74	28.48	43.62	24.98	44.68	10.08	79.90	96.67
⑫	OPUS-100	LAA _{dec,self} ⁻	103M	23.57	91.49	28.71	70.21	26.14	80.85	11.93	80.39	100.00
⑬	OPUS-100	LAA _{dec,self} ⁻ + LEE _{4,5} ⁻	103M	23.69	91.49	28.88	78.72	26.29	85.11	12.77	85.00	100.00
⑭	OPUS-100	⑬ + 24 layers	236M	26.82	98.94	32.31	100.00	29.56	99.47	15.08	84.59	100.00

Table 4: Experimental results without oversampling on the two datasets. The superscript "-" denotes that oversampling is removed.

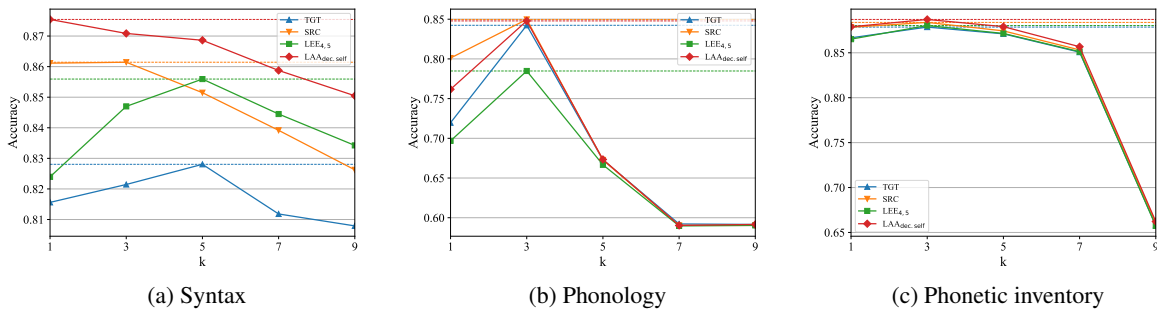


Figure 3: Prediction accuracy on syntax, phonology and phonetic inventory features using the language embeddings learned by Token_{tgt}, Token_{src}, LEE_{4,5} and LAA_{dec,self} which are trained on the TED-59 dataset.

gap except inferring the phonology features with the language embeddings trained on the TED-59 dataset (⑨, Table 6 in Appendix H).

5.4.2 Language Aware Multi-Head Attention

Figure 3 and 5 (the latter shown in Appendix I) show the prediction results of using the language representation learned by LAA for linguistic typology probing. We observe that the language representation learned in LAA achieves the best accuracy on syntax feature inference across the two datasets, which demonstrates its generalization ability to some extent. For phonology and phonetic inventory feature inference, there are no significant differences among the language representations learned in various MNMT models. We hypothesize that typological properties encoded in language representations are relied on the task that learns language representations. The translation task can force MNMT models to learn syntactic information to accommodate syntactic divergences

across languages for better translation. Additionally, the superior performance of LAA on many-to-many translation and linguistic typology prediction suggests that the language representations learned in LAA can benefit multilingual translation.

5.5 Effect of the Increased Parameters Introduced by LAA

We carried out experiments to study the effect of the increased parameters introduced by LAA. Specifically, we removed the language-specific properties of the introduced trainable matrices in LAA by randomly activating them with equal probabilities during training and inference. We denote the MNMT models with the random matrices as LAA^R. As these matrices become random, there is no explicit signal to guide the translation into the desired target languages. Hence we followed the prepending token strategies of Token_{tgt} and Token_{src} to navi-

ID	Dataset	Model	#Param	En → XX		XX → En		All		Zero-shot		
				BLEU	WR	BLEU	WR	BLEU	WR	BLEU	LangAcc	WR
①	TED-59	Token _{tgt}	77M	20.74	<i>ref</i>	24.08	<i>ref</i>	22.41	<i>ref</i>	2.42	37.62	<i>ref</i>
②	TED-59	Token _{src}	77M	21.24	91.38	23.77	15.52	22.50	53.45	10.50	71.82	98.91
③	TED-59	LAA _{dec.self}	92M	21.29	84.48	25.14	96.55	23.21	90.52	8.94	74.68	98.58
④	TED-59	LAA _{dec.self} + LEE _{4,5}	92M	21.28	87.93	25.29	100.00	23.28	93.97	9.48	74.69	99.09
⑤	TED-59	LAA _{dec.self} ^R + Token _{src}	92M	20.74±0.02	44.14±3.78	25.94±0.01	98.28±0.00	23.34±0.01	71.21±1.89	7.08±0.00	55.73±0.01	95.50±0.09
⑥	TED-59	LAA _{dec.self} ^R + Token _{tgt}	92M	20.18±0.01	10.34±1.22	26.25±0.02	98.28±0.00	23.22±0.01	54.31±0.61	0.92±0.00	6.76±0.00	2.75±0.09
⑦	OPUS-100	Token _{tgt}	77M	23.46	<i>ref</i>	29.49	<i>ref</i>	26.47	<i>ref</i>	5.81	55.92	<i>ref</i>
⑧	OPUS-100	Token _{src}	77M	24.04	81.91	28.74	8.51	26.39	45.21	4.06	34.53	23.33
⑨	OPUS-100	LAA _{dec.self}	103M	24.42	90.43	29.78	79.79	27.10	85.11	11.23	81.53	100.00
⑩	OPUS-100	LAA _{dec.self} + LEE _{4,5}	103M	24.41	88.30	29.83	76.60	27.12	82.45	11.31	81.70	100.00
⑪	OPUS-100	LAA _{dec.self} ^R + Token _{src}	103M	23.65±0.02	64.68±3.79	28.61±0.01	12.98±0.89	26.13±0.01	38.83±1.50	4.18±0.01	39.94±0.04	20.00±0.00
⑫	OPUS-100	LAA _{dec.self} ^R + Token _{tgt}	103M	22.80±0.02	16.81±1.39	29.27±0.01	22.13±2.43	26.04±0.01	19.47±1.22	6.16±0.01	68.00±0.11	50.00±2.36

