HW-TSC's Participation in the WMT 2024 QEAPE Task

Jiawei Yu[1]*, Xiaofeng Zhao[2], Min Zhang[2], Yanqing Zhao[2], Yuang Li[2], Chang Su[2], Xiaosong Qiao[2], Miaomiao Ma[2], Hao Yang[2]
[1]School of Informatics, Xiamen University, China
[2]Huawei Translation Services Center, Beijing, China
yujiawei@stu.xmu.edu.cn
{zhaoxiaofeng14,zhangmin186,yanghao30}@huawei.com

Abstract

The paper presents the submission by HW-TSC in the WMT 2024 Quality-informed Automatic Post Editing (QEAPE) shared task for the English-Hindi (En-Hi) and English-Tamil (En-Ta) language pair. We use LLM for En-Hi and Transformer for EN-ta respectively. For LLM, we first continue pertrain the Llama3, and then use the real APE data to SFT the pre-trained LLM. As for the transformer in En-Ta, we first pre-train a Machine Translation (MT) model by utilizing MT data collected from the web. Then, we fine-tune the model by employing real APE data. We also use the data augmentation method to enhance our model. Specifically, we incorporate candidate translations obtained from an external Machine Translation (MT) system. Given that APE systems tend to exhibit a tendency of 'over-correction', we employ a sentence-level Quality Estimation (QE) system to select the final output, deciding between the original translation and the corresponding output generated by the APE model. Our experiments demonstrate that pre-trained MT models are effective when being fine-tuned with the APE corpus of a limited size, and the performance can be further improved with external MT augmentation. our approach improves the HTER by -15.99 points and -0.47 points on En-Hi and En-Ta, respectively.

1 Introduction

Automatic Post-Editing (APE) is a post-processing task in a Machine Translation (MT) workflow, aiming to automatically identify and correct errors in MT outputs (Chatterjee et al., 2020a). WMT has been holding APE task competitions in different languages and fields since 2015. Different from previous years, this year's APE task is a subtask of the QE task, named Quality-informed automatic post-editing (QEAPE) (Zerva et al., 2024). It proposes to combine quality estimation and automatic post-editing in order to correct the output of machine translation. Participants are provided with a training set comprising 7,000 instances, a development set, and a test set, with each containing 1,000 instances. Each dataset consists of triplets — the source (*src*) sentences, the corresponding machine-translation (*mt*) outputs, and the human post-edited versions (*pe*) of the translations along with sentence-level QE annotations. Additionally, participants are permitted to utilize any additional data for systems training.

Typically, training an APE model requires large amount of training data. However, obtaining *pe* is an expensive task in terms of time and money. As a result, there exists a scarcity of large-scale APE datasets.

To address this challenge, numerous data augmentation techniques have been proposed (Junczys-Dowmunt and Grundkiewicz, 2016; Negri et al., 2018; Lee et al., 2020; Wei et al., 2020; Zhang et al., 2023). Wei et al. (2020) augment the APE training data with translations generated using a different MT system. Huang et al. (2022) train an external MT to obtain more datasets consistent with APE tasks. They also use Google translation to back translate the post-edits in the training set. Deoghare and Bhattacharyya (2022) augment the APE data by generating phrase-level APE triplets using SMT phrase tables. To ensure the quality of the synthetic data, they employ the LaBSE technique (Feng et al., 2022) to filter low-quality triplets.

We first collect our pre-training MT data from NLLB (Team et al., 2022), OpenSubtitles [1], TED2020 (Reimers and Gurevych, 2020), etc. To ensure the quality of the MT data, we use the LaBSE technique (Feng et al., 2022) and filter low-quality data. In our method, we use Google translation to back translate the post-edits in the training set. Subsequently, our dataset is structured as follows: the concatenation of source sentence, back

*Work done during internship at Huawei

[1]https://www.opensubtitles.org/en/search/subs

Proceedings of the Ninth Conference on Machine Translation (WMT), pages 541–546
November 15–16, 2024.

translation and machine translation as the input, while the post-edits serve as the reference output.

Chatterjee et al. (2020b) have proven that APE systems often make unnecessary edits to translation output. To mitigate this issue of over-correction, we employ a sentence-level QE system to determine the final output, selecting between the APE system's output and the original machine-translated (*mt*) version.

