

Netmarble AI Center’s WMT21 Automatic Post-Editing Shared Task Submission

Shinhyeok Oh*, Sion Jang*, Hu Xu*, Shounan An, Insoo Oh
Netmarble AI Center

Abstract

This paper describes Netmarble’s submission to WMT21 Automatic Post-Editing (APE) Shared Task for the English-German language pair. First, we propose a Curriculum Training Strategy in training stages. Facebook Fair’s WMT19 news translation model was chosen to engage the large and powerful pre-trained neural networks. Then, we post-train the translation model with different levels of data at each training stages. As the training stages go on, we make the system learn to solve multiple tasks by adding extra information at different training stages gradually. We also show a way to utilize the additional data in large volume for APE tasks. For further improvement, we apply Multi-Task Learning Strategy with the Dynamic Weight Average during the fine-tuning stage. To fine-tune the APE corpus with limited data, we add some related sub-tasks to learn a unified representation. Finally, for better performance, we leverage external translations as augmented machine translation (MT) during the post-training and fine-tuning. As experimental results show, our APE system significantly improves the translations of provided MT results by -2.848 and +3.74 on the development dataset in terms of TER and BLEU, respectively. It also demonstrates its effectiveness on the test dataset with higher quality than the development dataset.

1 Introduction

Automatic Post-Editing (APE) aims to improve the quality of an existing Machine Translation (MT) system by learning from human-edited samples (Chatterjee et al., 2019, 2020). With the continuous performance improvements of Neural Machine Translation (NMT) systems along with deep learning advancements, developing APE systems has faced a big challenge. Simple translation errors are hard to find in machine translation outputs, and

the remaining errors are still hard to solve. In recent years, transfer learning and data augmentation techniques have shown their efficiency when training models on datasets with limited size (Devlin et al., 2019). Therefore, such approaches are also adopted in APE tasks (Lopes et al., 2019).

Participants in WMT21 APE shared tasks are required to develop systems to automatically post-edit the translation outputs from an unknown MT system. In this year, the same data has been re-post-edited to improve the quality. As a result of performing statistics on the development set, the evaluation scores are 19.057 and 68.79 in terms of TER and BLEU, which are much higher than the scores of last year, 31.374 and 50.37, respectively. The central distribution of TER has shifted to the left compared to last year. We find that the section in the range of 5 to 10 has the most examples, which indicates that over-correction problems should be considered during the APE tasks. In addition, the dataset has been changed in terms of the domain (from IT to Wikipedia), which results in the change in data distribution. Therefore, directly using previous datasets or officially provided synthetic corpus (Junczys-Dowmunt and Grundkiewicz, 2016; Negri et al., 2018) to enlarge the training set of APE tasks might not be appropriate under such circumstances. In work by Yang et al. (2020), considering the change of data distribution, they select to use additional MT candidates as the data augmentation method to improve feature diversity in their APE systems, which significantly improves the APE performance.

Inspired by this idea, we decided to solve the APE task as NMT alike task and utilize the external MT at the fine-tuning stage. However, because of the limited size of the APE corpus and the improvement of MT quality, fine-tune the model only on the APE data, easily reach the performance ceiling in spite of using external translation. To solve the aforementioned issues, existing works for other

*These authors equally contributed to this work.

Natural Language Processing (NLP) tasks have adopted several Multi-task Learning (MTL) methods with the auxiliary task (Whang et al., 2021; Oh et al., 2021). We wondered whether it is possible to apply MTL mechanism with APE task to the fine-tuning stage since MTL trains the model to encourage representation sharing and improve generalization performance. Furthermore it aims to alleviate the data sparsity problem with a limited number of data in each task (Zhang and Yang, 2021). Therefore, we add some related NLP tasks along with the APE task. Our experiment results demonstrate that such approaches can further improve performance.