Table 5: Experiment results of LAA_{dec.self}^R on the two datasets. The superscript "R" denotes that the language-specific matrices in LAA_{dec.self} are stochastically activated during training and inference. We adopted Token_{tgt} or Token_{src} with LAA_{dec.self}^R to guide the MNMT models into the right translation directions. The results of ①②⑦⑧ are from Table 2. The results of ③④⑨⑩ are from Table 3.

gate the translation directions.⁹ We incorporated the stochastic matrices into the self-attention modules of the decoder since LAA_{dec.self} achieves superior performance on both supervised and zero-shot translation directions as shown in Table 3 and 4. For efficiency, the best checkpoint was selected according to the average BLEU on the validation sets with only one random seed. We evaluated models on the test sets with the best checkpoint and report the mean and standard deviation of the BLEU score with 5 different random seeds. Results with the temperature-based sampling strategy ($T = 5$) on the two datasets are shown in Table 5. We also conducted experiments on the raw data distribution and results are presented in Table 7 of Appendix J.

We observe that the performance of LAA_{dec.self}^R with different random seeds is stable as the standard deviations on various metrics are small. Additionally, the performance gap between LAA_{dec.self} and LAA_{dec.self}^R is small on supervised translation directions on the TED-59 dataset. Despite that, LAA_{dec.self} outperforms LAA_{dec.self}^R by a large margin on the OPUS-100 dataset. We hypothesize that the TED-59 dataset covers more restrict domains than the OPUS-100 dataset sampled from OPUS collection (Tiedemann, 2012),¹⁰ which may benefit LAA_{dec.self}^R as the knowledge from similar domains can be easily transferred across the stochastic matrices during training. Furthermore,

⁹We also attempted to make these stochastic matrices language-specific by gradually increasing the sampling probability on a specific matrix for each language during training but results are not satisfactory.

¹⁰There are various corpora in the OPUS collection, such as Wikipedia corpus, Bible corpus, UN corpus and TED corpus, etc. As the OPUS-100 dataset is constructed by randomly sampling sentence pairs from the corpora of OPUS collection for each language pair, it may cover more diverse domains.

LAA_{dec.self}^R consistently lags behind LAA_{dec.self} on zero-shot translation directions on the two datasets. **Given the same number of parameters in LAA_{dec.self}^R and LAA_{dec.self}, we can conclude that the performance improvement of LAA_{dec.self} is not only due to the parameter increase.**

6 Conclusion

To improve massively MNMT, we have presented two approaches to learning informative language representations, language embedding embodiment and language-aware multi-head attention. We find that the ways to inject target language information into MNMT models have a significant impact on the translation performance, especially on zero-shot translation. We validate the effectiveness of the two approaches on two public datasets. We probe the typological features encoded in language representations learned by LEE and LAA through linguistic typology feature prediction. The superior prediction performance of the matrix-based language representations learned by LAA on syntax features demonstrates its informativeness.

Acknowledgements

The present research was supported by the Key Research and Development Program of Yunnan Province (Grant No. 202203AA080004-2). We would like to thank the anonymous reviewers for their insightful comments.

