

# Can Automatic Post-Editing Improve NMT?

**Shamil Chollampatt**  
Rakuten, Inc.  
shamil.chollampatt@rakuten.com

**Raymond Hendy Susanto**  
Rakuten, Inc.  
raymhs91@gmail.com

**Liling Tan**  
Rakuten, Inc.  
liling.tan@rakuten.com

**Ewa Szymanska**  
Rakuten, Inc.  
ewa.szymanska@rakuten.com

## Abstract

Automatic post-editing (APE) aims to improve machine translations, thereby reducing human post-editing effort. APE has had notable success when used with statistical machine translation (SMT) systems but has not been as successful over neural machine translation (NMT) systems. This has raised questions on the relevance of APE task in the current scenario. However, the training of APE models has been heavily reliant on large-scale artificial corpora combined with only limited human post-edited data. We hypothesize that APE models have been underperforming in improving NMT translations due to the lack of adequate supervision. To ascertain our hypothesis, we compile a larger corpus of human post-edits of English to German NMT. We empirically show that a state-of-art neural APE model trained on this corpus can significantly improve a strong in-domain NMT system, challenging the current understanding in the field. We further investigate the effects of varying training data sizes, using artificial training data, and domain specificity for the APE task. We release this new corpus under CC BY-NC-SA 4.0 license at <https://github.com/shamilcm/pedra>.

## 1 Introduction

Automatic Post-Editing (APE) aims to reduce manual post-editing effort by automatically fixing errors in the machine-translated output. Knight and Chander (1994) first proposed APE to cope with systematic errors in selecting appropriate articles for Japanese to English translation. Earlier application of statistical phrase-based models for APE treated it as a monolingual re-writing task without considering the source sentence (Simard et al., 2007; Béchara et al., 2011). Modern APE models take the source text and machine-translated text as input and output the post-edited text in the target language (see Figure 1).

Source text (English): <b>Will he send the gifts to the house?</b>
Machine translated text (German): <b>Die Geschenke in mein Haus schicken?</b> <i>(The gifts) (to my) (house) (send)</i>
Post-edited text (German): <b>Wird er die Geschenke ins Haus schicken?</b> <i>(Will he) (the gifts) (to the) (house) (send)</i>

Figure 1: An example of post-editing given the source text in English and the translated text in German.

APE models are usually trained and evaluated in a *black-box* scenario where the underlying MT model and the decoding process are inaccessible making it difficult to improve the MT system directly. APE can be effective in this case to improve the MT output or to adapt its style or domain.

Recent advancement of APE has shown remarkable success on statistical machine translation (SMT) outputs (Junczys-Dowmunt and Grundkiewicz, 2018; Correia and Martins, 2019) even when trained with limited number of post-edited training instances (generally “triplets” consisting of *source*, *translated*, and *post-edited* segments), with or without additional large-scale artificial data (Junczys-Dowmunt and Grundkiewicz, 2016; Negri et al., 2018). Substantial improvements have been reported especially on English-German (EN-DE) WMT APE shared tasks on SMT (Bojar et al., 2017; Chatterjee et al., 2018), when models were trained with fewer than 25,000 human post-edited triplets. However, on NMT, strong APE models have failed to show any notable improvement (Chatterjee et al., 2018, 2019; Ive et al., 2020) when trained on similar-sized human post-edited data. This has led to questions regarding the usefulness of APE with current NMT systems that produce improved trans-

lations compared to SMT. Junczys-Dowmunt and Grundkiewicz (2018) concluded that the results of the WMT’18 APE (NMT) task “might constitute the end of neural automatic post-editing for strong neural in-domain systems” and that “neural-on-neural APE might not actually be useful”. Contrary to this belief, we hypothesize that a competitive neural APE model still has potential to further improve strong state-of-the-art in-domain NMT systems when trained on adequate human post-edited data.

We compile a new large post-edited corpus, *SubEdits*, which consists of actual human post-edits of translations of drama and movie subtitles produced by a strong in-domain proprietary NMT system. We use this corpus to train a state-of-the-art neural APE model (Correia and Martins, 2019), with the goal of answering the following three research questions to better assess the relevance of APE going forward:

- Can APE substantially improve in-domain NMT with adequate data size?
- How much does artificial APE data help?
- How significant is domain shift for APE?

**Spoilers:** Through automatic and human evaluation, we confirm our hypothesis that, in order to notably improve over the original NMT output (“do-nothing” baseline), state-of-the-art APE models need to be trained on a larger number of human post-edits, unlike the case with SMT. Training on datasets of sizes in the scale of those from the WMT APE tasks, even with large-scale in-domain artificial APE corpora, leads to underperformance. Our experimental results also highlight that APE models are highly sensitive to domain differences. To effectively exploit out-of-domain post-edited corpora such as *SubEdits* in other domains, it has to be carefully mixed with available in-domain data.