As mentioned above, large-volume data, such as news translation data and artificial synthetic data, can not be used to enlarge the APE corpus directly during the fine-tuning because of the large gap in data distribution. We wondered if there is a way to apply any learning method to the post-training so that we can utilize more data to train a more robust and powerful model. In work by (Xu et al., 2020), they applied Curriculum Learning according to the difficulty of each example on a single training stage. Inspired by the research, we try to apply Curriculum Learning across multiple training stages. As the training stage increases, we make the system learn to solve the different tasks by gradually providing extra information, described in Section 3 in detail. Extensive experiments show the effectiveness of applying the Curriculum Learning Strategy during the training phase. Finally, We combined these two approaches to make our final APE system, which significantly improves the performance of the APE task.

Our APE system is built based on Transformer (Vaswani et al., 2017) and is post-trained on WMT21 News-Translation Data (Koehn, 2005; Tiedemann, 2012; Rozis and Skadiņš, 2017; Bhatia et al., 2016; Tiedemann, 2012) and artificial synthetic data (Junczys-Dowmunt and Grundkiewicz, 2016; Negri et al., 2018) provided by APE Task with Curriculum Learning Strategy. For fine-tuning, MTL is applied with related NLP sub-tasks such as Part-Of-Speech (POS), Named Entity Recognition (NER), Masked Language Model (MLM), and Keep/Translate are added to the model to reduce the over-fitting as well as achieve better performance, described in Section 4 in detail. For better training efficiency, the Dynamic Weight Average (DWA) mechanism (Liu et al., 2019) is

applied during the MTL to keep the correct balance between these subtasks. Here we summarize our contributions as follows:

- We design Multi-task Learning Strategy (MLS) with DWA to the fine-tuning stage, which improves the training efficiency and the performance significantly.
- We adapt Curriculum Training Strategy (CTS) to our APE system during the post-training across the multiple training stages, which shows the effectiveness in performance. In addition, we showed a way to utilize the additional data in large volumes in APE tasks.

2 Base System

Our system is based on Facebook FAIR’s WMT19 News Translation Model (Ng et al., 2019), which used the big Transformer (Vaswani et al., 2017) and provided the pre-trained weights. We use both of them as our base system. In addition, we utilize data augmentation with external MT, which has been proposed by Yang et al. (2020) to generate the external translated sentence (mt_ext) and help generate the post-editing sentence (pe). An input sentence X that contains a source sentence (src), a translated sentence by the machine translation system (mt), and an external translated sentence (mt_ext) is defined as,

$$X = [src \langle SEP \rangle mt \langle SEP \rangle mt_ext], \quad (1)$$

and output a sequence, $H = [h_{src_0}, h_{src_1}, \dots, h_{src_n}, h_{\langle SEP \rangle}, h_{mt_0}, \dots, h_{mt_m}, h_{\langle SEP \rangle}, h_{mt_ext_0}, \dots, h_{mt_ext_l}] \in \mathbb{R}^{d_h \times (n+m+l+2)}$, where d_h represents a dimension of the encoder, and n, m, l represents the number of tokens for src, mt, mt_ext , respectively. We represent the parameters of the encoder as Θ_s . Then, H is fed into the decoder, and the decoder target is defined as $Y = [pe]$.

3 Curriculum Training Strategy (CTS)

CTS has been inspired by Curriculum Learning (Xu et al., 2020) that is applied according to the difficulty of each example on a single training stage, which has already been applied to our baseline architecture by Ng et al. (2019). In addition, we propose CTS, which applied Curriculum Learning across multiple training stages. CTS aims at step-by-step learning. In an early stage, the system learns to solve easy problems or something that needs to know beforehand and complex problems or target tasks in the later stages.

