

Levenshtein Training for Word-level Quality Estimation

Shuoyang Ding^{†*} Marcin Junczys-Dowmunt[‡] Matt Post^{†‡} Philipp Koehn[†]

[†] Center for Language and Speech Processing, Johns Hopkins University [‡] Microsoft
{dings, phi}@jhu.edu {marcin.junczysdowmunt, mattpost}@microsoft.com

Abstract

We propose a novel scheme to use the Levenshtein Transformer to perform the task of word-level quality estimation. A Levenshtein Transformer is a natural fit for this task: trained to perform decoding in an iterative manner, a Levenshtein Transformer can learn to post-edit without explicit supervision. To further minimize the mismatch between the translation task and the word-level QE task, we propose a two-stage transfer learning procedure on both augmented data and human post-editing data. We also propose heuristics to construct reference labels that are compatible with subword-level finetuning and inference. Results on WMT 2020 QE shared task dataset show that our proposed method has superior data efficiency under the data-constrained setting and competitive performance under the unconstrained setting.

1 Introduction

Quality estimation (QE) is the task of estimating the quality of translation without access to a human-generated reference. Most recent advances on quality estimation (Rei et al., 2020; Thompson and Post, 2020a; Ranasinghe et al., 2020, *inter alia*) focus on estimating the quality of translation on either the corpus or segment-level. However, in practice, the end-users of machine translation (MT) often call for quality signals on more fine-grained level—the level of individual words in a translation. Such signals are not only useful for more fine-grained triage of translation quality, but also open up the potential for targeted post-processing and faster human post-editing.

To automatically assess the quality of translations, it is natural to consider starting from a machine translation model, which has already acquired the translation knowledge. However, it is not

* Shuoyang Ding had a part-time affiliation with Microsoft at the time of this work.

Figure 1 consists of two parts, (a) and (b), illustrating edit operations on the sentence "Vaccinations put an end to the pandemic".

(a) TER-style edit tags: A table comparing the target and post-edit sequence. The target is "Vaccinations put an end to the pandemic" and the post-edit is "Vaccinations put to the pandemic". The TER row shows "OK" for "Vaccinations", "OK" for "put", "BAD" for "an", "BAD" for "end", "OK" for "to", "OK" for "the", and "OK" for "pandemic".

(b) Hypothetical Levenshtein Transformer edits: A table showing "prev. target" and "mask ins." for the same sentence. The "prev. target" row shows "NO" for "Vaccinations", "NO" for "put", "YES" for "an", "NO" for "end", "NO" for "to", "NO" for "the", and "NO" for "pandemic". The "mask ins." row shows "0" for "Vaccinations", "2" for "put", "0" for "an", and "0" for "end". Below this, the "word ins." row shows "Vaccinations", "put", "an end", "to", "the", "pandemic".

Figure 1: Figure (a) shows an example of TER-style edit tags used as reference for word-level quality estimation task. Figure (b) shows a series of hypothetical Levenshtein Transformer edits that generates the same sequence from the target input. The similarity of these edit operations motivates the study in this paper.

clear how to best transfer the translation knowledge to a word-level quality estimation setting, since autoregressive translation models only see the preceding context of the word they generate. Hence, they are not well-equipped to perform word-level quality assessments or edits on an existing translation, as there are both preceding and succeeding context for a translated word. Towards this goal, we leverage Levenshtein Transformer (LevT, Gu et al., 2019), a non-autoregressive neural machine translation (NMT) model trained to generate translations by starting with an empty output sequence and iteratively performing edits on the sequence. Because of this special training and decoding procedure, the model should have already learned to edit an existing translation sentence without supervision from post-edited translations. We then use multi-stage transfer learning to teach the model to perform the actual QE task, first on artificially-crafted pseudo post-editing, then on real human post-edited data.

We show that our method achieves better performance than the currently widely adopted Predictor-Estimator scheme (Kim et al., 2017; Kepler et al., 2019) under the data-constrained setting, while also being competitive when compared with the high-ranking submissions to the WMT 2020 word-level QE shared task (Specia et al., 2020) under the unconstrained setting.