## 2 SubEdits Corpus

Human post-edited corpora of NMT outputs from previous WMT APE shared tasks usually consist of fewer than 25,000 instances. Large-scale artificial corpora such as eSCAPE (Negri et al., 2018), do not adequately cater to the primary APE objective of correcting systematic errors of the MT outputs since the pseudo “post-edits” are independent human-translated references often differing

	Lang.	Size	Domain
<b>Human post-edited corpora</b>			
QT21 (Specia et al., 2017)	EN-LV	21K	Life Sciences
WMT’18 & ’19 APE (Chatterjee et al., 2018)	EN-DE	15K	IT
WMT’19 APE (Chatterjee et al., 2019)	EN-RU	17K	IT
APE-QUEST (Ive et al., 2020)	EN-NL EN-FR EN-PT	11K 10K 10K	Legal
SubEdits ( <i>this work</i> )	EN-DE	161K	Subtitles
<b>Artificial corpora</b>			
eSCAPE (Negri et al., 2018)	EN-DE EN-IT EN-RU	7.2M 3.3M 7.7M	Mixed

Table 1: APE corpora on NMT outputs and their sizes in terms of number of post-edited triplets.

greatly from the MT output. Table 1 lists the real and artificial APE corpora on NMT outputs. Due to the paucity of larger human post-edited corpora on NMT outputs, a study of APE performance under sufficient supervised training data conditions was not possible previously. To enable such a study, we introduce the *SubEdits* EN-DE post-editing corpus with over 161K triplets of source sentences, NMT translations, and human post-edits of NMT translations.

### 2.1 Corpus Collection

*SubEdits* corpus is collected from a database of subtitles of a popular video streaming platform, Rakuten Viki (<https://www.viki.com/>) Every subtitle segment had been originally manually transcribed and translated to English before translating it to German using a proprietary NMT system employed by the platform and specialized at translating subtitles. Viki community<sup>1</sup> members who volunteer as subtitle translators would then post-edit the machine-translated subtitles to further improve it, if necessary.

### 2.2 Corpus Filtering

We use an adaptation of Gale-Church filtering (Tan and Pal, 2014) used for machine translation for filtering the triplets. The global character mean ratio  $r_c$  is computed as the ratio between the number of characters in the source and machine translated portions of the entire corpus. We remove triplets ( $src$ ,  $mt$ ,  $pe$ ) from the corpus where the ratio between the number of characters of source ( $src$ ) and

<sup>1</sup><https://contribute.viki.com/>

	No. of triplets	No. of tokens		
		src	mt	pe
Train	141,413	1,432,247	1,395,211	1,423,257
Dev	10,000	101,330	98,581	100,795
Test	10,000	101,709	99,032	101,112

Table 2: Statistics of the SubEdits corpus

post-edit (*pe*) does not lie within a threshold range of  $(1 - t)r_c$  and  $(1 + t)r_c$  with  $t = 0.2$ . We normalize punctuation<sup>2</sup> and remove duplicate triplets. Among the triplets that share the same *src* and *mt* segments, we choose only the one with the longest *pe*. Finally, we remove triplets that are not correctly identified with the respective source and target language using a language identification tool<sup>3</sup> (Lui and Baldwin, 2012). We set aside 10,000 triplets as development set and 10,000 triplets as test set. The final statistics are shown in Table 2.

### 3 BERT Encoder-Decoder APE Model

BERT Encoder-Decoder APE (Correia and Martins, 2019) is a state-of-the-art neural APE model based on a Transformer model (Vaswani et al., 2017) with the encoder and decoder initialized with pre-trained multilingual BERT (Devlin et al., 2019) weights and fine-tuned on post-editing data.

A single encoder is used to encode both the source text and the machine-translated text by concatenating them with the separator token [SEP]. The encoder component of the model is identical to the original Transformer encoder initialized with pre-trained weights from the multilingual BERT. For the decoder, Correia and Martins (2019) initialized the context attention weights with the corresponding BERT self-attention weights. Also, the weights of the self-attention layers of the encoder and decoder are tied. All other weights are initialized with corresponding weights from the same multilingual BERT model as well.

BERT Encoder-Decoder APE was shown to outperform other state-of-the-art APE models (Tebbifakhr et al., 2018; Junczys-Dowmunt and Grundkiewicz, 2018) on SMT outputs even in the absence of additional large-scale artificial data that competing models have used. An improved variant of this model with additional in-domain artificial data, despite being the winning submission of the recent WMT’19 APE EN-DE (NMT) task

<sup>2</sup>Using Moses `normalize-punctuation.perl` script.

<sup>3</sup><https://github.com/saffsd/langid.py>

(Lopes et al., 2019), only performed marginally better than the baseline NMT output. For the purpose of this study, we base our experiments on the BERT Encoder-Decoder APE architecture (Correia and Martins, 2019).

## 4 Experimental Setup

### 4.1 Model Hyperparameters

For the BERT Encoder-Decoder model (BERT Enc-Dec), we use the implementation<sup>4</sup> and model hyperparameters used by Correia and Martins (2019) and initialize the encoder and decoder with cased multilingual BERT (base) from Transformers<sup>5</sup> library (Wolf et al., 2019). The encoder and decoder follow the architecture of BERT (base) with 12 layers and 12 attention heads, an embedding size of 768, and a feed-forward layer size of 3072. We set the effective batch size to 4096 tokens for parameter updates. We train BERT Enc-Dec on a single NVIDIA Quadro RTX6000 GPU. Training on our SubEdits corpus took approximately 5 hours to converge. We validate and save checkpoints at every 2000 steps and use early-stopping (patience of 4 checkpoints) to select the model based on best perplexity. We use a decoding beam size of 5.

As a control measure, we compare BERT Enc-Dec against two vanilla Transformer APE models using automatic metrics. The Transformer APE models use BERT vocabularies and tokenization, and employ a single encoder to encode the concatenation *src* and *mt*, but they are not initialized with pre-trained weights. The following are the descriptions of the two Transformer APE baselines:

**TF (base)** A Transformer (base) (Vaswani et al., 2017) model with 6 hidden layers implemented in OpenNMT-py.<sup>6</sup> The embedding size is 512 with 2048 feed-forward units. We use default learning parameters in OpenNMT-py: Adam optimizer with a learning rate of 2 and Noam scheduler.