References

- Roei Aharoni, Melvin Johnson, and Orhan Firat. 2019. [Massively multilingual neural machine translation](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 3874–3884. Association for Computational Linguistics.
- Naveen Arivazhagan, Ankur Bapna, Orhan Firat, Roei Aharoni, Melvin Johnson, and Wolfgang Macherey. 2019a. [The missing ingredient in zero-shot neural machine translation](#). *CoRR*, abs/1903.07091.
- Naveen Arivazhagan, Ankur Bapna, Orhan Firat, Dmitry Lepikhin, Melvin Johnson, Maxim Krikun, Mia Xu Chen, Yuan Cao, George F. Foster, Colin Cherry, Wolfgang Macherey, Zhifeng Chen, and Yonghui Wu. 2019b. [Massively multilingual neural machine translation in the wild: Findings and challenges](#). *CoRR*, abs/1907.05019.
- Ankur Bapna and Orhan Firat. 2019. [Simple, scalable adaptation for neural machine translation](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pages 1538–1548. Association for Computational Linguistics.
- Johannes Bjerva and Isabelle Augenstein. 2018. [From phonology to syntax: Unsupervised linguistic typology at different levels with language embeddings](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2018, New Orleans, Louisiana, USA, June 1-6, 2018, Volume 1 (Long Papers)*, pages 907–916. Association for Computational Linguistics.
- Johannes Bjerva, Robert Östling, Maria Han Veiga, Jörg Tiedemann, and Isabelle Augenstein. 2019. [What do language representations really represent?](#) *Comput. Linguistics*, 45(2):381–389.
- Graeme W. Blackwood, Miguel Ballesteros, and Todd Ward. 2018. [Multilingual neural machine translation with task-specific attention](#). In *Proceedings of the 27th International Conference on Computational Linguistics, COLING 2018, Santa Fe, New Mexico, USA, August 20-26, 2018*, pages 3112–3122. Association for Computational Linguistics.
- Ronan Collobert, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel P. Kuksa. 2011. [Natural language processing \(almost\) from scratch](#). *J. Mach. Learn. Res.*, 12:2493–2537.
- Alexis Conneau and Guillaume Lample. 2019. [Cross-lingual language model pretraining](#). In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 7057–7067.
- Daxiang Dong, Hua Wu, Wei He, Dianhai Yu, and Haifeng Wang. 2015. [Multi-task learning for multiple language translation](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing of the Asian Federation of Natural Language Processing, ACL 2015, July 26-31, 2015, Beijing, China, Volume 1: Long Papers*, pages 1723–1732. The Association for Computer Linguistics.
- Matthew S. Dryer and Martin Haspelmath, editors. 2013. *WALS Online*. Max Planck Institute for Evolutionary Anthropology, Leipzig.
- Angela Fan, Shruti Bhosale, Holger Schwenk, Zhiyi Ma, Ahmed El-Kishky, Siddharth Goyal, Mandeep Baines, Onur Celebi, Guillaume Wenzek, Vishrav Chaudhary, Naman Goyal, Tom Birch, Vitaliy Liptchinsky, Sergey Edunov, Michael Auli, and Armand Joulin. 2021. [Beyond english-centric multilingual machine translation](#). *J. Mach. Learn. Res.*, 22:107:1–107:48.
- William Fedus, Barret Zoph, and Noam Shazeer. 2021. [Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity](#). *CoRR*, abs/2101.03961.
- Orhan Firat, Kyunghyun Cho, and Yoshua Bengio. 2016a. [Multi-way, multilingual neural machine translation with a shared attention mechanism](#). In *NAACL HLT 2016, The 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, San Diego California, USA, June 12-17, 2016*, pages 866–875. The Association for Computational Linguistics.
- Orhan Firat, Baskaran Sankaran, Yaser Al-Onaizan, Fatos T. Yarman-Vural, and Kyunghyun Cho. 2016b. [Zero-resource translation with multi-lingual neural machine translation](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP 2016, Austin, Texas, USA, November 1-4, 2016*, pages 268–277. The Association for Computational Linguistics.
- Markus Freitag and Orhan Firat. 2020. [Complete multilingual neural machine translation](#). In *Proceedings of the Fifth Conference on Machine Translation, WMT@EMNLP 2020, Online, November 19-20, 2020*, pages 550–560. Association for Computational Linguistics.
- Hongyu Gong, Xian Li, and Dmitriy Genzel. 2021. [Adaptive sparse transformer for multilingual translation](#). *CoRR*, abs/2104.07358.
- Jiatao Gu, Hany Hassan, Jacob Devlin, and Victor O. K. Li. 2018. [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, NAACL-HLT 2018, New Orleans, Louisiana, USA, June 1-6, 2018, Volume 1 (Long Papers)*, pages 344–354. Association for Computational Linguistics.
- Jiatao Gu, Yong Wang, Kyunghyun Cho, and Victor O. K. Li. 2019. [Improved zero-shot neural machine translation via ignoring spurious correlations](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 1258–1268. Association for Computational Linguistics.
- Thanh-Le Ha, Jan Niehues, and Alex Waibel. 2016. [Toward multilingual neural machine translation with universal encoder and decoder](#). In *Proceedings of the 13th International Conference on Spoken Language Translation, IWSLT 2016, Seattle, WA, USA, December 8-9, 2016*. International Workshop on Spoken Language Translation.
- Harald Hammarström, Robert Forkel, Martin Haspelmath, and Sebastian Bank. 2021. *Glottolog 4.4*. Leipzig.
- Melvin Johnson, Mike Schuster, Quoc V. Le, Maxim Krikun, Yonghui Wu, Zhifeng Chen, Nikhil Thorat, Fernanda B. Viégas, Martin Wattenberg, Greg Corrado, Macduff Hughes, and Jeffrey Dean. 2017. [Google’s multilingual neural machine translation system: Enabling zero-shot translation](#). *Trans. Assoc. Comput. Linguistics*, 5:339–351.
- 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*.
- Xiang Kong, Adithya Renduchintala, James Cross, Yuqing Tang, Jiatao Gu, and Xian Li. 2021. [Multilingual neural machine translation with deep encoder and multiple shallow decoders](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume, EACL 2021, Online, April 19 - 23, 2021*, pages 1613–1624. Association for Computational Linguistics.
- Taku Kudo and John Richardson. 2018. [Sentencepiece: A simple and language independent subword tokenizer and detokenizer for neural text processing](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, EMNLP 2018: System Demonstrations, Brussels, Belgium, October 31 - November 4, 2018*, pages 66–71. Association for Computational Linguistics.
- Zehui Lin, Xiao Pan, Mingxuan Wang, Xipeng Qiu, Jiangtao Feng, Hao Zhou, and Lei Li. 2020. [Pre-training multilingual neural machine translation by leveraging alignment information](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 2649–2663. Association for Computational Linguistics.
- Zehui Lin, Liwei Wu, Mingxuan Wang, and Lei Li. 2021. [Learning language specific sub-network for multilingual machine translation](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, pages 293–305. Association for Computational Linguistics.
- Patrick Littell, David R. Mortensen, Ke Lin, Katherine Kairis, Carlisle Turner, and Lori S. Levin. 2017. [URIEL and lang2vec: Representing languages as typological, geographical, and phylogenetic vectors](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics, EACL 2017, Valencia, Spain, April 3-7, 2017, Volume 2: Short Papers*, pages 8–14. Association for Computational Linguistics.
- Minh-Thang Luong, Quoc V. Le, Ilya Sutskever, Oriol Vinyals, and Lukasz Kaiser. 2016. [Multi-task sequence to sequence learning](#). In *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*.
- Chaitanya Malaviya, Graham Neubig, and Patrick Littell. 2017. [Learning language representations for typology prediction](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, EMNLP 2017, Copenhagen, Denmark, September 9-11, 2017*, pages 2529–2535. Association for Computational Linguistics.
- Steven Moran and Daniel McCloy, editors. 2019. *PHOIBLE 2.0*. Max Planck Institute for the Science of Human History, Jena.
- 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, Brussels, Belgium, October 31 - November 4, 2018*, pages 875–880. Association for Computational Linguistics.
- Arturo Oncevay, Barry Haddow, and Alexandra Birch. 2020. [Bridging linguistic typology and multilingual machine translation with multi-view language representations](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 2391–2406. Association for Computational Linguistics.
- Robert Östling and Jörg Tiedemann. 2017. [Continuous multilinguality with language vectors](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics, EACL 2017, Valencia, Spain, April 3-7, 2017, Volume 2: Short Papers*, pages 644–649. Association for Computational Linguistics.