2 Levenshtein Training for Word-level Quality Estimation

2.1 Word-level Quality Estimation

We follow the problem formulation of word-level quality estimation (QE) in WMT QE shared task (Specia et al., 2020). Given the source-side input sentence and MT output pairs, the participants are asked to perform two binary classification tasks: (1) whether each word in the target-side translation is correct or not (*translation tag*), and (2) whether there are missing words in between each pair of output words (*gap tag*). The reference tags for such prediction are generated by performing human post-editing on the MT outputs and construct edit alignments with a Translation Error Rate (TER, Snover et al., 2006) computation tool. An example is shown in Figure 1a. Each submission is evaluated by the Matthews Correlation Coefficient (MCC, Matthews, 1975).

The state-of-the-art approach to this task is based on the Predictor-Estimator (PredEst) architecture (Kim et al., 2017; Kepler et al., 2019). At a very high level, the *predictor* training uses a cross-lingual masked language model (MLM) objective, which trains the model to predict a word in the target sentence given the source and both the left and right context on the target side. An *estimator* is then finetuned from the predictor model to predict word-level QE tags. In recent years, the top-ranking systems also incorporate large-scale pre-trained crosslingual encoder such as XLM-RoBERTa (Conneau et al., 2020), glass-box features (Moura et al., 2020) and pseudo post-editing data augmentation (Wei et al., 2020; Lee, 2020).

2.2 Levenshtein Transformer

Intuitively, translation knowledge is very beneficial for the word-level QE task. Hence, a natural choice for this task is to start from an NMT model and finetune it to produce word-level quality estimation outputs. However, there are two major limitations of NMT model that makes it unfit for this task: (1) Most NMT architectures are trained to perform inference in a left-to-right manner, and are therefore ill-equipped to perform edits on an existing translation output; (2) Most NMT architectures do not have a mechanism to predict whether there words missing at a given location. From our introduction below, the readers should notice that Levenshtein Transformer successfully addresses both of these limitations.

The Levenshtein Transformer (LevT, Gu et al., 2019) is a neural network architecture that can iteratively generate sequences in a non-autoregressive manner. Unlike normal autoregressive sequence models that have only one prediction head A_w to predict the next output words, LevT has two extra prediction heads A_{del} and A_{ins} that predicts deletion and insertion operations based on the output sequence from the previous iteration.

For translation generation, at the k -th iteration during decoding, with source-side input \mathbf{x} and target-side sequence input from the previous iteration $\mathbf{y}^{(k-1)}$ of length J , suppose the decoder block output is $\{\mathbf{h}_0, \mathbf{h}_1, \dots, \mathbf{h}_J\}$. The following predictions and edits will take place in order:

- deletion actions $d_j \in \mathcal{D}^{(k)}$:

$$p_{del}^{(k)}(d_j | \mathbf{x}, \mathbf{y}^{(k-1)}) = \text{softmax}(\mathbf{A}_{del}^T \mathbf{h}_j)$$

$$\text{for } j \in 1 \dots J_0 \text{ and } J_0 = |\mathbf{y}^{(k-1)}|$$

- mask insertion actions $\gamma_j \in \mathcal{S}^{(k)}$:

$$p_{ins}^{(k)}(\gamma_j | \mathbf{x}, \mathbf{y}') = \text{softmax}(\mathbf{A}_{ins}^T [\mathbf{h}_j; \mathbf{h}_{j+1}])$$

$$\text{for } j \in 0 \dots J_1, \mathbf{y}' = \mathcal{D}^{(k)}(\mathbf{y}^{(k-1)}) \text{ and } J_1 = |\mathbf{y}'|$$

- word prediction actions $w_j \in \mathcal{W}^{(k)}$:

$$p_w^{(k)}(w_j | \mathbf{x}, \mathbf{y}'') = \text{softmax}(\mathbf{A}_w^T \mathbf{h}_j)$$

$$\text{for } j \in 1 \dots J_2, \mathbf{y}'' = \mathcal{S}^{(k)}(\mathcal{D}^{(k)}(\mathbf{y}^{(k-1)})) \text{ and } J_2 = |\mathbf{y}''|$$

In the end, $\mathbf{y}^{(k)} = \mathcal{W}^{(k)}(\mathcal{S}^{(k)}(\mathcal{D}^{(k)}(\mathbf{y}^{(k-1)})))$. The iterative process will continue until $\mathbf{y}^{(k)} = \mathbf{y}^{(k-1)}$ or a maximum number of iterations is reached.