Reflecting on the historical development, 2023 is recognized as the inaugural year for large-scale models, with researchers transitioning a variety of tasks to these models, including APE. Notable studies include those that combine Neural Machine Translation (NMT) with Large Language Models (LLM) for APE (Koneru et al., 2024), and comprehensive multi-stage, multilingual large models such as Tower (Alves et al., 2024b), which integrate both MT and APE. Drawing inspiration from Tower, our evaluation utilizes the continued pretraining (CPT) and supervised fine-tuning (SFT) to explore the potential of LLM.

When being evaluated on the test set, our approach improves the HTER (Snover et al., 2006) by -15.99 points and -0.47 points on En-Hi and En-Ta, respectively.

The contributions of our work are as follows:

- We filter low-quality MT data from the collected data using LaBSE-based filtering.
- We propose an APE paradigm based on LLM, including CPT and SFT.
- We utilize Google translation to back translate the post-edits to get *src'* for data augmentation.
- We employ a sentence-level QE system to select the most appropriate output, choosing between the APE-generated output and the original translation.

2 Related Work

Last year's WMT23 APE shard task mainly focuses on transfer learning and data augmentation. Yu et al. (2023) use a Transformer pre-trained on the provided synthetic APE data and then fine-tuned on the real APE data. Additionally, they utilize an external MT system to generate back-translations (with Google Translate [2] run on the post-edits in the training set). They also integrate En-Mr parallel sentences from FLORES-200 (Costa-jussà et al., 2022). R-Drop (Liang et al., 2021), which regularizes the training inconsistency induced by dropout, is used to mitigate overfitting during the training phase. Besides, they use a sentence-level QE system to select the final output between the APE-generated output and the original translation.

Moon et al. (2023) center on data filtering techniques. With a focus on removing potentially harmful material from a model training perspective, the proposed method concentrates on eliminating the two extremes of the training data distribution: the (near-) perfect MT outputs on one side, and those that require complete rewriting on the other.

Another team "kaistai" is inspired by the recent surge of (LLMs) that have been successfully applied in a variety of language generation tasks. They use an LLM with specific prompts designed to generate either (a) post-edits or (b) post-edits along with the rationales behind them.

With experience in previous competitions, we also utilize an external MT system to generate back-translations in our transformer-based system. Additionally, we adopt a sentence-level QE system to select the final output.

3 Dataset

3.1 Data source

We first collect our MT data from the web, mainly from NLLB, OpenSubtitles, TED2020, etc. Then we filter the low-quality data using LABSE. After filtering, we get 3M En-Hi and 3M En-Ta parallel MT data. We first use our filtered MT data with 3M instances to pre-train our model. Then, we use the WMT24 official En-Hi and En-Ta APE datasets for fine-tuning, which consists of a training set and a development set. The training set for both language directions contains 7,000 APE triplets.

4 Method

4.1 LABSE filter

Before using the collected MT data to pretrain our model, we filter the low-quality parallel data by using the LaBSE-based filtering (Feng et al., 2022). We do this to ensure the quality of the MT data. To do so, we first generate embeddings of the En and Hi/Ta using the LaBSE model and normalize them. Then, we compute the cosine similarity between these normalized embeddings. We select the top

[2] https://translate.google.com

70% similarity parallel sentences as our filtered MT data.

4.2 LLM CPT + SFT

Due to the generative nature of the APE task, we believe that LLMs are well-suited for this purpose. Based on human evaluations, we have selected the Llama3-8B-Instruct model, which possesses proficiency in Hindi, as our foundational model. Inspired by the TowerInstruct (Alves et al., 2024a), we adopted a technical approach that combines CPT and SFT. Specifically, during the CPT phase, we utilized 3 million English-Hindi parallel corpora and employed LoRA training techniques. In the SFT phase, we created a customized prompt that, along with the training set provided by the organizers, constituted our SFT training dataset. Our prompt is as follows: "You are a post-editor. You improve translations from English to Hindi using the English source and Hindi translation. Do not provide any explanation or correction." The training paradigm is structured as [prompt: src <en2hi> mt <ape> response], where the response corresponds to the labels predicted by the model.