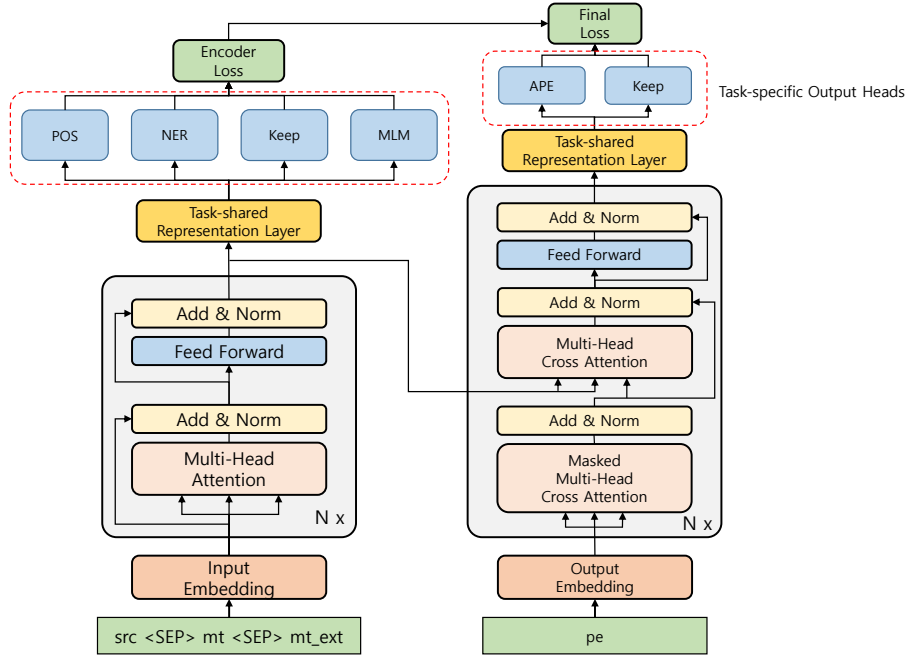


Figure 1: Overall architecture

3.1 Step 1: Understanding for Machine Translation

$$X = [src], \quad (2)$$

The APE task has to understand the machine translation system because the APE task modifies the mt results. Therefore, we designed the first step of the curriculum with the input as Equation 2 and the target as pe .

3.2 Step 2: Learning about Post-Editing

$$X = [src <SEP> mt], \quad (3)$$

After the first step, our system understands as the machine translation system. In this step, we make our system learn how to edit mt to pe with the input as Equation 3 and the target as pe .

3.3 Step 3: Post-Editing with External MT

For the second step, our system learns about the post-editing mechanism. In this step, we make the system learn to take the External MT into account with the input as Equation 1 and the target as pe .

3.4 Fine-Tuning

Finally, we fine-tune the APE system using the data given in the challenge with the input as Equation 1 and the target as pe .

4 Multi-task Learning Strategy (MLS)

Existing works for MTL propose jointly learning methods among related tasks. MTL aims to improve the generalization performance of the whole tasks by sharing knowledge representations of other tasks and can also alleviate the data sparsity problem where each task has limited labeled data (Zhang and Yang, 2021). Therefore, we utilize MLS for our system because WMT21 APE shared task provides only 7,000 train sentences. In NMT, existing works for MTL applied POS, NER, or MLM as subtasks and provided improved results (Chatterjee et al., 2017; Wang et al., 2020). Despite the impressive results, they applied only a few subtasks, such as one or two. Since we defined the APE task as NMT alike problem in our work, it would be helpful to leverage these subtasks into our work to achieve better performance. We find out that all these subtasks are cooperative with each other and benefit our system. Inspired by the word-level quality estimation task, we also add the Keep/Translate classification tasks for encoder and decoder to handle the high-quality APE task, which is described in Section 4.2 in detail. Since utilizing multiple subtasks, we have to consider the loss ratio between these subtasks. In our work, we apply the Dynamic Weight Average method described in Liu et al. (2019), and more details are described in Section 4.4. Our final system based on the model

post-trained using CTS with fine-tuning the APE data with MLS.

4.1 Architecture

Our architecture is described in Figure 1. The overall flow of the APE task is the same in Section 2. In this section, we explain five auxiliary subtasks consisting of POS, NER, MLM, Keep/Translate for the encoder, and Keep/Translate for the decoder. For the encoder, the encoding vector H is fed into Task-shared Representation Layer in Figure 1 like a Fully-connected Neural Network (FNN), and the output is represented as,

$$H_s = (W_1 H + b_1), \quad (4)$$

where $W_1 \in \mathbb{R}^{d_h^s \times d_h}$, and d_h^s represents a dimension of the Task-shared Representation Layer.

