

Multilingual Transfer and Domain Adaptation for Low-Resource Languages of Spain

Yuanchang Luo, Zhanglin Wu, Daimeng Wei, Hengchao Shang, Zongyao Li,
Jiixin Guo, Zhiqiang Rao, Shaojun Li, Jinlong Yang,
Yuhao Xie, Jiawei Zheng Bin Wei, Hao Yang

Huawei Translation Service Center, Beijing, China

{luoyuanchang1,wuzhanglin2,weidaimeng,shanghengchao,lizongyao,
guojiaxin1,raozhiqiang,lishaojun18,yangjinlong7,xieyuhao2,
zhengjiawei15,weibin29,yanghao30}@huawei.com

Abstract

This article introduces the submission status of the Translation into Low-Resource Languages of Spain task at (WMT 2024) by Huawei Translation Service Center (HW-TSC). We participated in three translation tasks: spanish to aragonese (es→arg), spanish to aranese (es→arn), and spanish to asturian (es→ast). For these three translation tasks, we use training strategies such as multilingual transfer, regularized dropout, forward translation and back translation, labse denoising, transduction ensemble learning and other strategies to neural machine translation (NMT) model based on training deep transformer-big architecture. By using these enhancement strategies, our submission achieved a competitive result in the final evaluation.

1 Introduction

Neural machine translation (MT) (Lyu et al., 2019; Bahdanau et al., 2014; Gehring et al., 2017) allows translation systems to be trained end-to-end without having to deal with issues like word alignment, translation rules, and complex decoding algorithms that characterize statistical machine translation systems (SMT) (Koehn et al., 2007). Recently, MT technology has evolved towards large language models (LLMs) (Guo et al., 2024). Although neural machine translation has developed rapidly in recent years, it relies heavily on big data - large-scale, high-quality bilingual corpora. Due to the cost and scarcity of real corpora, synthetic data plays an important role in improving translation quality. Existing methods for synthesizing data in NMT focus on leveraging monolingual data during training. Among them, forward translation (Abdulmumin et al., 2021), back translation (Abdulmumin et al., 2021) and data diversity (Nguyen et al., 2020) have been widely used to generate synthetic bilingual corpora. Such synthetic data can be used to improve the performance of NMT

models(Wu et al., 2023b). (Wei et al., 2023) also considers the style of the training data and exploits it to improve performance. Although synthetic data is efficient, synthetic data inevitably contains noise and erroneous translations. Denoising can prevent the training of NMT models from being interfered by noisy synthetic data by introducing high-quality real data as guidance. Another direction to improve the performance of NMT models is to use more efficient training strategies. For example, by mixing similar language data together to train a multi-language pre-training model (Li et al., 2022), due to the shared vocabulary, encoding layer and decoding layer parameters and language similarity, languages with less data can benefit from languages with more data. Regularized dropout (Wu et al., 2021) allows the NMT model to more effectively utilize limited data during the training process. Transduction ensemble learning (Wang et al., 2020) can aggregate the translation capabilities of multiple models into one model.

For the Translation into Low-Resource Languages of Spain task of WMT 2024, we participated in the es→arg, es→arn and es→ast language pair. We use training strategies such as multi-language pre-training models (Li et al., 2022), regularized dropout (Wu et al., 2021), forward translation (Abdulmumin et al., 2021), back translation (Abdulmumin et al., 2021), Labse denoise (Feng et al., 2020) and transduction ensemble learning (Wang et al., 2020) to train neural machine translation (NMT) models based on deep Transformer architecture.

Next, this article will expand on the details of our translation system in different translation tasks. The structure of the remaining sections is as follows: Section 2 introduces the data scale and data preprocessing process; Section 3 describes the overview of the NMT system; Section 4 gives the parameter settings, data processing results and experimental results; Section 6 gives System conclu-

sions were drawn.