4.3 Fine-tuned Transformer

We basically treat the APE task as an NMT-like problem, which takes *src* and *mt* as input and generates *pe* autoregressively. Following previous works, we use a special token <*s*> to concatenate src and mt to generate the input sentence: [*src,* <*s*>*, mt*], while the target sentence is *pe*. Initially, we pretrain the MT model using the standard Transformer (Vaswani et al., 2017) structure on 3M En-Ta MT training data. Furthermore, we fine-tune the MT model using the APE dataset with the APE training objective. To further solve the problem of data scarcity, following (Yu et al., 2023), we use the Google translation system to create the *src'* from the provided *pe* text. We simply concatenate the *src'* with the original *src* and *mt* to form the new input: [*src,* <*s*>*, src', <s>, mt*]. Then, we use it in the same way as before, aiming to have the model learn complementary information from *src* and *src'*. During inference, the same input [*src,* <*s*>*, src', <s>, mt*] is employed to generate the output, thereby enabling the utilization of the external information derived from *src'*. Since there is no *pe* during inference, we translate the given *mt* into *src'* using Google Translate.

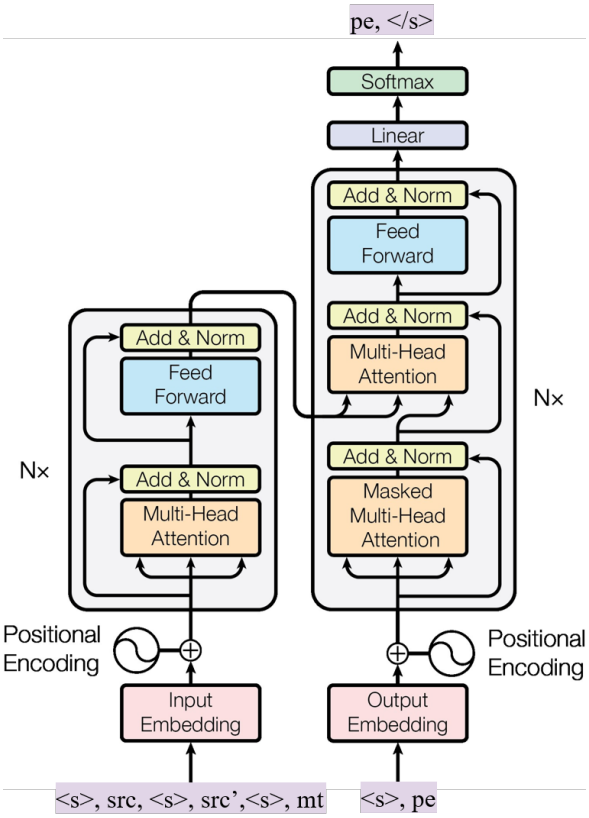

Figure 1: This figure, adapted from (Vaswani et al., 2017) shows the architecture of our model, where mt and augmented src' are concatenated with src before being input into the encoder, and post-edits are generated with the decoder.

4.4 Sentence-Level Quality Estimation

We use wmt22-cometkiwi-da (Rei et al., 2022) as our sentence-level QE model, which is a COMET quality estimation model. This model can be used for reference-free MT evaluation. It receives a source sentence and the respective translation and returns a single score between 0 and 1 that reflects the quality of the translation, where 1 represents a perfect translation. We use this model to rate both the original machine translation and the output generated by our APE system. We then compare the ratings for both sequences and select the one with a higher rating as the final output.

5 Experiment

5.1 Settings

Our transformer model on En-Ta is implemented with fairseq (Ott et al., 2019). Note that the vocabulary and encoder/decoder embeddings of our model are shared between two languages and contain 30K subtokens. We use the batch size of 30,720 tokens in the pre-training stage and 8,192 tokens in

System	En-Hi				En-Ta			
	BLEU↑	**HTER↓**	**ChrF↑**	**COMET↑**	**BLEU↑**	**HTER↓**	**ChrF↑**	**COMET↑**
Baseline (Do nothing)	39.28	46.36	59.48	0.81	70.16	24.71	81.80	0.91
Ours	54.50	30.37	71.06	0.85	69.64	**24.24**	82.36	0.92
swetaagrawal	58.38	**27.08**	73.45	0.86	70.05	24.54	82.30	0.92

Table 1: Results on the WMT24 QE-APE En-Hi and En-Ta test set. A situation with a higher BLEU score but a lower HTER indicates a better result. The official primary evaluation metric for this task is HTER.