4.2 Subtasks

POS & NER POS and NER task aims to predict parts of speech and named entities about an input sequence, respectively. Task-shared Representation Layer H_s is fed into Task-specific Output Heads on Figure 1 like a FNN, and the output is represented as,

$$\hat{Y}^{pos} = \text{softmax}(W_2 H_s + b_2), \quad (5)$$

where $W_2 \in \mathbb{R}^{C_{pos} \times d_h}$ is trainable parameters and C_{pos} is the number of class of POS task. The parameters of Task-specific Output Heads for POS task are represented as Θ_{pos} . Likewise, \hat{Y}^{ner} is obtained as in Equation 5 for NER task, where the parameters are represented as Θ_{ner} .

MLM In MLM task, we copy the input tokens from X to X^{mlm} , which is represented by $X^{mlm} = \{x_1, \dots, x_{n+m+l+2}\}$, where n, m, l represents the number of tokens for src, mt, mt_ext , respectively. Then, we randomly mask 15% of the tokens X^{mlm} using the special token $mask$, and define the target as original input tokens. X^{mlm} is fed into the encoder. Then, the output representation is used to the input for Task-specific Output Heads for MLM task as,

$$\begin{aligned} \hat{Y}^3 &= \text{softmax}(W_3 H_s + b_3), \\ \hat{Y}^{mlm} &= \{\hat{Y}_r^3 | x_r = mask, \\ &\quad \forall r \in \{0, \dots, n + m + l + 2\}\} \end{aligned} \quad (6)$$

where $W_3 \in \mathbb{R}^{C_{mlm} \times d_h}$ represents trainable parameters and C_{mlm} is the number of vocab for the encoder. The parameters of a linear projection layer are represented as Θ_{mlm} for MLM task.

Keep/Translate Considering the characteristics of the APE data with relatively low TER scores, we decide to add Keep/Translate classification subtask to both Encoder and Decoder in our APE system. Keep/Translate subtask aims to predict the labels of the input sequence, where is $\hat{Y}^{kt} \in \{Keep, Translate\}$. In this subtask, each token in the input will be labeled with *Keep* or *Translate*. For label generation, we apply to the pair of src_mt and src_mt . First, we use SimAlign (Jalili Sabet et al., 2020) to perform word alignment on the pe_mt pair. To each aligned word pair, we labeled them with *Keep* if they are equal. Otherwise, they will be marked as *Translate*. As for the pair of src_mt , we also do word alignment to find the correspondence between the source and target side. On the src side, the tokens are labeled with the same name as the corresponding words on the mt side. In our case, the same procedure on pe_mt is conducted for the pair of mt_ext and pe because we use the mt_ext as our data augmentation method. Figure 2 shows an example of label generation in the Keep/Translate task for better understanding. The output is represented as,

$$\hat{Y}^{kt} = \text{softmax}(W_4 H_s + b_4), \quad (7)$$

where $W_4 \in \mathbb{R}^{C_{kt} \times d_h}$ is trainable parameters and C_{kt} is the number of class of Keep/Translate task. The parameters of Task-specific Output Heads for Keep/Translate task are represented as Θ_{kt} and \hat{Y}^{kt} is obtained as in Equation 7 for Keep/Translate task.

4.3 Loss

As described above, five subtasks are used in our system, and most of them have data with imbalanced labels. The imbalanced ratio reaches 1:2160, 1:15, and 1:6 between minority and majority classes in POS Tagger, NER, and Keep/Translate subtasks, respectively. With such imbalanced data, the Cross-Entropy loss used in classification problems may result in performance degradation in some tasks. To improve the performance, the Focal loss (Lin et al., 2017) is considered as an alternative candidate because a Focal Loss function addresses class imbalance during training in tasks. It applies a modulating term to the cross-entropy loss in order to focus on learning the hard negative examples. It reduces the relative loss for well-classified examples ($p_t > 0.5$), putting more focus on hard, misclassified examples. Equation 8 describes the