- 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: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Demonstrations*, pages 48–53. Association for Computational Linguistics.
- Xiao Pan, Mingxuan Wang, Liwei Wu, and Lei Li. 2021. [Contrastive learning for many-to-many multilingual neural machine translation](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, pages 244–258. Association for Computational Linguistics.
- Ngoc-Quan Pham, Jan Niehues, Thanh-Le Ha, and Alexander Waibel. 2019. [Improving zero-shot translation with language-independent constraints](#). In *Proceedings of the Fourth Conference on Machine Translation, WMT 2019, Florence, Italy, August 1-2, 2019 - Volume 1: Research Papers*, pages 13–23. Association for Computational Linguistics.
- Jerin Philip, Alexandre Berard, Matthias Gallé, and Laurent Besacier. 2020. [Monolingual adapters for zero-shot neural machine translation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 4465–4470. 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, WMT 2018, Belgium, Brussels, October 31 - November 1, 2018*, pages 186–191. Association for Computational Linguistics.
- Ye Qi, Devendra Singh Sachan, Matthieu Felix, Sarguna Padmanabhan, and Graham Neubig. 2018. [When and why are pre-trained word embeddings useful for neural machine translation?](#) In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT, New Orleans, Louisiana, USA, June 1-6, 2018, Volume 2 (Short Papers)*, pages 529–535. Association for Computational Linguistics.
- Alessandro Raganato, Raúl Vázquez, Mathias Creutz, and Jörg Tiedemann. 2021. [An empirical investigation of word alignment supervision for zero-shot multilingual neural machine translation](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021*, pages 8449–8456. Association for Computational Linguistics.
- Annette Rios, Mathias Müller, and Rico Sennrich. 2020. [Subword segmentation and a single bridge language affect zero-shot neural machine translation](#). In *Proceedings of the Fifth Conference on Machine Translation, WMT@EMNLP 2020, Online, November 19-20, 2020*, pages 528–537. Association for Computational Linguistics.
- Devendra Singh Sachan and Graham Neubig. 2018. [Parameter sharing methods for multilingual self-attentional translation models](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers, WMT 2018, Belgium, Brussels, October 31 - November 1, 2018*, pages 261–271. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Neural machine translation of rare words with subword units](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers*. The Association for Computer Linguistics.
- Xu Tan, Jiale Chen, Di He, Yingce Xia, Tao Qin, and Tie-Yan Liu. 2019. [Multilingual neural machine translation with language clustering](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pages 963–973. Association for Computational Linguistics.
- Jörg Tiedemann. 2012. [Parallel data, tools and interfaces in OPUS](#). In *Proceedings of the Eighth International Conference on Language Resources and Evaluation, LREC 2012, Istanbul, Turkey, May 23-25, 2012*, pages 2214–2218. 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.
- Raúl Vázquez, Alessandro Raganato, Mathias Creutz, and Jörg Tiedemann. 2020. [A systematic study of inner-attention-based sentence representations in multilingual neural machine translation](#). *Comput. Linguistics*, 46(2):387–424.
- Hongyu Wang, Shuming Ma, Li Dong, Shaohan Huang, Dongdong Zhang, and Furu Wei. 2022. [Deepnet: Scaling transformers to 1, 000 layers](#). *CoRR*, abs/2203.00555.
- Qian Wang and Jiajun Zhang. 2022. [Parameter differentiation based multilingual neural machine translation](#). In *Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI 2022, Thirty-Fourth Conference on Innovative Applications of Artificial Intelligence,*