the fine-tuning stage. We leverage FP16 (mixed precision) training technique to accelerate training process. In all stages, we apply the Adam optimizer(Kingma and Ba, 2015) with $\beta_1 = 0.9$, $\beta_2 = 0.98$ to train the model, where the inverse square root schedule algorithm and warmup strategy are adopted for the learning rate. Concretely, We use a learning rate of 5e-4 with 20k warm-up steps in the pre-training stage and a learning rate of 5e-5 with 4k warm-up steps in the fine-tuning stage. Besides, we set the dropout to 0.1 in the pre-training stage, 0.3 in the fine-tuning stage, and the value of label smoothing to 0.1 in all stages. Early stopping is adopted with patience 10 and 30 epochs during pre-training and fine-tuning, respectively. During inference, we use beam search with a beam size of 10. Finally, We employ HTER (Snover et al., 2006), BLEU (Papineni et al., 2002), ChrF (Popovic, 2015), and COMET (Rei et al., 2022) as the evaluation metrics.

Our LLM on En-Hi is implemented with Llama-Factory(Zheng et al., 2024). The base model we used is Llama3-8B-Instruction. During the CPT phase, the batch size is set to 256, the learning rate to 1e-4, and training runs for 2 epochs with a precision of bf16. The maximum sequence length is 512 and pre-training is conducted using the LoRA method with a LoRA rank of 64.In the SFT phase, the batch size remains 256, the learning rate is adjusted to 1e-5, and training extends to 8 epochs with bf16 precision. We employ the AdamW optimizer, maintain a maximum sequence length of 512, and utilize PyTorch full_shard for training.

All our transformer models are trained on a Nvidia Tesla V100 GPU with 32GB memory and our LLMs are trained on 64 D910B with 32GB memory.

5.2 Result

Table 1 shows the experimental results evaluated on the test set, where the baseline result is produced by directly calculating scores between the provided MT and PE. We outperform the baseline on HTER for -15.99 and -0.47 points on the En-Hi and En-Ta language pair.

System	En-Hi	
	BLEU↑	**HTER↓**
Baseline (Do nothing)	30.52	58.44
Pretrain+finetune	**49.68**	**36.01**
+External MT	49.01	37.16
+Sentence-level QE	39.13	43.77

Table 2: Results on the WMT24 QE-APE En-Hi development set.

System	En->Ta	
	BLEU↑	**HTER↓**
Baseline (Do nothing)	65.31	29.63
Pretrain+finetune	26.33	57.12
+External MT	33.80	45.31
+Sentence-level QE	**66.11**	**27.66**

Table 3: Results on the WMT24 QE-APE En-Ta development set.

Table 2 shows the En-Hi experimental results evaluated on the dev set. The baseline denotes the test MT result. As illustrated in table 2, the HTER decreased from 58.44 to 36.01 after applying CPT+SFT, reflecting a reduction of 22.43. However, no performance improvement was observed with the addition of back-translation data. We hypothesize that this is due to the sufficiently robust performance of the CPT+SFT, which diminishes the impact of the back-translation data on further enhancement. Upon integrating QE labels, the HTER increased to 43.77 compared with CPT+SFT, an increase of 7.76. We think the QE label may not be accurate enough in En-Hi, resulting in performance loss.

Table 3 shows the En-Ta experimental results evaluated on the dev set. The first experiment is performed by fine-tuning all parameters of the pre-trained Transformer on the official training set, which increases by 27.49 in HTER compared with the baseline. Due to the lack of high-quality En-Ta MT data, the pre-training MT datasets we collected were mostly synthetic and of poor quality. This hinders the capabilities of MT models, which further results in fine-tuned APE models that also perform poorly. The experiment of adding external MT for data augmentation shows some improvement in performance. Toward the end, we utilize a sentence-level QE system to rate both the original translation and the APE output. We then select one of them with a higher rating as the final output of our APE system. With the combination of the APE model and sentence-level QE system, we see that the HTER decreases to 27.66, and the BLEU score increases to 66.11 points.

6 Conclusion

This paper presents our APE system submitted to the WMT 2024 QEAPE En-Hi and EN-Ta task. In our approach, we first filter low-quality MT data from the collected data using LaBSE-based filtering. Then we employ the data augmentation method to build the [*src, <s>, src', <s>, mt*] additional training datasets. Besides, We propose an APE paradigm based on LLM, including CPT and SFT. Moreover, we explore the sentence-level QE system to discard low-quality APE outputs. Evaluation of our APE system shows that our approach achieves gains on the WMT-24 APE development and test sets.