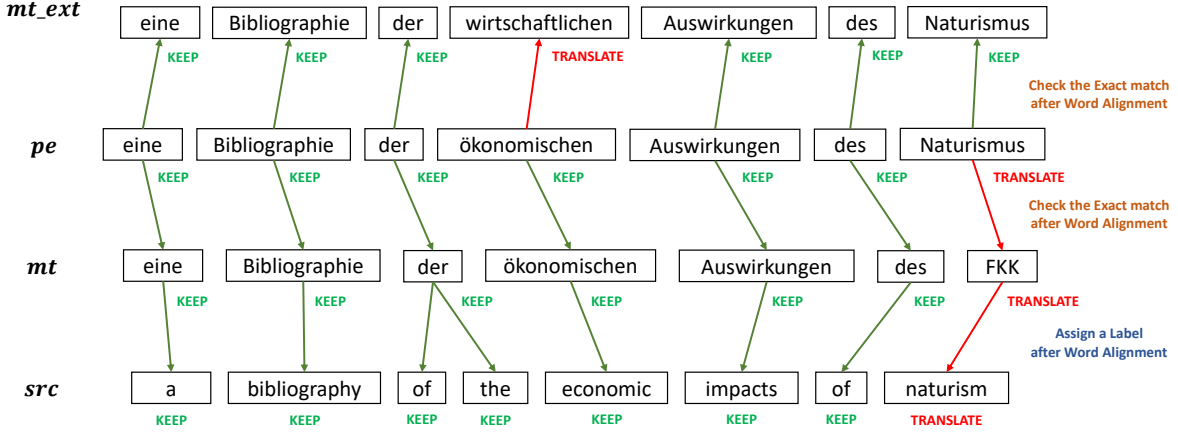


Figure 2: A label generation example in the Keep/Translate task

Focal Loss, where p_t is the probability of each class predicted by the model and γ represents the focusing parameter. Considering the imbalanced property of each task, we apply the Focal Loss to three of our subtasks, such as POS Tagger, NER, and Keep/Translate in the decoder.

$$FL(p_t) = -(1 - p_t)^\gamma \log(p_t) \quad (8)$$

Class Balanced Loss is designed to use a re-weighting scheme that uses the effective number of samples for each class to re-balance the loss, thereby yielding a class-balanced loss (Cui et al., 2019). As the number of samples increases, there is information overlap among data. Therefore, the marginal benefit that a model can extract from the data diminishes. The effective number of samples, which played as the expected volume of samples, is used to capture the diminishing marginal benefits by using more data points of a class. For Keep/Translate task in the decoder, it just considered the PE as input, so we applied the Focal Loss to the subtask. However, for Keep/Translate task in the encoder, as one of the data augmentation methods, the external MT is also considered as input along with the *src* and *mt*. As the information of input increases, we think it may cause information overlap among data because *mt* and the *mt_ext* have the most in common. Therefore, we apply the Class-Balanced Loss as our loss function in Keep/Translate subtask in the encoder. Equation 10 describes the Class-Balanced Loss (L_{cb}), where C is the total number of classes, z_y is the output from the model for class y , n_y is the number of samples in the ground-truth class and $\beta \in [0, 1)$ is a hyperparameter which can be calculated in Equation 9. In Equation 9, i denotes the class

index, $i \in \{1, 2, \dots, C\}$, and N is the number of samples.

As for the MLM task, since it does not suffer from the data imbalance problem, we use the Cross-Entropy loss in our work as other works do.

$$\begin{aligned} N_i &= N, \\ \beta_i &= \beta = (N - 1)/N \end{aligned} \quad (9)$$

$$L_{cb} = -\frac{1 - \beta}{1 - \beta^{n_y}} \log \left(\frac{\exp(z_y)}{\sum_{j=1}^C \exp(z_j)} \right) \quad (10)$$

4.4 Dynamic Weight Average

For most Multi-Task learning networks, it's difficult to find the best ratio between each task in subtasks manually. Therefore, we apply the Dynamic Weight Average (DWA) (Liu et al., 2019) to our work, which adapts the task weighting over time by considering the rate of change of the loss for each task.