2 Dataset

2.1 Data Size

In accordance with the requirements of the WMT 2024 outline, on the Translation into Low-Resource Languages of Spain machine translation task, we used the officially provided data to train the NMT system from scratch. Table 1 shows the training data size for each language pair of the bilingual machine translation task. These language pairs include Spanish to Aragonese (es→arg), Spanish to Arabic (es→arn) and Spanish to Asturian (es→ast).

	es→arg	es→arn	es→ast
Bilingual	0.06M	2.04M	13.36M
Source Monolingual	0.4M	8M	8M
Target Monolingual	0.26M	6M	3M

Table 1: Data size for each bilingual machine translation task

2.2 Data Pre-processing

The data pre-processing process is as follows:

- Remove duplicate sentences or sentence pairs.
- Remove invisible characters and xml escape characters.
- Convert full-width symbols to half-width symbols.
- Use fast_align (Dyer et al., 2013) to filter poorly aligned sentence pairs.
- Filter out sentences with more than 80 tokens in bilingual data.
- Remove sentences with duplicate tokens.
- When performing subword segmentation, joint sentencepiece (Kudo and Richardson, 2018) is used for es→arg, es→arn and es→ast translation tasks.

3 NMT System

3.1 System Overview

Transformer is the state-of-the-art model structure in recent MT evaluations. There are two parts of research to improve this kind: the first part uses wide networks (eg: Transformer-Big (Vaswani, 2017)),

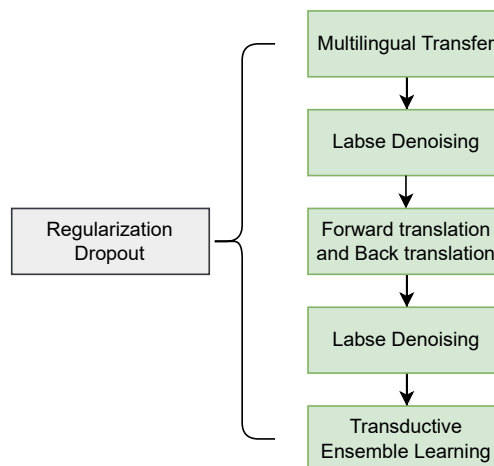


Figure 1: The overall training flow chart of our NMT system on the different translation tasks.

and the other part uses deeper language representations (eg: Deep Transformer (Wang et al., 2019)). For all MT tasks, we combine these two improvements, adopting the Deep Transformer-Big (Wu et al., 2023a) model structure to train the NMT system. Deep Transformer-Big uses pre-layer normalization, features 25-layer encoder, 6-layer decoder, 16-heads self-attention, 1024-dimensional word embedding and 4096-dimensional ffn embedding.

Fig. 1 shows the overall training flow chart of our NMT system on the Translation into Low-Resource Languages of Spain task, we use multilingual transfer (Li et al., 2022), regularization dropout (Wu et al., 2021), forward translation (Abdulmumin et al., 2021), back translation (Abdulmumin et al., 2021), Labse denoise (Feng et al., 2020) and transduction ensemble learning (Wang et al., 2020) and other training strategies are used to train neural machine translation (NMT) models based on deep Transformer-big architecture.

3.2 Multilingual Transfer

Recent researches have shown that multilingual models outperform their bilingual counterparts, particularly when the number of languages in the system is limited and those languages are related (Li et al., 2022). This is mainly due to the capability of the model to learn interlingual knowledge (shared semantic representation between languages). Transfer learning using pre-trained multilingual model has shown very promising results

for low resource tasks. In this task, we first select a multilingual system as the base system, then fine-tune the system with low resource language pairs.

Specifically, we add the "<arg>" tag to the Spanish side of the es→arg bilingual data, the "<arn>" tag to the Spanish side of the es→arn bilingual data, and the "<ast>" tag to the Spanish side of the es→ast bilingual data, and sample them. Mix shuf to train a one-to-many pre-training model; sample the es→arg, es→arn and es→ast original bilingual data and then mix shuf to train a many-to-one pre-training model. Then, the one-to-many pre-training model and the many-to-one pre-training model are trained by using the original bilingual data, and three translation models from Spanish to Aragonese, Arabic, and Asturian and three translation models from Aragonese, Arabic, and Asturian to Spanish are obtained.