For the word-level QE task, we only perform one iteration of the above process – we use the word deletion head A_{del} to predict quality labels for MT words and the mask insertion head A_{ins} to predict quality labels for gaps. We perform neither the word prediction with A_w nor multiple iterations of prediction. Still, one should note the similarity between the function of those prediction heads for translation prediction and for word-level QE, which is our motivation for choosing this specific architecture for word-level QE.

It should also be pointed out that using a Levenshtein Transformer translation model as a pre-trained model shares some similar spirit to ELECTRA (Clark et al., 2020), which performs pre-training by learning to detect corrupted tokens generated from a masked language model.

2.3 Pre-trained Model

To achieve optimal performance and take advantage of multilinguality, we would also like to take advantage of the large-scale pre-training. Because there is no pre-trained LevT translation model available, we choose to incorporate M2M-100 (Fan et al., 2020), a large-scale pre-trained multilingual autoregressive transformer model. Recall that the main architectural difference between LevT and a standard transformer model is the two extra prediction heads on LevT. Hence, to adapt it into a LevT model, we first need to add randomly-initialized extra prediction heads of LevT to the pre-trained checkpoint. These randomly-initialized prediction heads are then trained with the rest of the model on parallel data in order to adapt the autoregressive translation model to a LevT-style non-autoregressive translation model.

3 Transfer Learning from Translation to Word-level Quality Estimation

By training for the translation task, LevT already acquired some initial knowledge of post-editing. However, there is still some train/inference time mismatch for the word-level QE task regarding (1) the target-side context and (2) the edit tags. In terms of target-side context, during training, the target-side context is a noisy version of the real target sentence in the training set; during inference, the target-side context is a translation output from an NMT system. In terms of edit tags, during training, we want the model to predict LCS edit tags that correct the noisy version of the target sentence; during inference, we want the model to predict TER-style (Levenshtein) edit tags that correspond to human post-editing of NMT outputs.

To address such mismatch, we adopt a two-step transfer learning scheme. For both stages of transfer learning, we need *translation triplets* (source, MT output, post-edited output) to perform fine-tuning. However, in practice, human post-editing resources are quite scarce. Hence, like some previous work (Lee, 2020; Wang et al., 2020), we start by performing transfer learning on synthetic translation triplets, followed by real translation triplets constructed with human post-editing.

Synthetic Data Construction We explore four different methods for translation triplet synthesis:

- **src-mt-ref** Take a parallel dataset (src, ref) and translate the source sentence with an MT model (mt).

- **bt-rt-tgt** Take a target-side monolingual dataset (tgt). Translate the target sentence with an backward MT model (bt) and then translate bt again with an forward MT model, thus creating a round-trip translated output (rt). Use rt as the MT output and the original tgt as the pseudo post-edited output.
- **src-mt1-mt2** Take a source-side monolingual dataset (src). Translate the source sentence with a weaker MT model (mt1) and a stronger MT model (mt2). Use mt1 as the MT output and mt2 as the pseudo post-edited output.
- **mvppe** Take the source-side of a parallel dataset (src). Translate the source sentence with a multilingual MT model as the MT output (mt) and build a pseudo post-edited output (pe) with multiview pseudo post-edit (MVPPE) decoding, as described below.

Multiview Pseudo Post-Editing (MVPPE) Inspired by Thompson and Post (2020b) which used a multilingual translation system as a zero-shot paraphraser, we propose a novel pseudo post-editing method to build synthetic post-editing dataset from parallel corpus (src, tgt). The first step is to translate the source side of the parallel corpus with a multilingual translation system as the MT output (mt) in the triplet. We then generate the pseudo-post-edited output by ensembling two different *views* of the same model. These two views are:

- the translation output distribution $p_t(\text{pe} \mid \text{src})$, with src as the model input;
- the paraphrase output distribution $p_p(\text{pe} \mid \text{tgt})$, with tgt as the model input.

Note that both views will create a distribution in the target language space, which can be ensembled in the same way as standard MT model ensembles, forming an interpolated distribution:

$$p(\text{pe} \mid \text{src}, \text{tgt}) \triangleq \lambda_t p_t(\text{pe} \mid \text{src}) + \lambda_p p_p(\text{pe} \mid \text{tgt})$$

with λ_t and λ_p as the interpolation weights. Similarly, beam search can also be performed on top of the ensemble. The intuition behind this idea is that such ensemble should create a target sentence that is semantically equivalent to the target side of the parallel corpus, while being close to the original MT output as much as possible, imitating the way humans perform the post-editing task.