Equations 11 and 12 describe DWA. Here, $\lambda_k(\cdot)$ represents the weighting for task k , $w_k(\cdot)$ calculates the relative descending loss rate for each task in each epoch, t is an iteration index, and T represents a temperature that controls the softness of task weighting. \mathcal{L} in Equation 12 is the loss value, calculated as the average loss in each epoch over several iterations.

$$\lambda_k(t) := \frac{\text{Kexp}(\omega_k(t-1)/T)}{\sum_i \text{Kexp}(\omega_i(t-1)/T)} \quad (11)$$

$$\omega_k(t-1) = \frac{\mathcal{L}_k(t-1)}{\mathcal{L}_k(t-2)} \quad (12)$$

	TER	BLEU
CTS-best (ensemble)	16.44	71.88
CTS-best (single)	16.46	71.94
<i>w/o step 2</i>	16.70	71.84
<i>w/o step 3</i>	17.33	70.24
<i>w/o step 2 & 3</i>	17.28	70.88
<i>baseline</i>	19.06	68.79

Table 1: CTS results on WMT21 APE development dataset. CTS-best (ensemble) is built by two similar single models. we submitted CTS-best (ensemble) as CONTRASTIVE result.

4.5 Joint Learning Procedure

All tasks are jointly trained, and the objective is defined as,

$$\mathcal{L} = \frac{1}{K} \sum_i^K \lambda_i \mathbb{L}(Y_i, f(X_i)), \quad (13)$$

where λ is a dynamic weight determining the degree of subtasks and f is the training classifier. Note that the parameter K is the number of subtasks. $\mathbb{L}(Y, f(X))$ is the loss of f w.r.t. the target Y .

5 Experiments

5.1 Datasets

Following existing works, we utilize additional resources (Junczys-Dowmunt and Grundkiewicz, 2016; Negri et al., 2018), which have source sentences (*src*), machine translation sentences (*mt*), and post-editing sentences (*pe*). Moreover, we also utilize some of News-Translation data for the WMT21 (Koehn, 2005; Tiedemann, 2012; Rozis and Skadiņš, 2017; Bhatia et al., 2016; Tiedemann, 2012), which has source sentences (*src*) and translated sentences that can be used as *pe*. For evaluation and fine-tuning, we use the data for WMT21 automatic post-editing shared task. Moreover, we utilize translated sentences using Google Translate and Quality Estimation NMT Model (Fomicheva et al., 2020). The former is used to make *mt_ext* from the additional resources and the data for WMT21 automatic post-editing. The latter is used to make *mt* from News-Translation data. We filtered all the training data based on and number checking logic, which filters the pairs with different numbers in source and target side.

	TER	BLEU
MLS w DWA	16.21	72.53
MLS w/o DWA	16.37	72.34

Table 2: Ablation analysis of DWA on the WMT21 APE development dataset.

5.2 Experimental Settings

For the first step of CTS, we utilize WMT19 en-de weights by Fairseq (Ng et al., 2019). In the second step, we utilize News-Translation data with translated sentences with Quality Estimation NMT Model as *mt*. In the third step, we make our system learn with Junczys-Dowmunt and Grundkiewicz (2016); Negri et al. (2018) and Google Translate as *mt_ext*. Finally, when learning the fine-tuning step, which contains MLS, we utilize the data for WMT21 Automatic Post-Editing shared task.

5.3 Results: CTS

To study the effectiveness of CTS, we conduct ablation experiments on WMT21 Automatic Post-Editing development dataset. We set the *baseline*, which is a system that leaves all the test instances unmodified. As shown in Table 1, we can observe that the *step 3* is more effective than the *step 2*, and that using only *step 2* doesn’t help APE. As our system is learning step by step with CTS, it allows that our system has strengths in the APE task.

5.4 Results: MLS

Table 2 presents the ablation analysis about DWA when fine-tuning with MLS on WMT21 APE development set. From the result, we can observe that MLS with DWA has better performance than the one without applying it. For that reason, we adopt DWA at a fine-tuning stage with MLS in our APE Task.