**TF (BERT size.)** A bigger Transformer with the same number of layers, attention heads, embedding dimensions, hidden, and feed-forward dimensions as BERT Enc-Dec, but without any pre-training and tying of self-attention layers. All learning hyperparameters follow that of TF (base) model.

<sup>4</sup><https://github.com/deep-spin/OpenNMT-APE>

<sup>5</sup><https://github.com/huggingface/transformers>

<sup>6</sup><https://github.com/OpenNMT/OpenNMT-py>

## 4.2 Pre-processing and Post-processing

SubEdits corpus contains HTML tags such as line breaks (`<br>`) and italic tags (`<i>`), and symbols denoting musical notes ( $\♪$ ,  $\♩$ ) and segments often begin with hyphens (-). We applied several processing steps to make the data as close as possible to natural sentences on which BERT has been pre-trained on. The triplets with multi-line *src*, *mt*, and *pe* containing `<br>` tags are split into separate training instances<sup>7</sup> and we remove italics and other HTML tags, musical note symbols, and leading hyphens. Thereafter, the input is tokenized with the BERT tokenization and word-piece segmentation in the Transformers library. During test time, we keep track of the changes made to input such as deletion of leading hyphens, music symbols, and italics tags, and splitting at `<br>` tags. After decoding, the outputs are detokenized and post-processed to re-introduce the tracked changes and evaluated.

## 4.3 Evaluation

We evaluate the models using three different automatic metrics: BLEU (Papineni et al., 2002), ChrF (Popović, 2015), and TER (Snover et al., 2006). For our evaluation on SubEdits test set, differing from WMT APE task evaluation, we post-process and detokenize the outputs and use SacreBLEU<sup>8</sup> (Post, 2018) to evaluate BLEU and ChrF, and TERCOM<sup>9</sup> to compute TER with normalization. Significance test is done by bootstrap re-sampling on BLEU with 1000 samples (Koehn, 2004). Additionally, we conduct human evaluation to ascertain the improvement of the BERT Enc-Dec APE model and to determine the human upper-bound performance for the SubEdits benchmark (see Section 5.3).

We also compare the APE model on the canonical WMT APE dataset (Section 5.6 and Table 7). We follow their evaluation method and use the released tokenized post-edited reference to compute BLEU, ChrF, and TER on the tokenized output.

# 5 Results and Discussion

## 5.1 Proprietary In-domain NMT

We first assess the quality of an proprietary in-domain NMT system that is used for compiling

<sup>7</sup>We only separate at `<br>` when the *src*, *mt*, and *pe* contains same number of `<br>` symbols.

<sup>8</sup><https://github.com/mjpost/sacreBLEU>

<sup>9</sup><http://www.cs.umd.edu/~snover/tercom/>

	BLEU↑	ChrF↑	TER↓
Proprietary NMT	46.83	63.81	37.20
Google Translate	40.96	59.20	41.91
Microsoft Translator	38.78	57.68	43.72
SYSSTRAN	38.06	56.74	44.37

Table 3: Comparison of the proprietary NMT to leading commercial MT systems on an in-domain test set.

the SubEdits corpus. We use it as a black-box system and use the evaluation results from Table 3 to demonstrate that it is a strong baseline for studying APE performance on NMT outputs.

We compare the proprietary NMT system to three leading commercial EN-DE NMT systems: Google Translate, Microsoft Translator, and SYSSTRAN, on a separate in-domain EN-DE test set of 5,136 subtitle segments with independent reference translations (i.e., not post-edits of any system) fetched from the same video streaming platform as the SubEdits corpus. The results (as of May 2020) are summarized in Table 3. Unsurprisingly, the proprietary NMT system specialized at translating drama subtitles substantially outperforms other general MT systems.

## 5.2 APE Performance on SubEdits

Table 4 reports the performance of vanilla transformer and BERT Enc-Dec APE models and compares it the *do-nothing* NMT baseline (the output produced by the proprietary in-domain NMT system). TF (base) APE improves over the *do-nothing* NMT baseline output ( $p < 0.05$ ), particularly on TER scores. However, TF (BERT size) APE shows a smaller improvement on ChrF and TER scores and a drop in BLEU. Even with the SubEdits corpus, large networks such as TF (BERT size) tends to overfit. However, with pre-trained BERT initialization, BERT Enc-Dec APE shows substantial improvement across all metrics. Unlike previous studies that report marginal improvements (Chatterjee et al., 2018, 2019), our results show that a strong APE model trained on large human post-edits can significantly outperform ( $p < 0.001$ ) a strong in-domain NMT system.

## 5.3 Human Evaluation

To validate the improvement in automatic evaluation scores and to estimate the human upper-bound performance on SubEdits, we conducted human evaluation. We hired five German native freelance translators who are also proficient in English and

	No. of Params	Dev			Test		
		BLEU↑	ChrF↑	TER↓	BLEU↑	ChrF↑	TER↓
<i>do-nothing</i> NMT		62.07	71.66	27.68	61.88	71.33	28.06
w/ TF (Base) APE	105.5M	62.47	72.26	25.65	62.26	71.97	25.94
w/ TF (BERT size.) APE	290.4M	62.04	72.04	25.73	61.62	71.65	26.14
w/ BERT Enc-Dec APE	262.4M	<b>64.88</b>	<b>74.94</b>	<b>23.29</b>	<b>64.53</b>	<b>74.71</b>	<b>23.72</b>

Table 4: Performance of APE models on the SubEdits test set.