Compatibility with Subwords To the best of our knowledge, previous work on word-level quality estimation either builds models that directly

output word-level tags (Lee, 2020; Hu et al., 2020; Moura et al., 2020) or uses simple heuristics to re-assign word-level tags to the first subword token (Wang et al., 2020). Since LevT predicts edits on a subword-level starting from translation training, we need to: (1) for inference, convert subword-level tags predicted by the model to word-level tags for evaluation, and (2) for both finetuning stages, build subword-level reference tags.

For inference, the conversion can be easily done by heuristics. For finetuning, a *naive subword-level tag reference* can be built by running TER alignments on MT and post-edited text after subword tokenization. However, a preliminary analysis shows that such reference introduces a 10% error after being converted back to word-level. Hence, we introduce another heuristic to create *heuristic subword-level tag references*. The high-level idea is to interpolate the word-level and the naive subword-level references to ensure that the interpolated subword-level tag reference can be perfectly converted back to the word-level references. Details for the subword tag conversions can be found in Appendix A and Algorithm 1 and 2 in the appendix.

4 Experiments

4.1 Setup

Our experiments are based on the setup of the WMT 2020 QE shared task (Specia et al., 2020). The results are reported under two settings: the constrained setting and the unconstrained setting. Apart from the official metric MCC, we also report the F1 score of the OK and BAD tags (F1-OK and F1-BAD in the result tables).

In the constrained setting, we focus on the data efficiency of our model and use only the human-labeled dataset, the NMT model, and the parallel data (used to train the NMT model) provided by the shared task, with neither large-scale pre-trained model nor synthetic data finetuning. In the unconstrained setting, we additionally use some extra resources we have access to. For en-de, we use WMT20 en-de parallel data to train LevT model instead of the smaller parallel data from the shared task, as in the constrained setting. For en-zh, we use the same dataset because it is already close to the WMT data scale. We also experiment with the M2M-100-small initialization (Fan et al., 2020, 418M parameters) as described in Section 2.3. Note that M2M-100 directly applies sentencepiece on untokenized data, a tokenization

	MCC	F1-OK	F1-BAD
en-de			
OpenKiwi	0.358	0.879	0.468
LevT w/o KD	0.441	0.926	0.498
LevT	0.477	0.929	0.535
en-zh			
OpenKiwi	0.509	0.849	0.658
LevT	0.629	0.885	0.741

Table 1: Constrained setting. All LevT models here are transformer-base models. F1-OK and F1-BAD are F1 scores of the OK and BAD tags, respectively. Higher is better for all metrics in this table.

scheme that is incompatible with the shared task setting. For our experiments with M2M-100-small, we proceed with applying sentencepiece on tokenized data during finetuning. We also experiment with synthetic data finetuning with different data synthesis methods on en-de language pair, while for en-zh, because we don’t have access to an extra high-quality MT system, we only experiment with mvppe method. Details for our data synthesis setup can be found in Appendix B.1.

4.2 Results

Table 1 shows results under the data-constrained setting. Even without knowledge distillation (KD) during LevT training, our model already scores much higher than the baseline OpenKiwi system. When training with KD data generated with the shared task NMT system (trained on the same parallel data), the advantage of the LevT expands even more. This shows that our proposal to build a word-level QE system from LevT translation models has higher data efficiency than the widely adopted PreEst approach used by the OpenKiwi baseline.

Table 2 shows results under the unconstrained setting. We first notice that finetuning with synthetic data before human post-edited data almost always helps. With the transformer-base model on en-de language pair, we experimented with all four data synthesis methods, and we find that src-mt1-mt2 performs the best, closely followed by mvppe. This might be related to the fact that all LevT models are trained with KD data. Because of this, the model is better posed to fit synthesized data with MT-like output as pseudo post-edited data, instead of human-generated translations. Also, initializing with the M2M-100 model is helpful despite the tokenization scheme mismatch, although the performance gain is much more modest on en-

Init.	LevT	Data Synth.	MCC	F1-OK	F1-BAD
en-de					
N	base	N	0.539	0.925	0.613
N	base	src-mt-ref	0.542	0.925	0.616
N	base	bt-rt-tgt	0.535	0.925	0.609
N	base	mvppe	0.548	0.926	0.620
N	base	src-mt1-mt2	0.549	0.926	0.622
N	big	N	0.551	0.927	0.623
N	big	src-mt1-mt2	0.562	0.939	0.617
M2M	418M	N	0.583	0.932	0.650
M2M	418M	src-mt1-mt2	0.589	0.934	0.654
en-zh					
N	base	N	0.629	0.885	0.741
N	big	N	0.625	0.885	0.738
M2M	418M	N	0.633	0.884	0.744
M2M	418M	mvppe	0.646	0.892	0.752
WMT20					
en-de best			0.597	0.935	0.662
en-zh best			0.610	0.887	0.623

Table 2: Unconstrained setting. base and big stand for the transformer-base and transformer-big architecture. 418M is the M2M-100-small model.