To find the best combination of subtasks in MLS, we conducted an ablation analysis on the same development dataset. Vanilla in the table is a system without adding any subtasks. We add the subtasks one by one during the fine-tuning to see the effect of each subtask on the performance. As shown in Table 3, the one using all the subtasks performs best among all the combinations, which means that these subtasks are cooperative in the APE task.

	TER	BLEU
Vanilla	16.71	71.75
w/ POS	16.49	72.12
w/ NER	16.52	72.19
w/ MLM	16.55	72.00
w/ Keep/Translate	16.45	72.32
Fine-tuned with MLS	16.21	72.53

Table 3: The Multi-task Learning results on WMT21 APE validation dataset. Fine-tuned with MLS using all subtasks model is submitted as PRIMARY result.

	TER	BLEU
Netmarble_CONTRASTIVE	17.28	71.55
Netmarble_PRIMARY	17.97	70.53
<i>baseline</i>	18.05	71.07

Table 4: Official results on WMT21 APE test dataset.

5.5 Official Results

Table 4 shows the official results of our proposed methods on WMT21 test dataset. The test dataset has baseline scores of 18.05 and 71.07, which is higher than the development dataset with 19.06 and 68.79 in terms of TER and BLEU, respectively. Despite its high quality, our proposed methods showed effectiveness on this test dataset.

5.6 Implementation Details

We set the batch size to 256 for the *step 2* and *step 3* in CTS at each GPU, 16 for the fine-tuning and MLS. We set the initial learning rate to $1e-4$ using scheduler in Fairseq (Ng et al., 2019) for all experiments. The average runtime of one epoch for each approach was about 360 minutes for the *step 2*, 90 minutes for the *step 3*, and 40 seconds for MLS. We train our models using AdamW (Loshchilov and Hutter, 2019) optimizer and conduct experiments with 16 Tesla A100 GPUs for CTS, Tesla V100 GPU for MLS.

6 Conclusion

In this paper, we propose an APE system based on CTS and MLS. CTS allows understanding between machine translation and automatic post-editing, and shows a way using additional data in large volume in APE task. MLS learns a shared unified representation from related subtasks to improve the performance. We submitted the system, which

Fine-tunes with MLS, as our primary version and the ensembled CTS as our contrastive version. The experimental results show that our system is able to effectively detect and correct the errors made by a high-quality NMT system, improving the score by -2.848 and $+3.74$ on the development dataset in terms of TER and BLEU, respectively. Our proposed methods also achieved performance improvement on the test dataset with higher quality.

References

- K. Bhatia, K. Dahiya, H. Jain, P. Kar, A. Mittal, Y. Prabhu, and M. Varma. 2016. [The extreme classification repository: Multi-label datasets and code](#).
- Rajen Chatterjee, Christian Federmann, Matteo Negri, and Marco Turchi. 2019. [Findings of the WMT 2019 shared task on automatic post-editing](#). In *Proceedings of the Fourth Conference on Machine Translation (Volume 3: Shared Task Papers, Day 2)*, pages 11–28, Florence, Italy. Association for Computational Linguistics.
- Rajen Chatterjee, Markus Freitag, Matteo Negri, and Marco Turchi. 2020. [Findings of the WMT 2020 shared task on automatic post-editing](#). In *Proceedings of the Fifth Conference on Machine Translation*, pages 646–659, Online. Association for Computational Linguistics.
- Rajen Chatterjee, Matteo Negri, Marco Turchi, Marcello Federico, Lucia Specia, and Frédéric Blain. 2017. [Guiding neural machine translation decoding with external knowledge](#). In *Proceedings of the Second Conference on Machine Translation, Volume 1: Research Papers*, pages 157–168, Copenhagen, Denmark. Association for Computational Linguistics.
- Yin Cui, Menglin Jia, Tsung-Yi Lin, Yang Song, and Serge Belongie. 2019. [Class-balanced loss based on effective number of samples](#). In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9260–9269.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Marina Fomicheva, Shuo Sun, Lisa Yankovskaya, Frédéric Blain, Francisco Guzmán, Mark Fishel, Nikolaos Aletras, Vishrav Chaudhary, and Lucia Specia. 2020. [Unsupervised quality estimation for neural machine translation](#). *Transactions of the Association for Computational Linguistics*, 8:539–555.