- IAAI 2022, *The Twelfth Symposium on Educational Advances in Artificial Intelligence, EAAI 2022 Virtual Event, February 22 - March 1, 2022*, pages 11440–11448. AAAI Press.
- Yining Wang, Jiajun Zhang, Feifei Zhai, Jingfang Xu, and Chengqing Zong. 2018. [Three strategies to improve one-to-many multilingual translation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 2955–2960. Association for Computational Linguistics.
- Yining Wang, Long Zhou, Jiajun Zhang, Feifei Zhai, Jingfang Xu, and Chengqing Zong. 2019. [A compact and language-sensitive multilingual translation method](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 1213–1223. Association for Computational Linguistics.
- Zirui Wang, Yulia Tsvetkov, Orhan Firat, and Yuan Cao. 2021. [Gradient vaccine: Investigating and improving multi-task optimization in massively multilingual models](#). In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Liwei Wu, Shanbo Cheng, Mingxuan Wang, and Lei Li. 2021. [Language tags matter for zero-shot neural machine translation](#). In *Findings of the Association for Computational Linguistics: ACL/IJCNLP 2021, Online Event, August 1-6, 2021*, volume ACL/IJCNLP 2021 of *Findings of ACL*, pages 3001–3007. Association for Computational Linguistics.
- Wanying Xie, Yang Feng, Shuhao Gu, and Dong Yu. 2021. [Importance-based neuron allocation for multilingual neural machine translation](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, pages 5725–5737. Association for Computational Linguistics.
- Hongfei Xu, Qihui Liu, Josef van Genabith, and Deyi Xiong. 2021. [Modeling task-aware MIMO cardinality for efficient multilingual neural machine translation](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 2: Short Papers), Virtual Event, August 1-6, 2021*, pages 361–367. Association for Computational Linguistics.
- Yilin Yang, Akiko Eriguchi, Alexandre Muzio, Prasad Tadepalli, Stefan Lee, and Hany Hassan. 2021. [Improving multilingual translation by representation and gradient regularization](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021*, pages 7266–7279. Association for Computational Linguistics.
- Dian Yu, Taiqi He, and Kenji Sagae. 2021. [Language embeddings for typology and cross-lingual transfer learning](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, pages 7210–7225. Association for Computational Linguistics.
- Biao Zhang, Ankur Bapna, Rico Sennrich, and Orhan Firat. 2021. [Share or not? learning to schedule language-specific capacity for multilingual translation](#). In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Biao Zhang, Philip Williams, Ivan Titov, and Rico Sennrich. 2020. [Improving massively multilingual neural machine translation and zero-shot translation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 1628–1639. Association for Computational Linguistics.
- Yaoming Zhu, Jiangtao Feng, Chengqi Zhao, Mingxuan Wang, and Lei Li. 2021. [Counter-interference adapter for multilingual machine translation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 16-20 November, 2021*, pages 2812–2823. Association for Computational Linguistics.
- Barret Zoph and Kevin Knight. 2016. [Multi-source neural translation](#). In *NAACL HLT 2016, The 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, San Diego California, USA, June 12-17, 2016*, pages 30–34. The Association for Computational Linguistics.

A Efficient Implementation of LAA

Suppose that the number of target languages is l and the dimensions of Q, K, V are $b \times n \times d_{model}$, where b is the batch size and n is the sequence length. Since a minibatch usually contains samples from various language pairs during training, each minibatch is required to be split into smaller batches according to the target language so as to compute Eq. (3) and (5). And the output of Eq. (5) is required to be reorganized to preserve the original order of samples within the minibatch, which brings extra computational cost and is inefficient for computing on modern parallel hardware like GPU. To avoid this, we maintain a matrix $\widetilde{W} \in \mathbb{R}^{l \times d_{model} \times d_{model}}$ which consists of language representations from all languages. We select $\overline{W} \in \mathbb{R}^{b \times d_{model} \times d_{model}}$ from \widetilde{W} , which contains all language representations for samples in the minibatch. This selection can be efficiently executed by the `PyTorch` toolkit.¹¹ To avoid broadcasting¹² W^Q, W^K, W^V, W^O into larger dimensions when added with \overline{W} , we reformulate Eq. (3) and (5) as follows:

$$\begin{aligned} q &= QW^Q + Q\overline{W} \\ k &= KW^K + K\overline{W} \\ v &= VW^V + V\overline{W} \end{aligned} \quad (6)$$

$$Z = zW^O + z\overline{W}^T \quad (7)$$

Note that we omit the subscript for head index in original formulas as all heads are computed in parallel. As there are always duplicate elements in $\{Q, K, V\}$, the second term (i.e. $\{Q, K, V\}\overline{W}$) in Eq. (6) can be computed first and then cached to avoid redundant computation.

B Dataset for Many-to-Many Translation

We removed the `__en__` tag prepended to non-English sentences in the original TED-59 dataset.¹³

¹¹https://pytorch.org/docs/stable/generated/torch.index_select.html?highlight=index_select#torch.index_select

¹²<https://pytorch.org/docs/stable/notes/broadcasting.html?highlight=broadcasting>

¹³An artificial English tag `__en__` is prepended to every non-English sentence in the raw dataset, which may affect model training and bias BLEU. Nevertheless, previous works on this dataset usually do not elaborate this procedure, which may make our results on this dataset not directly comparable to theirs.

We trained BPE model (Sennrich et al., 2016) using SentencePiece (Kudo and Richardson, 2018) to get subword units with a joint vocabulary of size 64K.

C Model & Training for Many-to-Many Translation

We implement our MNMT models based on Fairseq (Ott et al., 2019). We set the dimension of word embeddings and FFN layer to 512/2048. Embeddings were shared for the encoder, decoder and the output projection. To prevent overfitting, we set dropout rate to 0.1 on the OPUS-100 dataset and 0.2 on the TED-59 dataset due to its relatively small data size. We adopted the cross-entropy loss with a label smoothing of 0.1 as the training objective. We used Adam ($\beta_1 = 0.9, \beta_2 = 0.98$) (Kingma and Ba, 2015) to optimize model parameters. We varied the learning rate according to the `inverse_square_root` schedule (Vaswani et al., 2017) with a warm-up step of 4000 and a peak learning rate of 0.0005. We trained all MNMT models for 30 epochs and each minibatch contains a maximum of 4096 tokens. The training data of each epoch is composed of the training datasets from all translation directions. For the OPUS-100 dataset, the number of total training steps for the model with and without oversampling are $\sim 933,000$ and $\sim 617,000$ respectively. For the TED-59 dataset, the number of total training steps with and without oversampling are $\sim 226,000$ and $\sim 142,000$ respectively.¹⁴ Parallel sentences in the training data sets where the number of subwords on either the source or target side exceeds 100 were removed.