Ablation Configuration	MCC	F1-OK	F1-BAD
best	0.589	0.934	0.654
-LevT	0.555	0.938	0.610
-LevT +lang-adapt	0.565	0.930	0.635
-LevT -synth.	0.321	0.915	0.380
-LevT -synth. +lang-adapt	0.451	0.946	0.498
-m2m	0.562	0.939	0.617
-m2m -KD	0.526	0.933	0.589
-m2m -heuristic tag	0.551	0.936	0.610
-m2m -synth.	0.551	0.927	0.623
-m2m -synth. -heuristic tag	0.539	0.925	0.613

Table 3: Ablation analysis. All results trained with synthetic data in this table use the src-mt1-mt2 data synthesis method. +lang-adapt stands for adding an extra autoregressive MT training step using the same parallel training data as LevT training, so the M2M-100 model is adapted to translating a specific language pair.

zh language pair. This is possibly influenced by the relatively low translation quality of M2M-100 model on en-zh language pair, as pointed out by Fan et al. (2020).

With all our techniques applied, our best Target MCC result is only slightly behind the winning system on en-de language pair, while being significantly better than the winning system on en-zh language pair. Most notably, for en-zh language pair, even our smallest LevT model are able to beat the state-of-the-art. It should also be pointed out that all of our results are achieved without any model ensemble, and our pre-trained model architecture is just a transformer-big counterpart, while other participating teams deployed larger models.

To confirm that each component of our training scheme is necessary, we conducted a comprehensive ablation study on en-de language pair, shown in Table 3. The upper part of the table demonstrates that LevT training is necessary, and we do so by conducting the finetuning directly on M2M-100-small initialization. Despite the strength of the M2M-100 model as a translation model, there is still a significant performance drop without LevT Training, and more so without synthetic finetuning. To rule out the effect of bilingual knowledge introduced with LevT training, we also experimented with continue-training the M2M-100 model (+lang-adapt in Table 3) with the same parallel data used for LevT training, but the performance gap remains. On the other hand, the lower part of the table highlights the effect of various other training techniques, where we use the best system without M2M-100-small initialization as the base. We can conclude that KD is crucial for optimal performance and that finetuning with heuristic subword-level tag reference is responsible for a small but stable performance improvement.

5 Conclusion

In this work, we proposed to use Levenshtein Transformer to substitute the usual MLM-style training in the Predictor-Estimator framework as the initial training step. We also proposed a series of techniques to effectively transfer the translation knowledge to the word-level QE task, including data synthesis, heuristic subword-level reference, and incorporating pre-trained translation models. Our results demonstrate superior data efficiency under the data-constrained setting and competitive performance under the unconstrained setting. We also hope this work can inspire further exploration for other uses of Levenshtein Transformer apart from the non-autoregressive translation.

Acknowledgements

This material is based upon work supported by the United States Air Force under Contract No. FA8750-19-C-0098. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the United States Air Force and DARPA.