- Masoud Jalili Sabet, Philipp Dufter, François Yvon, and Hinrich Schütze. 2020. [SimAlign: High quality word alignments without parallel training data using static and contextualized embeddings](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1627–1643, Online. Association for Computational Linguistics.
- Marcin Junczys-Dowmunt and Roman Grundkiewicz. 2016. [Log-linear combinations of monolingual and bilingual neural machine translation models for automatic post-editing](#). In *Proceedings of the First Conference on Machine Translation: Volume 2, Shared Task Papers*, pages 751–758, Berlin, Germany. Association for Computational Linguistics.
- Philipp Koehn. 2005. [Europarl: A Parallel Corpus for Statistical Machine Translation](#). In *Conference Proceedings: the tenth Machine Translation Summit*, pages 79–86, Phuket, Thailand. AAMT, AAMT.
- Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. 2017. [Focal loss for dense object detection](#). In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 2999–3007.
- Shikun Liu, Edward Johns, and Andrew J. Davison. 2019. [End-to-end multi-task learning with attention](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- António V. Lopes, M. Amin Farajian, Gonçalo M. Correia, Jonay Trénous, and André F. T. Martins. 2019. [Unbabel’s submission to the WMT2019 APE shared task: BERT-based encoder-decoder for automatic post-editing](#). In *Proceedings of the Fourth Conference on Machine Translation (Volume 3: Shared Task Papers, Day 2)*, pages 118–123, Florence, Italy. Association for Computational Linguistics.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Matteo Negri, Marco Turchi, Rajen Chatterjee, and Nicola Bertoldi. 2018. [ESCAPE: a large-scale synthetic corpus for automatic post-editing](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Nathan Ng, Kyra Yee, Alexei Baevski, Myle Ott, Michael Auli, and Sergey Edunov. 2019. [Facebook fair’s wmt19 news translation task submission](#). In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 314–319, Florence, Italy. Association for Computational Linguistics.
- Shinhyeok Oh, Dongyub Lee, Taesun Whang, IINam Park, Seo Gaeun, EungGyun Kim, and Harksoo Kim. 2021. [Deep context- and relation-aware learning for aspect-based sentiment analysis](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 495–503, Online. Association for Computational Linguistics.
- Roberts Rozis and Raivis Skadiņš. 2017. [Tilde MODEL - multilingual open data for EU languages](#). In *Proceedings of the 21st Nordic Conference on Computational Linguistics*, pages 263–265, Gothenburg, Sweden. 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’12)*, pages 2214–2218, Istanbul, Turkey. European Language Resources Association (ELRA).
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Yiren Wang, ChengXiang Zhai, and Hany Hassan. 2020. [Multi-task learning for multilingual neural machine translation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1022–1034, Online. Association for Computational Linguistics.
- Taesun Whang, Dongyub Lee, Dongsuk Oh, Chanhee Lee, Kijong Han, Dong-hun Lee, and Saebyeok Lee. 2021. [Do response selection models really know what’s next? utterance manipulation strategies for multi-turn response selection](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(16):14041–14049.
- Benfeng Xu, Licheng Zhang, Zhendong Mao, Quan Wang, Hongtao Xie, and Yongdong Zhang. 2020. [Curriculum learning for natural language understanding](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6095–6104, Online. Association for Computational Linguistics.
- Hao Yang, Minghan Wang, Daimeng Wei, Hengchao Shang, Jiabin Guo, Zongyao Li, Lizhi Lei, Ying Qin, Shimin Tao, Shiliang Sun, and Yimeng Chen. 2020. [Hw-tsc’s participation at wmt 2020 automatic post editing shared task](#). In *Proceedings of the Fifth Conference on Machine Translation*, pages 797–802, Online. Association for Computational Linguistics.
- Yu Zhang and Qiang Yang. 2021. [A survey on multi-task learning](#). *IEEE Transactions on Knowledge and Data Engineering*, pages 1–1.