D Evaluation of MNMT Model

We performed beam search decoding with a beam size of 5 and length penalty of 1.0 during inference. As the TED-59 dataset has already been tokenized, we detokenized reference and system translations with `sacremoses`¹⁵ toolkit before computing BLEU. We chose the best checkpoint according to

¹⁴To reduce the number of padding tokens in the minibatch, Fairseq sorts the training data by comparing the target sentence length first and then the source sentence length by default. Additionally, the number of tokens in each training sample is computed as the maximum of source and target sentence length, which is used to enforce the minibatch size. As a result, the language tag prepending strategy will affect the number of minibatches. Because of this, we report the approximate number of training steps.

¹⁵<https://github.com/alvations/sacremoses>

the average BLEU on the validation sets and then evaluated it on the test sets. Considering that there are no validation sets for zero-shot translation directions, we used the checkpoint selected on the supervised translation directions for zero-shot translation. To be comparable with (Zhang et al., 2020), we employed `langdetect`¹⁶ toolkit for language identification on the OPUS-100 dataset. For TED-59 dataset, we used the `langid.id` toolkit¹⁷ instead as it supports more languages than `langdetect`. We disregarded languages that cannot be detected by `langid.id` toolkit such as Canadian French (fr-ca), Brazilian Portuguese (pt-br) and Burmese (my).

E Dataset for Linguistic Typology Prediction

URIEL is a typological compendium which accommodates diverse linguistic resources from several typological databases such as WALS (Dryer and Haspelmath, 2013), PHOIBLE (Moran and McCloy, 2019) and Glottolog (Hammarström et al., 2021). We used `lang2vec`¹⁸ library to query URIEL database which provides uniform interface to access various linguistic features.

F Prediction Method for Typology Features

We inferred typology features of languages from the language representations derived from the trained MNMT model which achieves the best performance on the validation sets. Some previous works train logistic regression classifiers to perform typology prediction (Malaviya et al., 2017; Oncevay et al., 2020). Nevertheless, logistic regression is a parameterized algorithm and the number of its parameters increase with the feature dimensions of input data, which makes it difficult to handle data scarcity with high-dimensional inputs such as matrix. We hence adopted k -nearest neighbors approach (k-NN), a non-parametric method, for linguistic typology prediction. We employed cosine similarity as the distance measure for LEE. For the language representations in LAA, we computed the average cosine similarity for all rows in the matrix as the distance metric. We set k as odd numbers and varied k in $\{1, 3, 5, 7, 9\}$. We left one language

¹⁶<https://github.com/Mimino666/langdetect>

¹⁷<https://github.com/saffsd/langid.py>

¹⁸<https://github.com/antonisa/lang2vec>

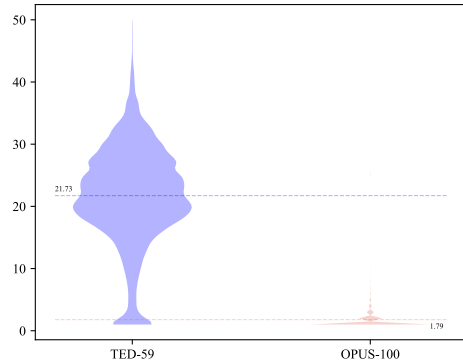


Figure 4: The distribution of the numbers of translations in other languages for each unique English sentence on the TED-59 and OPUS-100 training data. The dashed lines show the mean values.

out and took the remaining languages as training examples to make predictions. This procedure was repeated for each language and the average prediction accuracies on all languages are reported.

G Detailed Many-to-Many Translation Results for LEE

Our findings summarized from Table 2 can be shown as follows:

The ways to indicate the desired target language to MNMT models can significantly affect the translation performance, especially on zero-shot translation directions. There are 0.33 and 5.88 maximum average BLEU differences on supervised (④ vs. ⑥) and zero-shot (⑥ vs. ⑨) translation directions respectively on the TED-59 dataset among the variations of LEE with the same number of parameters. Similarly, although LEE does not use extra parameters compared to $\text{Token}_{\text{src}}$ and $\text{Token}_{\text{tgt}}$, the average BLEU difference on zero-shot translation directions between LEE and $\text{Token}_{\text{src}}$ can be significant on both datasets (6.56 on the TED-59, ⑥ vs. ②) (3.31 on the OPUS-100 dataset, ⑪ vs. ⑫). These indicate the importance of the effective target language information injecting method for MNMT models.

The performance gap between $\text{Token}_{\text{src}}$ and $\text{Token}_{\text{tgt}}$ on zero-shot translation directions varies substantially across different datasets.

$\text{Token}_{\text{src}}$ obtains the best performance in terms of both average BLEU and WR and outperforms $\text{Token}_{\text{tgt}}$ by 8.08 BLEU on zero-shot translation directions on the TED-59 dataset (① vs. ②). However, it lags behind $\text{Token}_{\text{tgt}}$ on zero-shot transla-

tion directions on the OPUS-100 dataset ((10) vs. (11)). In order to find the main reasons behind the performance discrepancy between $\text{Token}_{\text{src}}$ and $\text{Token}_{\text{tgt}}$ on different datasets, we analyzed the two datasets from the perspective of their multi-way alignment since TED talks have been translated into many languages. Specifically, we counted the numbers of translations in other languages for each unique English sentence in the training data of the two datasets.¹⁹ The results are visualized in Figure 4. Although both TED-59 and OPUS-100 are English-centric datasets, the majority of English sentences in the TED-59 dataset have more than one translations in other languages. We hypothesize that identical English sentences on the source side paired with distinct target sentences in different languages may encourage MNMT models to capture the correlations between the prepended tokens of $\text{Token}_{\text{src}}$ and the target languages, which benefits zero-shot translation and is consistent with the finding of Wu et al. (2021).