7 Acknowledgements

We would like to thank the anonymous reviewers. Their insightful comments helped us in improving the current version of the paper.

References

Duarte M. Alves, José Pombal, Nuno Miguel Guerreiro, Pedro Henrique Martins, João Alves, M. Amin Farajian, Ben Peters, Ricardo Rei, Patrick Fernandes, Sweta Agrawal, Pierre Colombo, José G. C. de Souza, and André F. T. Martins. 2024a. Tower: An open multilingual large language model for translation-related tasks.

Duarte M. Alves, José Pombal, Nuno M. Guerreiro, Pedro H. Martins, João Alves, Amin Farajian, Ben Peters, Ricardo Rei, Patrick Fernandes, Sweta Agrawal, Pierre Colombo, José G. C. de Souza, and André F. T. Martins. 2024b. Tower: An open multilingual large language model for translation-related tasks.

Rajen Chatterjee, Markus Freitag, Matteo Negri, and Marco Turchi. 2020a. Findings of the WMT 2020 shared task on automatic post-editing. In *Proc. of WMT@EMNLP*.

Rajen Chatterjee, Markus Freitag, Matteo Negri, and Marco Turchi. 2020b. Findings of the WMT 2020 shared task on automatic post-editing. In *Proc. of WMT@EMNLP*.

Marta R. Costa-jussà, James Cross, Onur Çelebi, Maha Elbayad, Kenneth Heafield, Kevin Heffernan, Elahe Kalbassi, Janice Lam, Daniel Licht, Jean Maillard, Anna Y. Sun, Skyler Wang, Guillaume Wenzek, Al Youngblood, Bapi Akula, Loïc Barrault, Gabriel Mejia Gonzalez, Prangthip Hansanti, John Hoffman, Semarley Jarrett, Kaushik Ram Sadagopan, Dirk Rowe, Shannon Spruit, Chau Tran, Pierre Andrews, Necip Fazil Ayan, Shruti Bhosale, Sergey Edunov, Angela Fan, Cynthia Gao, Vedanuj Goswami, Francisco Guzmán, Philipp Koehn, Alexandre Mourachko, Christophe Ropers, Safiyyah Saleem, Holger Schwenk, and Jeff Wang. 2022. No language left behind: Scaling human-centered machine translation. *CoRR*, abs/2207.04672.

Sourabh Dattatray Deoghare and Pushpak Bhattacharyya. 2022. IIT bombay's WMT22 automatic post-editing shared task submission. In *Proc. of WMT*.

Fangxiaoyu Feng, Yinfei Yang, Daniel Cer, Naveen Arivazhagan, and Wei Wang. 2022. Language-agnostic BERT sentence embedding. In *Proc. of ACL*.

Xiaoying Huang, Xingrui Lou, Fan Zhang, and Tu Mei. 2022. Lul's WMT22 automatic post-editing shared task submission. In *Proc. of WMT*.

Marcin Junczys-Dowmunt and Roman Grundkiewicz. 2016. Log-linear combinations of monolingual and bilingual neural machine translation models for automatic post-editing. In *Proc. of WMT*.

Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *Proc. of ICLR*.

Sai Koneru, Miriam Exel, Matthias Huck, and Jan Niehues. 2024. Contextual refinement of translations: Large language models for sentence and document-level post-editing. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 2711–2725, Mexico City, Mexico. Association for Computational Linguistics.

WonKee Lee, Jaehun Shin, Baikjin Jung, Jihyung Lee, and Jong-Hyeok Lee. 2020. Noising scheme for data augmentation in automatic post-editing. In *Proc. of WMT@EMNLP*.

Xiaobo Liang, Lijun Wu, Juntao Li, Yue Wang, Qi Meng, Tao Qin, Wei Chen, Min Zhang, and Tie-Yan Liu. 2021. R-drop: Regularized dropout for neural networks. In *Proc. of NeurIPS*.

Hyeonseok Moon, Seugnjun Lee, Chanjun Park, Jaehyung Seo, Sugyeong Eo, and Heuiseok Lim. 2023. What is the resultful data?: Empirical study on the adaptability of the automatic post-editing training data. In *Proceedings of the Eighth Conference on Machine Translation, WMT 2023, Singapore, December 6-7, 2023*. Association for Computational Linguistics.

Matteo Negri, Marco Turchi, Rajen Chatterjee, and Nicola Bertoldi. 2018. ESCAPE: a large-scale synthetic corpus for automatic post-editing. In *Proc. of LREC*.