References

- Kevin Clark, Minh-Thang Luong, Quoc V. Le, and Christopher D. Manning. 2020. **ELECTRA: pre-training text encoders as discriminators rather than generators**. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. **Unsupervised cross-lingual representation learning at scale**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Angela Fan, Shruti Bhosale, Holger Schwenk, Zhiyi Ma, Ahmed El-Kishky, Siddharth Goyal, Man-deep Baines, Onur Celebi, Guillaume Wenzek, Vishrav Chaudhary, Naman Goyal, Tom Birch, Vitaliy Liptchinsky, Sergey Edunov, Edouard Grave, Michael Auli, and Armand Joulin. 2020. **Beyond english-centric multilingual machine translation**. *arXiv preprint*.
- Jiatao Gu, Changhan Wang, and Junbo Zhao. 2019. **Levenshtein transformer**. 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 11179–11189.
- Chi Hu, Hui Liu, Kai Feng, Chen Xu, Nuo Xu, Zefan Zhou, Shiqin Yan, Yingfeng Luo, Chenglong Wang, Xia Meng, Tong Xiao, and Jingbo Zhu. 2020. **The NiuTrans system for the WMT20 quality estimation shared task**. In *Proceedings of the Fifth Conference on Machine Translation*, pages 1018–1023, Online. Association for Computational Linguistics.
- Fabio Kepler, Jonay Trénous, Marcos Treviso, Miguel Vera, and André F. T. Martins. 2019. **OpenKiwi: An open source framework for quality estimation**. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 117–122, Florence, Italy. Association for Computational Linguistics.
- Hyun Kim, Hun-Young Jung, Hong-Seok Kwon, Jong-Hyeok Lee, and Seung-Hoon Na. 2017. **Predictor-estimator: Neural quality estimation based on target word prediction for machine translation**. *ACM Trans. Asian Low Resour. Lang. Inf. Process.*, 17(1):3:1–3:22.
- Dongjun Lee. 2020. **Two-phase cross-lingual language model fine-tuning for machine translation quality estimation**. In *Proceedings of the Fifth Conference on Machine Translation*, pages 1024–1028, Online. Association for Computational Linguistics.
- Brian W Matthews. 1975. **Comparison of the predicted and observed secondary structure of t4 phage lysozyme**. *Biochimica et Biophysica Acta (BBA)-Protein Structure*, 405(2):442–451.
- João Moura, Miguel Vera, Daan van Stigt, Fabio Kepler, and André F. T. Martins. 2020. **IST-unbabel participation in the WMT20 quality estimation shared task**. In *Proceedings of the Fifth Conference on Machine Translation*, pages 1029–1036, Online. Association for Computational Linguistics.
- 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.
- Tharindu Ranasinghe, Constantin Orasan, and Ruslan Mitkov. 2020. **TransQuest at WMT2020: Sentence-level direct assessment**. In *Proceedings of the Fifth Conference on Machine Translation*, pages 1049–1055, Online. Association for Computational Linguistics.
- Ricardo Rei, Craig Stewart, Ana C Farinha, and Alon Lavie. 2020. **COMET: A neural framework for MT evaluation**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2685–2702, Online. Association for Computational Linguistics.
- Matthew Snover, Bonnie Dorr, Richard Schwartz, Linnea Micciulla, and John Makhoul. 2006. **A study of translation edit rate with targeted human annotation**. In *Proceedings of association for machine translation in the Americas*, volume 200. Citeseer.
- Lucia Specia, Frédéric Blain, Marina Fomicheva, Erick Fonseca, Vishrav Chaudhary, Francisco Guzmán, and André F. T. Martins. 2020. **Findings of the WMT 2020 shared task on quality estimation**. In *Proceedings of the Fifth Conference on Machine Translation*, pages 743–764, Online. Association for Computational Linguistics.
- Brian Thompson and Matt Post. 2020a. **Automatic machine translation evaluation in many languages via zero-shot paraphrasing**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 90–121, Online. Association for Computational Linguistics.
- Brian Thompson and Matt Post. 2020b. **Paraphrase generation as zero-shot multilingual translation: Disentangling semantic similarity from lexical and syntactic diversity**. In *Proceedings of the Fifth Conference on Machine Translation*, pages 561–570, Online. Association for Computational Linguistics.
- Minghan Wang, Hao Yang, Hengchao Shang, Daimeng Wei, Jiaxin Guo, Lizhi Lei, Ying Qin, Shimin Tao, Shiliang Sun, Yimeng Chen, and Liangyou Li. 2020. **HW-TSC’s participation at WMT 2020 quality estimation shared task**. In *Proceedings of the Fifth Conference on Machine Translation*, pages 1056–1061, Online. Association for Computational Linguistics.

Daimeng Wei, Hengchao Shang, Zhanglin Wu, Zhengzhe Yu, Liangyou Li, Jiabin 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 *Proceedings of the Fifth Conference on Machine Translation*, pages 293–299, Online. Association for Computational Linguistics.

A Details on Edit Tag Conversion to/from Subword-Level

Algorithm 1 shows the tag conversion algorithm from subword-level to word-level (for inference). Algorithm 2 shows the tag conversion algorithm from word-level to subword-level (for finetuning).