Score each of the three German translations for the original English text from 1 to 5 based on how adequately they express the meaning of the original English text.

Original English Text:  
I hope you all can understand.

German Translation A:  
Ich hoffe, du kannst verstehen,

German Translation B:  
Ich hoffe, ihr alle könnt das verstehen.

German Translation C:  
Ich hoffe, dass ihr alle verstehen könnt.

---

**Scores:**  
1 = Complete nonsense | 2 = Very little meaning is captured | 3 = Some meaning is captured | 4 = Most of the meaning is captured | 5 = Perfect.  
(You may ignore minor grammatical issues that do not affect the meaning)

Translation A \*  
 1  2  3  4  5  I can't decide

Translation B \*  
 1  2  3  4  5  I can't decide

Translation C \*  
 1  2  3  4  5  I can't decide

Figure 2: Interface used to rate the translations.

had prior experience with English/German translation.

Given the original English text, the annotators were asked to rate the adequacy (from 1 to 5) for three German translations: (1) the *do-nothing* baseline output (NMT), (2) BERT Enc-Dec APE output (APE), and (3) the human post-edited text (Human). Figure 2 shows the interface presented to the annotators for rating the translations. The three translations are presented on the same screen in random order and the annotators are unaware of their origin.

Following recent WMT APE tasks (Bojar et al., 2017; Chatterjee et al., 2018, 2019), our human evaluation is also based solely on adequacy assessments. Previous studies reported a high correlation of fluency judgments with adequacy (Callison-Burch et al., 2007) making the fluency annotations

Annotator	NMT	APE	Human	# Eval.
A	3.7	4.2	4.5	593 / 603
B	3.5	4.0	4.4	594 / 603
C	3.7	4.3	4.4	603 / 603
D	2.8	3.4	3.8	587 / 603
E	3.3	3.8	4.3	602 / 603
A-E	3.4	3.9	4.3	2979 / 3015

Table 5: Average adequacy scores (1-5) rated by annotators (A to E). Overall average is shown in the last row (A-E).

superfluous (Przybocki et al., 2009). Unlike the recent WMT APE tasks, we did not opt for direct assessments (Graham et al., 2013) since we wanted to evaluate the degradation or improvement in the quality of the NMT output due to APE and human post-edits on the same English source segments.

We elicit judgments for all test set instances where the APE model modified the NMT output beyond simple edits on punctuation, HTML tags, spacing, or casing. 2,815 out of the 10,000 instances in our test set contains non-simple edits. A set of 50 instances out of 2,815 was evaluated by all annotators to compute inter-annotator agreement.<sup>10</sup>

After evaluation, we filtered out the instances where the annotator was unable to decide a score for any of the three translations. The average scores by each annotator (A to E) and the overall average scores are shown in Table 5. The numerator of the “# Eval.” column indicates the number of evaluations used for the average score computation after filtering out the “*I can't decide*” annotations. The results of our human evaluation (Table 5) show that all five annotators rate the APE output better than baseline NMT output by at least +0.5 on average, reaching an overall score of 3.9. All the five annotators rated the human post-edited output substantially better than the NMT output and the APE output, which indicates that quality of the post-edits in the SubEdits corpus is high. Human post-edits received an overall average score of 4.3.

<sup>10</sup>Each annotator scored 603 test instances.

Using the repeated set of 46 instances,<sup>11</sup> we compute inter-annotator agreement using average pairwise Cohen’s Kappa  $\kappa$  (Cohen, 1960) to be 0.27 which is considered to be fair (Landis and Koch, 1977) and similar to that observed for adequacy judgments in WMT tasks (Callison-Burch et al., 2007). However, the ranges of scores used by the annotators differ considerably (especially, annotator ‘D’). Hence, measures such as a weighted Kappa  $\kappa_w$  (Cohen, 1968), which assigns partial credit to smaller disagreements and works better with ordinal data (such as our adequacy judgments), is more suitable. We compute the average pairwise quadratically weighted Kappa  $\kappa_w$  to be 0.50, and consider their agreement to be moderate.

#### 5.4 Can APE substantially improve in-domain NMT with adequate data size?

To analyze the effect of training data size with respect to APE performance, we train BERT Enc-Dec APE with varying sizes of training data from the SubEdits corpus and evaluated the models on the SubEdits development set. For each training data size, ranging from 6,250 to 125,000, we train three models on three random samples of the respective size from the SubEdits training set. Each point in Figure 3 denotes the *mean* score of the three models (the vertical error bars at each point denote the *minimum* and *maximum* scores). The *do-nothing* NMT baseline score is represented by a *horizontal* dotted line. As a reference, we mark the size equivalent to that of WMT’18 APE EN-DE (NMT) training set (13,441 triplets) with the *vertical* dotted line. The rightmost point on each graph represents the score if the full training corpus is used.

Although the sizes of WMT APE dataset and the SubEdits corpus are not directly comparable, we see that size does matter for better APE performance. When the APE model was trained on a subset of SubEdits corpus that is of the same size as the WMT’18 APE training data, it performs worse than the baseline in terms of BLEU score and only marginally improves in ChrF and TER scores (see intersection points of the vertical and horizontal lines in Figure 3).

Interestingly, doubling the amount of training data from 12,500 to 25,000 provides slight BLEU

<sup>11</sup>We removed 4 instances out of the 50, where one or more annotators chose the “I can’t decide” option.

gains above the *do-nothing* baseline and increasing the data size to 50,000 training instances improves the model further by +1 BLEU. The curves continue to show an increasing trend. After 100,000 training instances, the data size effect on score improvement slows down. This experiment shows the possibility that previous work on APE for NMT outputs might have reached a plateau simply due to the lack of human post-edited data rather than the limited usefulness of APE models.