3.3 Regularization Dropout

Dropout (Srivastava et al., 2014) is a widely used technique for regularizing deep neural network training, which is crucial to prevent over-fitting and improve the generalization ability of deep models. Dropout performs implicit ensemble by simply dropping a certain proportion of hidden units from the neural network during training, which may cause an unnegligible inconsistency between training and inference. Regularized Dropout (R-Drop) (Wu et al., 2021) is a simple yet more effective alternative to regularize the training inconsistency induced by dropout. Concretely, in each mini-batch training, each data sample goes through the forward pass twice, and each pass is processed by a different sub model by randomly dropping out some hidden units. R-Drop forces the two distributions for the same data sample outputted by the two sub models to be consistent with each other, through minimizing the bidirectional Kullback-Leibler (KL) divergence (Van Erven and Harremos, 2014) between the two distributions. In this way, the inconsistency between the training and inference stage can be alleviated.

3.4 Forward translation and Back translation

Forward translation, also known as self-training (Abdulmumin et al., 2021), is one of the most commonly used data augmentation methods. FT has proven effective for improving NMT performance by augmenting model training with synthetic parallel data. Generally, FT is performed in three steps:

(1) randomly sample a subset from the large-scale source monolingual data; (2) use a "teacher" NMT model to translate the subset data into the target language to construct the synthetic parallel data; (3) combine the synthetic and authentic parallel data to train a "student" NMT model.

Apertium is a free/open-source rule-based architecture for MT that consists of a pipeline of modules performing part-of-speech disambiguation and tagging, lexical transfer, lexical selection, chunk-level or recursive structural transfer, and morphological generation. To make our model better, we use Apertium as a "teacher" model to produce pseudo-corpus.

Back translation (BT) (Abdulmumin et al., 2021) refers to translating the target monolingual data into the source language, and then using the synthetic data to increase the training data size. This method has been proven effective to improve the NMT model performance.

We use the machine translation model obtained by Multilingual Transfer to produce back translation synthetic parallel data, and mix it with forward translation synthetic parallel data and authentic parallel data for training, which can achieve better results than FT or BT.

3.5 Labse Denoising

Due to the low quality of our bilingual data, we use LaBSE (Feng et al., 2020) to calculate the semantic similarity of each bilingual sentence pair and exclude bilingual sentence pairs with similarity scores below 0.7 from our training corpus. Use these clean data to better train the model.

3.6 Transductive Ensemble Learning

Ensemble learning (Garmash and Monz, 2016), which aggregates multiple diverse models for inference, is a common practice to improve the accuracy of machine learning tasks. However, it has been observed that the conventional ensemble methods only bring marginal improvement for NMT when individual models are strong or there are a large number of individual models. Transductive Ensemble Learning (TEL) (Wang et al., 2020) study how to effectively aggregate multiple NMT models under the transductive setting where the source sentences of the test set are known. TEL uses dev sets finetune a strong model, which boosts strong individual models with significant improvement and benefits a lot from more individual models.

	BLEU			ChrF++		
	es→arg	es→arn	es→ast	es→arg	es→arn	es→ast
FLORES+ dev sets	es→arg	es→arn	es→ast	es→arg	es→arn	es→ast
NMT baseline	38.5	8.5	17.3	64.6	34.3	46.6
+ FT & BT	41.7	9.5	16.9	64.8	34.9	45.5
+ Labse denoising	48	10.1	17.5	72.4	38.8	47.5
FLORES+ devtest sets	es→arg	es→arn	es→ast	es→arg	es→arn	es→ast
+ TEL	63	26.3	19.8	80.3	47.9	52.2

Table 2: BLEU and ChrF++ scores of es→arg, es→arn and es→ast NMT systems

4 Experiment

4.1 Setup

We use the open-source fairseq (Ott et al., 2019) to train NMT models, and then use SacreBLEU (Post, 2018) and Chrf++ to measure system performance. The main parameters are as follows: each model is trained using 8 V100 GPUs, batch size is 4096, parameter update frequency is 1, and learning rate is $5e-4$. The number of warmup steps is 4000, and model is saved every 1000 steps. The architecture we used is described in section 3.1. We adopt dropout, and the rate varies across different training phases. R-Drop (Srivastava et al., 2014) is used in model training, and we set λ to 5.