Algorithm 1: Conversion of subword-level tags to word-level tags

Input: subword-level token sequence \mathbf{y}^{sw} ,
word-level token sequence \mathbf{y}^w ,
subword-level tag sequence \mathbf{q}^{sw}

Output: word-level tag sequence \mathbf{q}^w

```
 $\mathbf{q}^w \leftarrow [];$ 
for each word  $w_k$  in  $\mathbf{y}^w$  do
  find subword index span  $(s_k, e_k)$  in  $\mathbf{y}^{sw}$ 
  that corresponds to  $w_k$ ;
   $\mathbf{q}_k^{sw} \leftarrow$  subword-level translation and
  gap tags within span  $(s_k, e_k)$ ;
   $g_{s_k}^{sw} \leftarrow$  subword-level gap tag before
  span  $(s_k, e_k)$ ; //  $g_{s_k-1}^{sw} \in \mathbf{q}^{sw}$ 
  if  $\forall \mathbf{q}_k^{sw}$  are OK then
    |  $\mathbf{q}^w += [g_{s_k}^{sw}, \text{OK}];$ 
  else
    |  $\mathbf{q}^w += [g_{s_k}^{sw}, \text{BAD}];$ 
  end
end
 $\mathbf{q}^w += [\mathbf{q}^{sw}[-1]];$  // add ending gap tag
return  $\mathbf{q}^w;$ 
```

B Details of Experimental Setup

B.1 Data Synthesis

For all en-de experiments, we took the first 1 million line from the Europarl corpus and conduct data synthesis.

- `src-mt-ref` We translate the source side of the parallel data with the NMT system provided by the shared task organizer.
- `bt-rt-tgt` Both the back-translation and the round-trip translation are performed with M2M-100-mid (1.2B) model.
- `src-mt1-mt2` We take the source side of the parallel data and translate it with the NMT system from the shared task (weaker system, mt1) and the Facebook winning system for the WMT19 en-de news translation (Ng et al., 2019, stronger system, mt2). We remove all the cases where mt1 and mt2 are identical.

- `mvppe` The MVPPE decoding is conducted with M2M-100-mid (1.2B) model.

For en-zh experiments, we take the shared task en-zh parallel data but exclude the UN data for MVPPE data synthesis. The same multilingual translation model is used.

We also experimented with using larger synthetic data for en-de with some synthesis method, but didn't observe a significant performance difference compared to this smaller dataset.

B.2 Misc.

We preprocess our data by first tokenizing with Moses tokenizer, and then applying subword segmentation. For all the LevT models without M2M-100 initialization, we use the same BPE model and source/target-side vocabulary as the official NMT checkpoint provided by the WMT20 QE shared task. For models with M2M-100 initialization, we use the M2M-100 sentencepiece model.

Under the constrained setting, we use the NMT checkpoint supplied by the shared task to generate the knowledge distillation data for LevT translation training, both for en-de and en-zh. Under the unconstrained setting, for en-de, we use the Facebook winning system for the WMT19 en-de news translation, and for en-zh, we use our own Transformer-base en-zh model trained on WMT17 en-zh data.

All of our implementations are based on the Fairseq toolkit. We use the same hyperparameter for LevT translation model training as the document provided in Fairseq¹. For both synthetic and human post-edited data finetuning, we use Adam optimizer with a learning rate 2e-5 with warmup (4000 updates for synthetic finetuning and 2000 for human post-edited data finetuning), and we use the shared task development set to select the best checkpoint.

For all the mvppe experiments, we use $\lambda_t = 2.0$ and $\lambda_p = 1.0$, after doing a grid search over $\lambda_t = \{1.0, 2.0, 3.0\}$ and $\lambda_p = \{1.0, 1.2, 1.5\}$ with a goal to match the TER distribution of human post-editing obtained from the en-de human PE dev data.

All the word-level and subword-level tags we use as the reference for finetuning are computed using our own TER implementation², but we stick to the

¹https://github.com/pytorch/fairseq/blob/master/examples/nonautoregressive_translation/README.md

²<https://github.com/marian-nmt/moses-scorers>

original reference tags in the test set for evaluation to avoid potential result mismatch. Our evaluations are done with the official evaluation scripts³ from the shared task. The script computes the Matthews Correlation Coefficient (MCC, Matthews, 1975), which is formulated as follows:

$$S = \frac{TP + FN}{N} \quad (1)$$

$$P = \frac{TP + FP}{N} \quad (2)$$

$$MCC = \frac{TP/N - S \times P}{\sqrt{PS(1-S)(1-P)}} \quad (3)$$

where TP/FP stands for true/false positives and TN/FN stands for true/false negatives. N stands for the number of examples in the dataset. The script also computes F1 scores of the OK and BAD tags.