#### 5.5 How much does artificial APE data help?

Previous work using strong neural APE models (Junczys-Dowmunt and Grundkiewicz, 2018; Tebbifakhr et al., 2018) relied predominantly on artificial corpora such as that released by Junczys-Dowmunt and Grundkiewicz (2016) and the eSCAPE corpora (Negri et al., 2018). However, artificial post-edits are either generated from monolingual corpora or independent reference translations and they do not directly address the errors made by the MT system that is to be fixed by APE.

We compare the APE model performance when trained on large-scale in-domain and out-of-domain artificial data (in the order of millions of triplets) to training on the human post-edited SubEdits corpus (over 141K human post-edits). As out-of-domain artificial data, we use the eSCAPE EN-DE NMT corpus and filter sentences that have between 0 and 200 characters resulting in 5.3 million triplets. As in-domain artificial data, we generated an artificial APE corpus using the same approach used to create the eSCAPE corpus by decoding the source sentences from the OpenSubtitles2016 parallel corpus (Lison and Tiedemann, 2016), which is also from the subtitle domain<sup>12</sup> using the same proprietary NMT system we use to create the SubEdits corpus; the corresponding references translations become the artificial post-edits. We use the same filtering criteria and pre-processing methods for SubEdits (Section 2.2 and 4.2) resulting in 5.6 million artificial triplets. We set aside 10,000 triplets from each artificial corpus and use it as a development set when training solely on the corresponding corpus. We refer to this artificial corpus as SubEscape.

We compare the performance of the BERT Enc-Dec APE trained on SubEdits corpus to that when

<sup>12</sup>Although both SubEdits and SubEscape are from the subtitle domain, the translations in SubEscape are from [www.opensubtitles.org/](http://www.opensubtitles.org/) whereas the SubEdits post-edits are compiled from Rakuten Viki.




Figure 3: Performance of BERT Enc-Dec APE model with varying training data size in terms of BLEU, ChrF, and TER metrics on the SubEdits dev set. The vertical dotted line in each figure shows the data size used for WMT APE EN-DE (NMT) task (13,441 triplets) and the horizontal dotted line shows the NMT Baseline results.

	BLEU $\uparrow$	ChrF $\uparrow$	TER $\downarrow$
<i>do-nothing</i> NMT	61.88	71.33	28.06
<b>w/ BERT Enc-Dec APE trained on:</b>			
SubEdits (R)	64.53	74.71	23.72
eSCAPE (A)	52.35	65.65	31.95
SubEscape (A)	50.51	65.89	32.78
+ SubEdits 10 $\times$ (A+R)	<b>64.59</b>	<b>75.09</b>	<b>23.41</b>

Table 6: APE performance on SubEdits test set when trained with real (R) and artificial (A) training corpora.

trained on the artificial corpora in Table 6. We find that training on artificial corpora alone, irrespective of their domain, cannot improve over the *do-nothing* baseline and in fact, degrades the performance substantially. However, when we combine SubEscape with up-sampled (10 $\times$ ) SubEdits corpus, we get a small improvement, particularly in terms of ChrF and TER.

## 5.6 How significant is domain shift for APE?

While NMT performance has been known to be particularly domain-dependant (Chu and Wang, 2018), domain shift between NMT and APE training has not been investigated previously. To assess this, we evaluate BERT Enc-Dec APE on the canonical WMT’18 APE EN-DE (NMT) dataset.<sup>13</sup> The baseline NMT system and datasets used for the WMT’18 task is from the Information Technology (IT) domain and is notably different from the domain of SubEdits. We experiment with different methods of combining SubEdits (out-domain) with the WMT APE training data (in-domain). For all experiments, we use 1,000 instances held out from the WMT’18 APE training

<sup>13</sup>WMT’19 APE task also used the same dataset for benchmarking EN-DE APE systems

	BLEU $\uparrow$	ChrF $\uparrow$	TER $\downarrow$
<i>do-nothing</i> NMT	74.73	85.89	16.84
<b>w/ BERT Enc-Dec APE trained on:</b>			
WMT’18 APE (I)	75.08	85.81	16.88
SubEdits (O)	49.05	69.48	39.30
+WMT’18 APE (O+I)	74.93	85.90	16.92
+WMT’18 APE 10 $\times$ (O+I)	<b>75.27</b>	<b>86.08</b>	<b>16.62</b>

Table 7: APE performance with in (I) and out-of-domain (O) training data on WMT APE NMT test set.

data as the validation set. The results are reported in Table 7. When trained on SubEdits alone, despite its size, we see that there is a drastic drop in performance compared to training the much smaller WMT APE data alone. When we combine SubEdits with 10 $\times$  upsampled WMT APE training data, we observe some improvement, particularly in terms of BLEU ( $p < 0.05$ ), over training with WMT APE data alone. These results show that in-domain training data is crucial to training APE models to improve in-domain NMT.

## 6 Analysis

### 6.1 Impact of APE with varying NMT quality

To study the impact of APE with varying quality of NMT output, we conduct analysis on subsets of our development set with varying translation qualities (Figure 4). We split the SubEdits development set into 10 subsets by aggregating those triplets with the NMT output scoring  $> 90$  TER (lowest quality),  $90 - 81$  TER,  $\dots$ ,  $20 - 11$  TER, and  $\leq 10$  (highest quality). They are ordered from left to right in the  $x$ -axis in Figure 4 according to increasing MT quality.  $y$ -axis denotes the difference ( $\Delta$ ) between the TER score of APE output and NMT output for each subset.