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.

Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics, July 6-12, 2002, Philadelphia, PA, USA*, pages 311–318. ACL.

Maja Popovic. 2015. chrf: character n-gram f-score for automatic MT evaluation. In *Proc. of WMT@EMNLP*.

Ricardo Rei, Marcos V. Treviso, Nuno Miguel Guerreiro, Chrysoula Zerva, Ana C. Farinha, Christine Maroti, José G. C. de Souza, Taisiya Glushkova, Duarte M. Alves, Luísa Coheur, Alon Lavie, and André F. T. Martins. 2022. Cometkiwi: Ist-unbabel 2022 submission for the quality estimation shared task. In *Proc. of WMT*.

Nils Reimers and Iryna Gurevych. 2020. Making monolingual sentence embeddings multilingual using knowledge distillation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.

Matthew G. Snover, Bonnie J. Dorr, Richard M. Schwartz, Linnea Micciulla, and John Makhoul. 2006. A study of translation edit rate with targeted human annotation. In *Proc. of AMTA*. Association for Machine Translation in the Americas.

NLLB Team, Marta R. Costa-jussà, James Cross, Onur Çelebi, Maha Elbayad, Kenneth Heafield, Kevin Heffernan, Elahe Kalbassi, Janice Lam, Daniel Licht, Jean Maillard, Anna Sun, Skyler Wang, Guillaume Wenzek, Al Youngblood, Bapi Akula, Loic Barrault, Gabriel Mejia Gonzalez, Prangthip Hansanti, John Hoffman, Semarley Jarrett, Kaushik Ram Sadagopan, Dirk Rowe, Shannon Spruit, Chau Tran, Pierre Andrews, Necip Fazil Ayan, Shruti Bhosale, Sergey Edunov, Angela Fan, Cynthia Gao, Vedanuj Goswami, Francisco Guzmán, Philipp Koehn, Alexandre Mourachko, Christophe Ropers, Safiyyah Saleem, Holger Schwenk, and Jeff Wang. 2022. No language left behind: Scaling human-centered machine translation.

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 *Proc. of NeurIPS*.

Daimeng Wei, Hengchao Shang, Zhanglin Wu, Zhengzhe Yu, Liangyou Li, Jiaxin Guo, Minghan Wang, Hao Yang, Lizhi Lei, Ying Qin, and Shiliang Sun. 2020. Hw-tsc's participation in the WMT 2020 news translation shared task. In *Proc. of WMT@EMNLP*.

Jiawei Yu, Min Zhang, Yanqing Zhao, Xiaofeng Zhao, Yuang Li, Chang Su, Yinglu Li, Miaomiao Ma, Shimin Tao, and Hao Yang. 2023. Hw-tsc's participation in the WMT 2023 automatic post editing shared task. In *Proceedings of the Eighth Conference on Machine Translation, WMT 2023, Singapore, December 6-7, 2023*, pages 926–930. Association for Computational Linguistics.

Chrysoula Zerva, Frederic Blain, Jos'e G. C. de Souza, Diptesh Kanojia, Sourabh Deoghare, Nuno M. Guerreiro, Giuseppe Attanasio, Ricardo Rei, Constantin Orasan, Matteo Negri, Marco Turchi, Rajen Chatterjee, Pushpak Bhattacharyya, Markus Freitag, and Andr'e Martins. 2024. Findings of the quality estimation shared task at wmt 2024. are llms closing the gap in qe? In *Proceedings of the Ninth Conference on Machine Translation*. Association for Computational Linguistics.

Min Zhang, Xiaofeng Zhao, Zhao Yanqing, Hao Yang, Xiaosong Qiao, Junhao Zhu, Wenbing Ma, Su Chang, Yilun Liu, Yinglu Li, Minghan Wang, Song Peng, Shimin Tao, and Yanfei Jiang. 2023. Leveraging chatgpt and multilingual knowledge graph for automatic post-editing. In *International Conference on Human-Informed Translation and Interpreting Technology (HiT-IT 2023)*. Accepted for publication.

Yaowei Zheng, Richong Zhang, Junhao Zhang, Yanhan Ye, Zheyan Luo, and Yongqiang Ma. 2024. Llamafactory: Unified efficient fine-tuning of 100+ language models. *CoRR*, abs/2403.13372.