4.2 Data processing

	es→arg	es→arn	es→ast
Bilingual	0.06M	2.04M	13.36M
Data Pre-processing	0.04M	1.51M	3.91M
Labse Filter	0.03M	1.16M	1.92M
Upsampling	0.56M	1.74M	1.92M

Table 3: Data size for each bilingual machine translation task after data pre-processing

Due to the poor quality of bilingual data in low-resource languages, after the rule cleaning mentioned in section 2.2 and the labse model cleaning mentioned in section 3.2, the amount of data is smaller, and the data amount of es→arg, es→arn and es→ast is quite different. When training one-to-many and many-to-one pre-training models, if the amount of bilingual data for a certain language direction is too small, the translation quality will be extremely poor. Therefore, Following (Conneau and Lample, 2019; Liu et al., 2020) we re-balance the training set by upsampling data from each language l with a ratio:

$$\lambda_l = \frac{1}{p_l} \frac{p_l^{1/T}}{\sum_{l=1}^n p_l^{1/T}} \quad \text{with} \quad p_l = \frac{n_l}{\sum_{l=1}^n n_l}$$

where, T is the temperature parameter and we set T to 2. n_l is the number of utterances for language l in the training set. The data amount changes as shown in the following table 3.

4.3 Results

Tables 2 shows the evaluation results of es→arg, es→arn and es→ast NMT systems on the brand new FLORES+ dev sets and devtest sets, the results of dev test sets are obtained through OCELoT submission. We use Multilingual Transfer and R-Drop to build a strong baseline, then use FT and BT for data enhancement, and use Labse denoising for more efficient training, and finally use Transductive Ensemble Learning to ensemble multiple models ability.

As can be seen from the table above, after FT & BT and Labse denoising, the translation quality from Spanish to three directions has been improved to varying degrees. This shows that for low-resource scenarios, these two strategies can expand the amount of data and improve the quality of the data. Enhance the translation quality of machine translation models. Among them, the improvement of both strategies in the es→arg direction is higher than that of the other two directions, and the bilingual data of es→arg is also the least. This shows that FT & BT’s strategy of expanding the amount of data and labse denoising’s strategy of improving data quality are both in situations where the amount of bilingual data is small, The effect is more obvious.

In addition, after Transductive Ensemble Learning, the BLEU value of FLORES+ devtest sets has been greatly improved compared to the FLORES+ dev sets test set. Although it is not the same test set, the BLEU value has improved across latitudes, which shows that The fields of dev sets and devtest sets are very consistent, and Transductive Ensemble Learning, a strategy that utilizes dev sets, can maximize the translation effect of the model on the

test set in the same field.

5 Conclusion

This paper presents HW-TSC’s submission to the Translation into Low-Resource Languages of Spain task of WMT 2024. For both translation tasks, we use a series of training strategies to train NMT models based on the deep Transformer-big architecture. By using these enhancement strategies, our submission achieves a competitive result in the final evaluation. For example, #607 in the spanish to aragonese constrained submissions, #608 in the spanish to aranese constrained submissions, and #606 in the spanish to asturian constrained submissions.