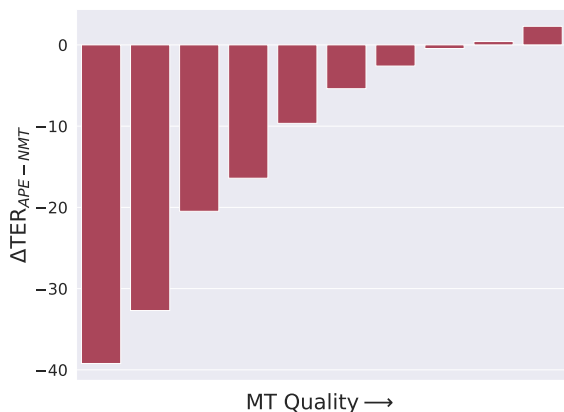


Figure 4: Translation quality difference due to APE ( $y$ -axis) shown by the  $\Delta\text{TER}_{\text{APE-NMT}}$  with increasing MT quality ( $x$ -axis). Negative  $\Delta\text{TER}$  indicates improvement in performance.

The more negative  $\Delta\text{TER}$  indicates a larger improvement due to APE. We find that on the lower quality subsets, APE improves over NMT substantially. This improvement margin reduces with improving NMT quality and can deteriorate the NMT output when NMT quality is at the highest. This experiment shows that APE contributes to improving overall NMT performance by predominantly fixing poorer quality NMT outputs. The APE model’s error will dominate and APE can become counter-productive when NMT output is nearly perfect (i.e., when there are very few or no post-edits done on them as indicated by sentence-level TER scores of  $< 10$ ). APE task remains relevant until NMT systems achieve this state, which is still not the case even for strong in-domain NMT systems as indicated by our experiments.

## 6.2 Qualitative Analysis

We qualitatively analyze the output produced by APE on the SubEdits development set to better understand the improvements and errors made by the APE model. Table 8 shows three example outputs produced by the APE model along with the original English text (SRC), the *do-nothing* baseline output (NMT), and the human post-edits (Human).

APE is able to fix incorrect named-entity translations made by the NMT system. Example 1 demonstrates an example (“Zhongyuan Palast”→“Palast Zhongcui”) where the incorrect entity is corrected by the APE model to match the human post-edits.

NMT often under-translates and misses phrases

<i>Example 1: Incorrect named entities</i>	
SRC	Go to <b>Zhongcui Palace!</b>
NMT	Geh zum <b>Zhongyuan Palast!</b>
APE	Geh zum <b>Palast Zhongcui!</b>
Human	Geht zum <b>Palast Zhongcui!</b>
<i>Example 2: Missing phrases</i>	
SRC	Let’s go back <b>to the resort</b> and we’ll talk it out.
NMT	Geh zurck und wir werden reden.
APE	Geh zurck <b>zum Resort</b> und wir werden reden.
Human	Lass uns zurck <b>zum Resort</b> gehen und darber reden.
<i>Example 3: Requires more context</i>	
SRC	Before coming, City Master negotiated with me.
NMT	Bevor <b>er</b> gekommen ist, hat der Stadtmeister ml cit mir verhandelt.
APE	Bevor <b>wir</b> kommen, hat die Stadtmeisterin mit mir verhandelt.
Human	Bevor <b>ich</b> kam, hat die Stadtmeisterin mit mir verhandelt.

Table 8: Examples where the APE model proposes changes to the NMT output on the SubEdits test set. The original sentence in English (SRC) and the human post-edit (Human) is also shown.

and the APE models usually can patch these under-translations, e.g. Example 2 where the prepositional phrase “to the resort”→“zum Resort” was missing in the MT outputs and the APE model was able to mend the translation.

As much as sentence-level APE works well empirically, the lack of context results in erroneous translation by the NMT system where it tries to infer a wrong pronoun and the APE model attempts to assume yet another wrong pronoun, e.g. translating a pronoun-dropped source text in Example 3. Often, the prior or future context from video, audio, or other subtitle instances is necessary to fill these contextual gaps. Sentence-level APE cannot address these issues robustly, which calls for further research on multimodal (Deena et al., 2017; Caglayan et al., 2019) and document-level (Hardmeier et al., 2015; Voita et al., 2019) translation and post-editing, especially for subtitles.

## 7 Related Work

Until 2018, APE models were benchmarked on SMT outputs through various WMT APE tasks (Bojar et al., 2015, 2016, 2017). The scale of post-edited data provided by these tasks was in the order of 10,000 to 25,000 triplets. The largest collection of human post-edits, released by Zhechev (2012), however, was on SMT and consisted of 30,000 to 410,000 triplets across 12 language pairs. On SMT output, participating systems showed impressive gains even with small



training datasets from WMT APE tasks (Junczys-Dowmunt and Junczys-Dowmunt, 2017; Tebbifakhr et al., 2018). The results of subsequent APE (NMT) tasks were not as promising with only marginal improvements on English-German and no improvement on English-Russian (Chatterjee et al., 2019).