Prepending a target language token to the source side ($\text{Token}_{\text{src}}$) and to the target side ($\text{Token}_{\text{tgt}}$) benefit $\text{En} \rightarrow \text{XX}$ and $\text{XX} \rightarrow \text{En}$ translation directions respectively The average 0.5 BLEU gain and 91.38% WR on the TED-59 dataset ((1) vs. (2)) together with the average 0.58 BLEU gain and 81.91% WR on the OPUS-100 dataset ((10) vs. (11)) on the $\text{En} \rightarrow \text{XX}$ translation directions indicate that $\text{Token}_{\text{src}}$ is more preferable for $\text{En} \rightarrow \text{XX}$ language pairs than $\text{Token}_{\text{tgt}}$. Moreover, $\text{Token}_{\text{src}}$ achieves the highest WR on $\text{En} \rightarrow \text{XX}$ translation directions across the two datasets. Despite the leading performance of $\text{Token}_{\text{src}}$ on $\text{En} \rightarrow \text{XX}$ translation directions, it suffers from inferior average BLEU and WR on $\text{XX} \rightarrow \text{En}$ translation directions, which results in comparable average BLEU with $\text{Token}_{\text{tgt}}$ over all supervised translation directions. In contrast, $\text{Token}_{\text{tgt}}$ achieves better performance on $\text{XX} \rightarrow \text{En}$ translation directions than $\text{Token}_{\text{src}}$.

Embodying the target language embedding in the decoder is preferable for supervised translation directions The candidate positions for embodying language information are scattered across the encoder (position 1, 2) and decoder (position 3, 4, 5, 6). Table 2 shows that embodying the language embedding into positions in the decoder

¹⁹We collected English sentences by concatenating all parallel corpora of English and eliminating duplicate entries.

can achieve improvements in average BLEU ((1)(2) vs. (5)(6)(7)(8), (10)(11) vs. (12)), while incorporating the language embedding to positions in the encoder will cause performance drop ((1)(2) vs. (3)(4)). This suggests that embodying the language embedding for the target language into positions closer to target translations is the key ingredient to improve model performance on supervised translation directions for LEE.

H Linguistic Typology Prediction Results for LEE

ID	Feature	Dataset	Model	k					Max
				1	3	5	7	9	
(1)	Syntax	TED-59	$\text{Token}_{\text{tgt}}$	81.56	82.15	82.81	81.18	80.79	82.81
(2)		TED-59	$\text{Token}_{\text{src}}$	86.11	86.14	85.15	83.92	82.63	86.14
(3)		TED-59	$\text{LEE}_{4,5}$	82.40	84.70	85.59	84.45	83.42	85.59
(4)		OPUS-100	$\text{Token}_{\text{tgt}}$	82.66	83.97	83.49	83.25	82.98	83.97
(5)		OPUS-100	$\text{Token}_{\text{src}}$	84.80	84.32	83.70	83.93	83.12	84.80
(6)		OPUS-100	$\text{LEE}_{4,5}$	84.42	83.43	83.80	82.34	82.07	84.42
(7)	Phonology	TED-59	$\text{Token}_{\text{tgt}}$	71.96	84.23	67.28	59.21	59.15	84.23
(8)		TED-59	$\text{Token}_{\text{src}}$	80.11	84.97	67.28	59.03	59.03	84.97
(9)		TED-59	$\text{LEE}_{4,5}$	69.68	78.48	66.66	58.97	59.03	78.48
(10)		OPUS-100	$\text{Token}_{\text{tgt}}$	79.57	85.94	81.38	73.78	73.82	85.94
(11)		OPUS-100	$\text{Token}_{\text{src}}$	86.76	86.50	83.94	73.95	73.95	86.76
(12)		OPUS-100	$\text{LEE}_{4,5}$	85.79	80.74	73.82	73.87	74.03	85.79
(13)	Phonetic inventory	TED-59	$\text{Token}_{\text{tgt}}$	86.67	87.86	87.11	85.07	65.80	87.86
(14)		TED-59	$\text{Token}_{\text{src}}$	87.85	88.38	87.44	85.32	66.06	88.38
(15)		TED-59	$\text{LEE}_{4,5}$	86.52	88.03	87.15	85.12	65.73	88.03
(16)		OPUS-100	$\text{Token}_{\text{tgt}}$	87.89	87.68	72.78	59.41	59.47	87.89
(17)		OPUS-100	$\text{Token}_{\text{src}}$	87.86	87.80	72.96	59.47	59.42	87.86
(18)		OPUS-100	$\text{LEE}_{4,5}$	88.38	87.76	72.80	59.42	59.24	88.38

Table 6: Linguistic typology prediction accuracies on syntax, phonology and phonetic inventory features using the language embedding learned by $\text{Token}_{\text{tgt}}$, $\text{Token}_{\text{src}}$ and $\text{LEE}_{4,5}$ which are trained on the TED-59 and OPUS-100 datasets respectively. k denotes the number of nearest neighbors in k-NN. Max denotes the maximum accuracy when k varies in {1, 3, 5, 7, 9}.