References

- Idris Abdulmumin, Bashir Shehu Galadanci, and Abubakar Isa. 2021. Enhanced back-translation for low resource neural machine translation using self-training. In *Information and Communication Technology and Applications: Third International Conference, ICTA 2020, Minna, Nigeria, November 24–27, 2020, Revised Selected Papers 3*, pages 355–371. Springer.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Alexis Conneau and Guillaume Lample. 2019. Cross-lingual language model pretraining. *Advances in neural information processing systems*, 32.
- Chris Dyer, Victor Chahuneau, and Noah A Smith. 2013. A simple, fast, and effective reparameterization of ibm model 2. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 644–648.
- Fangxiaoyu Feng, Yinfei Yang, Daniel Cer, Naveen Arivazhagan, and Wei Wang. 2020. Language-agnostic bert sentence embedding. *arXiv preprint arXiv:2007.01852*.
- Ekaterina Garmash and Christof Monz. 2016. Ensemble learning for multi-source neural machine translation. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 1409–1418.
- Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N Dauphin. 2017. Convolutional sequence to sequence learning. In *International conference on machine learning*, pages 1243–1252. PMLR.
- Jiaxin Guo, Hao Yang, Zongyao Li, Daimeng Wei, Hengchao Shang, and Xiaoyu Chen. 2024. A novel paradigm boosting translation capabilities of large language models. In *Findings of the Association for Computational Linguistics: NAACL 2024, Mexico City, Mexico, June 16-21, 2024*, pages 639–649. 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, et al. 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. Association for Computational Linguistics.
- Taku Kudo and John Richardson. 2018. Sentencepiece: A simple and language independent subword tokenizer and detokenizer for neural text processing. *arXiv preprint arXiv:1808.06226*.
- Shaojun Li, Yuanchang Luo, Daimeng Wei, Zongyao Li, Hengchao Shang, Xiaoyu Chen, Zhanglin Wu, Jinlong Yang, Zhiqiang Rao, Zhengzhe Yu, et al. 2022. Hw-tsc systems for wmt22 very low resource supervised mt task. In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 1098–1103.
- Yinhan Liu, Jiatao Gu, Naman Goyal, Xian Li, Sergey Edunov, Marjan Ghazvininejad, Mike Lewis, and Luke Zettlemoyer. 2020. Multilingual denoising pre-training for neural machine translation. *Transactions of the Association for Computational Linguistics*, 8:726–742.
- He Lyu, Ningyu Sha, Shuyang Qin, Ming Yan, Yuying Xie, and Rongrong Wang. 2019. Advances in neural information processing systems. *Advances in neural information processing systems*, 32.
- Xuan-Phi Nguyen, Shafiq Joty, Kui Wu, and Ai Ti Aw. 2020. Data diversification: A simple strategy for neural machine translation. *Advances in Neural Information Processing Systems*, 33:10018–10029.
- 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.
- 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.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks

- from overfitting. *The journal of machine learning research*, 15(1):1929–1958.
- Tim Van Erven and Peter Harremo. 2014. Rényi divergence and kullback-leibler divergence. *IEEE Transactions on Information Theory*, 60(7):3797–3820.
- Ashish Vaswani. 2017. Attention is all you need. *arXiv preprint arXiv:1706.03762*.
- Qiang Wang, Bei Li, Tong Xiao, Jingbo Zhu, Changliang Li, Derek F Wong, and Lidia S Chao. 2019. Learning deep transformer models for machine translation. *arXiv preprint arXiv:1906.01787*.
- Yiren Wang, Lijun Wu, Yingce Xia, Tao Qin, ChengXiang Zhai, and Tie-Yan Liu. 2020. Transductive ensemble learning for neural machine translation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 6291–6298.
- Daimeng Wei, Zhanglin Wu, Hengchao Shang, Zongyao Li, Minghan Wang, Jiaxin Guo, Xiaoyu Chen, Zhengzhe Yu, and Hao Yang. 2023. **Text style transfer back-translation**. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2023, Toronto, Canada, July 9-14, 2023*, pages 7944–7959. Association for Computational Linguistics.
- Lijun Wu, Juntao Li, Yue Wang, Qi Meng, Tao Qin, Wei Chen, Min Zhang, Tie-Yan Liu, et al. 2021. R-drop: Regularized dropout for neural networks. *Advances in Neural Information Processing Systems*, 34:10890–10905.
- Zhanglin Wu, Daimeng Wei, Zongyao Li, Zhengzhe Yu, Shaojun Li, Xiaoyu Chen, Hengchao Shang, Jiaxin Guo, Yuhao Xie, Lizhi Lei, et al. 2023a. Treating general mt shared task as a multi-domain adaptation problem: Hw-tsc’s submission to the wmt23 general mt shared task. In *Proceedings of the Eighth Conference on Machine Translation*, pages 170–174.
- Zhanglin Wu, Zhengzhe Yu, Zongyao Li, Daimeng Wei, Yuhao Xie, Xiaoyu Chen, Hengchao Shang, Jiaxin Guo, Zhiqiang Rao, Shaojun Li, et al. 2023b. Hw-tsc’s neural machine translation system for ccmt 2023. In *China Conference on Machine Translation*, pages 13–27. Springer.