Previously, there was no study to assess the necessity of larger human post-edited training data on APE performance on NMT outputs which we address in this paper. APE models were predominantly trained on large-scale artificial data combined with a few thousand human post-edits. Junczys-Dowmunt and Grundkiewicz (2016) proposed generation of large-scale artificial APE training data via round-trip translation approach inspired from *back-translation* (Sennrich et al., 2016). They combined artificial training data with real data provided by WMT APE tasks to train their model. Using a similar approach of generating artificial APE data, Freitag et al. (2019) trained a monolingual re-writing APE model trained on the generated artificial training data alone. Contrary to the round-trip translation approach, large-scale artificial APE data was generated by simply translating source sentences using NMT and SMT systems and using the reference translations as the “pseudo” post-edits to create eSCAPE corpus (Negri et al., 2018). Using the eSCAPE English-Italian APE corpus, Negri et al. (2017) assessed the performance of an online APE model in a simulated environment where the APE model is updated at test time with new user inputs. They found that their online APE models trained on eSCAPE found it difficult to improve specialized in-domain NMT systems.

Such an analysis by training on artificial corpora may not adequately assess the actual potential of APE since these corpora do not fully cater to the task and can be noisy. The “synthetic” post-edits are independent or loosely coupled with the MT outputs, and are often drastically different from the MT output. This makes analyzing APE performance over competitive NMT systems on actual post-edited data an important step in understanding the potential of APE research. Contrary to previous conclusions, our analysis shows that a competitive in-domain NMT system can be markedly improved by a strong neural APE model when trained on sufficient human post-edited training data.

## 8 Conclusion

APE has been an effective option to fix systematic MT errors and improve translations from black-box MT services. However, on NMT outputs, APE has shown hardly any improvement since training has been done on limited human post-edited data. The newly collected SubEdits corpus is the largest corpus of NMT human post-edits collected so far. We reassessed the usefulness of APE on NMT using this corpus.

We showed that with a larger human post-edited corpus, a strong neural APE model can substantially improve a strong in-domain NMT system. While artificial APE corpora help, we showed that the APE model performs better when trained on adequate human post-edited data (SubEdits) compared to large-scale artificial corpora. Finally, our experiments comparing in and out-domain APE show that domain-specificity of training affects APE performance drastically and a combination of in and out-of-domain data with certain upscaling alleviates the domain-shift problem for APE. We find that APE mostly contributes to improving NMT performance by fixing the poorer-quality outputs that still exist with strong in-domain NMT systems. We release the post-editing datasets used in this paper (SubEscape and SubEdits) along with pre/post-processing scripts at PEDRa GitHub repository (<https://github.com/shamilcm/pedra>)

## Acknowledgements

We thank the anonymous reviewers for their useful comments. We also thank Rakuten Viki community members who had contributed subtitle post-edits that helped building the SubEdits dataset.

## References

- Hanna Béchara, Yanjun Ma, and Josef van Genabith. 2011. *Statistical post-editing for a statistical MT system*. In *Proceedings of the 13th Machine Translation Summit*.
- Ondřej Bojar, Rajen Chatterjee, Christian Federmann, Yvette Graham, Barry Haddow, Shujian Huang, Matthias Huck, Philipp Koehn, Qun Liu, Varvara Logacheva, Christof Monz, Matteo Negri, Matt Post, Raphael Rubino, Lucia Specia, and Marco Turchi. 2017. *Findings of the 2017 conference on machine translation (WMT17)*. In *Proceedings of the Second Conference on Machine Translation*.

- Ondřej Bojar, Rajen Chatterjee, Christian Federmann, Yvette Graham, Barry Haddow, Matthias Huck, Antonio Jimeno Yebes, Philipp Koehn, Varvara Logacheva, Christof Monz, Matteo Negri, Aurélie Névél, Mariana Neves, Martin Popel, Matt Post, Raphael Rubino, Carolina Scarton, Lucia Specia, Marco Turchi, Karin Verspoor, and Marcos Zampieri. 2016. [Findings of the 2016 conference on machine translation](#). In *Proceedings of the First Conference on Machine Translation: Volume 2, Shared Task Papers*.
- Ondřej Bojar, Rajen Chatterjee, Christian Federmann, Barry Haddow, Matthias Huck, Chris Hokamp, Philipp Koehn, Varvara Logacheva, Christof Monz, Matteo Negri, Matt Post, Carolina Scarton, Lucia Specia, and Marco Turchi. 2015. [Findings of the 2015 Workshop on Statistical Machine Translation](#). In *Proceedings of the 10th Workshop on Statistical Machine Translation*.
- Ozan Caglayan, Pranava Madhyastha, Lucia Specia, and Loïc Barrault. 2019. [Probing the need for visual context in multimodal machine translation](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*.
- Chris Callison-Burch, Cameron Fordyce, Philipp Koehn, Christof Monz, and Josh Schroeder. 2007. [\(meta-\) evaluation of machine translation](#). In *Proceedings of the Second Workshop on Statistical Machine Translation*.
- 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: Shared Task Papers*.
- Rajen Chatterjee, Matteo Negri, Raphael Rubino, and Marco Turchi. 2018. [Findings of the WMT 2018 shared task on automatic post-editing](#). In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*.
- Chenhui Chu and Rui Wang. 2018. [A survey of domain adaptation for neural machine translation](#). In *Proceedings of the 27th International Conference on Computational Linguistics*.
- Jacob Cohen. 1960. A coefficient of agreement for nominal scales. *Educational and psychological measurement*, 20(1):37–46.
- Jacob Cohen. 1968. Weighted Kappa: Nominal scale agreement provision for scaled disagreement or partial credit. *Psychological bulletin*, 70(4):213.
- Gonçalo M. Correia and André F. T. Martins. 2019. [A simple and effective approach to automatic post-editing with transfer learning](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*.
- Salil Deena, Raymond WM Ng, Pranava Madhyastha, Lucia Specia, and Thomas Hain. 2017. Exploring the use of acoustic embeddings in neural machine translation. In *Proceedings of the 2017 IEEE Automatic Speech Recognition and Understanding Workshop*, pages 450–457.
- 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)*.
- Markus Freitag, Isaac Caswell, and Scott Roy. 2019. [APE at scale and its implications on MT evaluation biases](#). In *Proceedings of the Fourth Conference on Machine Translation (Volume 1: Research Papers)*.
- Yvette Graham, Timothy Baldwin, Alistair Moffat, and Justin Zobel. 2013. [Continuous measurement scales in human evaluation of machine translation](#). In *Proceedings of the 7th Linguistic Annotation Workshop and Interoperability with Discourse*.
- Christian Hardmeier, Preslav Nakov, Sara Stymne, Jörg Tiedemann, Yannick Versley, and Mauro Cettolo. 2015. Pronoun-focused mt and cross-lingual pronoun prediction: Findings of the 2015 discomt shared task on pronoun translation. In *Proceedings of the Second Workshop on Discourse in Machine Translation (DiscoMT)*.
- Julia Ive, Lucia Specia, Sara Szoc, Tom Vanallemeersch, Joachim Van den Bogaert, Eduardo Farah, Christine Maroti, Artur Ventura, and Maxim Khalilov. 2020. [A post-editing dataset in the legal domain: Do we underestimate neural machine translation quality?](#) In *Proceedings of The 12th Language Resources and Evaluation Conference*.
- 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*.
- Marcin Junczys-Dowmunt and Roman Grundkiewicz. 2018. [MS-UEdin submission to the WMT2018 APE shared task: Dual-source transformer for automatic post-editing](#). In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*.
- Marcin Junczys-Dowmunt and Marcin Junczys-Dowmunt. 2017. [The AMU-UEdin submission to the WMT 2017 shared task on automatic post-editing](#). In *Proceedings of the Second Conference on Machine Translation*.
- Kevin Knight and Ishwar Chander. 1994. [Automated postediting of documents](#). In *Proceedings of the 12th AAAI National Conference on Artificial Intelligence*.