Table 4 shows some statistics of the data we use in our experiments.

B.3 Extra Results on the Updated Dataset

As of Sep. 2021, there is an updated version of the WMT20 shared task dataset (train/dev/test) with the same MT output but different human post-edited output. For reference of future work, we provide results on this updated dataset from some of our experiment configurations in Table 5.

Algorithm 2: Construction of heuristic subword-level tags

Input: subword-level token sequence \mathbf{y}^{sw} ,
word-level token sequence \mathbf{y}^w , naive
subword-level tag sequence \mathbf{q}^{sw} ,
word-level tag sequence \mathbf{q}^w

Output: heuristic subword-level tag
sequence $\widetilde{\mathbf{q}}^{sw}$

```

 $\widetilde{\mathbf{q}}^{sw} \leftarrow []$ ;
for each word  $w_k$  in  $\mathbf{y}^w$  do
  find subword index span  $(s_k, e_k)$  in  $\mathbf{y}^{sw}$ 
  that corresponds to  $w_k$ ;
   $\mathbf{q}_k^{sw} \leftarrow$  subword-level translation and
  gap tags within span  $(s_k, e_k)$ ;
   $t_k^w \leftarrow$  word-level translation tag for  $w_k$ ;
  //  $t_k^w \in \mathbf{q}^w$ 
   $g_k^w \leftarrow$  word-level gap tag before  $w_k$ ;
  //  $g_k^w \in \mathbf{q}^w$ 
   $\widetilde{\mathbf{q}}^{sw} += [g_k^w]$ ; // copy left gap tag
   $n = |\mathbf{q}_k^{sw}|$ ; // # subwords for  $w_k$ 
  if  $t_k^w$  is OK then
    /* word is OK, so all subwords
    and gaps between them
    should be OK */
     $\widetilde{\mathbf{q}}^{sw} += [\text{OK}] * (2n - 1)$ 
  else if  $\exists$  BAD in  $\mathbf{q}_k^{sw}$  then
    /* no conflict between
    subword-level tag and
    word-level tag - copy  $\mathbf{q}_k^{sw}$ 
    as-is */
     $\widetilde{\mathbf{q}}^{sw} += \mathbf{q}_k^{sw}$ ;
  else
    /* subword-level tag disagrees
    with word-level tag - force
    it as all-BAD to guarantee
    perfect conversion */
     $\widetilde{\mathbf{q}}^{sw} += [\text{BAD}] * (2n - 1)$ 
end
 $\widetilde{\mathbf{q}}^{sw} += [\mathbf{q}^w[-1]]$ ; // add ending gap tag
return  $\widetilde{\mathbf{q}}^{sw}$ ;

```

³<https://github.com/sheffieldnlp/qe-eval-scripts>

	# sentence pairs/triplets
shared task en-de parallel	3.96M
shared task en-zh parallel	20.3M
WMT20 en-de parallel	44.2M
en-de src-mt-ref synthetic	945K
en-de src-mt1-mt2 synthetic	808K
en-de bt-rt-tgt synthetic	998K
en-de mvppe synthetic	993K
en-zh mvppe synthetic	9.26M
shared task en-de human PE train	7K
shared task en-zh human PE train	7K
shared task en-de human PE dev	1K
shared task en-zh human PE dev	1K
shared task en-de human PE test	1K
shared task en-zh human PE test	1K

Table 4: Data Statistics for Our Experiments

	MCC	F1-OK	F1-BAD
en-de			
LevT base (constr.)	0.363	0.951	0.410
LevT base	0.434	0.954	0.479
LevT big	0.459	0.958	0.498
LevT m2m	0.489	0.955	0.533
LevT m2m + synth.	0.500	0.956	0.544
en-zh			
LevT base	0.463	0.907	0.554
LevT big	0.460	0.903	0.552
LevT m2m	0.453	0.898	0.548
LevT m2m + synth.	0.459	0.900	0.552

Table 5: Results on the updated dataset