I Linguistic Typology Prediction Results for Models Trained on OPUS-100

Figure 5 shows the linguistic typology prediction results for language representations extracted from MNMT models which are trained on the OPUS-100 dataset.

J Many-to-Many Translation Results of $\text{LAA}_{\text{dec.self}}^R$ on Raw Data Distribution

Table 7 presents the many-to-many translation results of $\text{LAA}_{\text{dec.self}}^R$ when oversampling is removed.

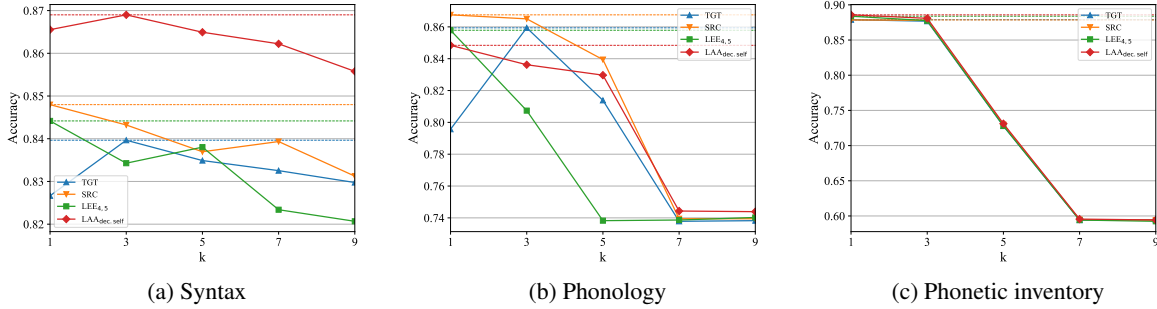


Figure 5: Prediction accuracy on syntax, phonology and phonetic inventory features using the language embeddings learned by $\text{Token}_{\text{tgt}}$, $\text{Token}_{\text{src}}$, $\text{LEE}_{4,5}$ and $\text{LAA}_{\text{dec,self}}$ which are trained on the OPUS-100 dataset.

ID	Dataset	Model	#Param	En \rightarrow XX		XX \rightarrow En		All		Zero-shot		
				BLEU	WR	BLEU	WR	BLEU	WR	BLEU	LangAcc	WR
①	TED-59	$\text{Token}_{\text{tgt}}^-$	77M	19.54	<i>ref</i>	24.23	<i>ref</i>	21.89	<i>ref</i>	2.84	40.94	<i>ref</i>
②	TED-59	$\text{Token}_{\text{src}}^-$	77M	20.25	96.55	23.65	8.62	21.95	52.59	9.65	65.45	96.77
③	TED-59	$\text{LAA}_{\text{dec,self}}^-$	92M	20.64	96.55	25.16	98.28	22.90	97.41	9.93	75.04	97.67
④	TED-59	$\text{LAA}_{\text{dec,self}}^- + \text{LEE}_{4,5}^-$	92M	20.67	94.83	25.05	93.10	22.86	93.97	10.02	75.26	97.82
⑤	TED-59	$\text{LAA}_{\text{dec,self}}^- + \text{Token}_{\text{src}}^-$	92M	19.65 \pm 0.01	61.38 \pm 2.61	25.75 \pm 0.01	90.34 \pm 1.54	22.70 \pm 0.01	75.86 \pm 0.86	9.89 \pm 0.00	73.30 \pm 0.02	95.91 \pm 0.06
⑥	TED-59	$\text{LAA}_{\text{dec,self}}^- + \text{Token}_{\text{tgt}}^-$	92M	18.53 \pm 0.01	1.03 \pm 0.94	26.30\pm0.02	93.10 \pm 0.00	22.41 \pm 0.01	47.07 \pm 0.47	1.02 \pm 0.00	11.37 \pm 0.01	1.06 \pm 0.07
⑦	OPUS-100	$\text{Token}_{\text{tgt}}^-$	77M	21.82	<i>ref</i>	28.45	<i>ref</i>	25.14	<i>ref</i>	6.63	58.76	<i>ref</i>
⑧	OPUS-100	$\text{Token}_{\text{src}}^-$	77M	22.15	74.47	27.68	10.64	24.91	42.55	4.91	37.80	30.00
⑨	OPUS-100	$\text{LAA}_{\text{dec,self}}^-$	103M	23.57	91.49	28.71	70.21	26.14	80.85	11.93	80.39	100.00
⑩	OPUS-100	$\text{LAA}_{\text{dec,self}}^- + \text{LEE}_{4,5}^-$	103M	23.69	91.49	28.88	78.72	26.29	85.11	12.77	85.00	100.00
⑪	OPUS-100	$\text{LAA}_{\text{dec,self}}^- + \text{Token}_{\text{src}}^-$	103M	21.63 \pm 0.01	50.21 \pm 1.39	27.07 \pm 0.02	0.64 \pm 0.58	24.35 \pm 0.01	25.43 \pm 0.69	4.96 \pm 0.02	41.21 \pm 0.05	31.33 \pm 1.83
⑫	OPUS-100	$\text{LAA}_{\text{dec,self}}^- + \text{Token}_{\text{tgt}}^-$	103M	20.75 \pm 0.02	9.15 \pm 1.61	28.11 \pm 0.03	20.00 \pm 1.75	24.43 \pm 0.02	14.57 \pm 0.97	6.80 \pm 0.02	66.74 \pm 0.07	52.00 \pm 1.83

Table 7: Experiment results of $\text{LAA}_{\text{dec,self}}^R$ on the two datasets. The superscript "-" denotes that there is no oversampling. The results of ①②③④⑦⑧⑨⑩ are from Table 4.