- Philipp Koehn. 2004. [Statistical significance tests for machine translation evaluation](#). In *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing*.
- J. Richard Landis and Gary G. Koch. 1977. The measurement of observer agreement for categorical data. *Biometrics*, 33(1):159–174.
- Pierre Lison and Jörg Tiedemann. 2016. [OpenSubtitles2016: Extracting large parallel corpora from movie and TV subtitles](#). In *Proceedings of the 10th International Conference on Language Resources and Evaluation*.
- 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)*.
- Marco Lui and Timothy Baldwin. 2012. [langid.py: An off-the-shelf language identification tool](#). In *Proceedings of the ACL 2012 System Demonstrations*.
- Matteo Negri, Marco Turchi, Nicola Bertoldi, and Marcello Federico. 2017. [Online neural automatic post-editing for neural machine translation](#). In *Proceedings of the Fifth Italian Conference on Computational Linguistics*.
- Matteo Negri, Marco Turchi, Rajen Chatterjee, and Nicola Bertoldi. 2018. [eSCAPE: a large-scale synthetic corpus for automatic post-editing](#). In *Proceedings of the 11th International Conference on Language Resources and Evaluation*.
- 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*.
- Maja Popović. 2015. [chrF: character n-gram F-score for automatic MT evaluation](#). In *Proceedings of the 10th Workshop on Statistical Machine Translation*. Association for Computational Linguistics.
- Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*.
- Mark Przybocki, Kay Peterson, Sébastien Bronsart, and Gregory Sanders. 2009. The NIST 2008 Metrics for machine translation challenge overview, methodology, metrics, and results. *Machine Translation*, 23(2-3):71–103.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Improving neural machine translation models with monolingual data](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.
- Michel Simard, Cyril Goutte, and Pierre Isabelle. 2007. [Statistical phrase-based post-editing](#). In *Proceedings of Human Language Technologies: The 2007 Annual Conference of the North American Chapter of the 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 the Seventh Conference of the Association for Machine Translation in the Americas*.
- Lucia Specia, Kim Harris, Frédéric Blain, Aljoscha Buchardt, Viviven Macketanz, Inguna Skadia, Matteo Negri, , and Marco Turchi. 2017. Translation quality and productivity: A study on rich morphology languages. In *Proceedings of Machine Translation Summit XVI*.
- Liling Tan and Santanu Pal. 2014. [Manawi: Using multi-word expressions and named entities to improve machine translation](#). In *Proceedings of the Ninth Workshop on Statistical Machine Translation*.
- Amirhossein Tebbifakhr, Ruchit Agrawal, Matteo Negri, and Marco Turchi. 2018. [Multi-source transformer with combined losses for automatic post editing](#). In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*.
- 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 30*.
- Elena Voita, Rico Sennrich, and Ivan Titov. 2019. [Context-aware monolingual repair for neural machine translation](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the Ninth International Joint Conference on Natural Language Processing*.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, R’emi Louf, Morgan Funtowicz, and Jamie Brew. 2019. [HuggingFace’s Transformers: State-of-the-art natural language processing](#). *arXiv preprint*, arXiv:1910.03771.
- Ventsislav Zhechev. 2012. [Machine translation infrastructure and post-editing performance at Autodesk](#). In *Proceedings of the AMTA 2012 Workshop on Post-Editing Technology and Practice*.