

APE at Scale and its Implications on MT Evaluation Biases

Markus Freitag, Isaac Caswell, Scott Roy

Google Research

{freitag, icaswell, hsr}@google.com

Abstract

In this work, we train an Automatic Post-Editing (APE) model and use it to reveal biases in standard Machine Translation (MT) evaluation procedures. The goal of our APE model is to correct typical errors introduced by the translation process, and convert the “translationese” output into natural text. Our APE model is trained entirely on monolingual data that has been round-trip translated through English, to mimic errors that are similar to the ones introduced by NMT. We apply our model to the output of existing NMT systems, and demonstrate that, while the human-judged quality improves in all cases, BLEU scores drop with forward-translated test sets. We verify these results for the WMT18 English→German, WMT15 English→French, and WMT16 English→Romanian tasks. Furthermore, we selectively apply our APE model on the output of the top submissions of the most recent WMT evaluation campaigns. We see quality improvements on all tasks of up to 2.5 BLEU points.

1 Introduction

Neural Machine Translation (NMT) (Bahdanau et al., 2015; Gehring et al., 2017; Vaswani et al., 2017) is currently the most popular approach in Machine Translation leading to state-of-the-art performance for many tasks. NMT relies mainly on parallel training data, which can be an expensive and scarce resource. There are several approaches to leverage monolingual data for NMT: Language model fusion for both phrase-based (Brants et al., 2007) and neural MT (Gülçehre et al., 2015, 2017), back-translation (Sennrich et al., 2016b), unsupervised NMT (Lample et al., 2018a; Artetxe et al., 2018a), dual learning (Cheng et al., 2016; He et al., 2016; Xia et al., 2017), and multi-task learning (Domhan and Hieber, 2017).

In this paper, we present a different approach to leverage monolingual data, which can be used as a post-processor for any existing translation. The idea is to train an Automatic Post-Editing (APE) system that is only trained on a large amount of synthetic data, to fix typical errors introduced by the translation process. During training, our model uses a noisy version of each sentence as input and learns how to reconstruct the original sentence. In this work, we model the noise with round-trip translations (RTT) through English, translating a sentence in the target language into English, then translating the result back into the original language. We train our APE model with a standard transformer model on the WMT18 English→German, WMT15 English→French and WMT16 English→Romanian monolingual News Crawl data and apply this model on the output of NMT models that are either trained on all available bitext or trained on a combination of bitext and back-translated monolingual data. Furthermore, we show that our APE model can be used as a post-processor for the best output of the recent WMT evaluation campaigns, where it improves even the output of these well engineered translation systems.

In addition to measuring quality in terms of BLEU scores on the standard WMT test sets, we split each test set into two subsets based on whether the source or target is the original sentence (each sentence is either originally written in the source or target language and human-translated into the other). We call these the source-language-original and target-language-original halves, respectively. We find that evaluating our post-edited output on the source-language-original half actually decreases the BLEU scores, whereas the BLEU scores improve for the target-language-original half. This is in line with results from Koppel and Ordan (2011), who demonstrate

that the mere fact of being translated plays a crucial role in the makeup of a translated text, making the actual (human) translation a less natural example of the target language. We hypothesize that, given these findings, the consistent decreases in BLEU scores on test sets whose source side are natural text does not mean that the actual output is of lower quality. To verify this hypothesis, we run human evaluations for different outputs with and without APE. The human ratings demonstrate that the output of the APE model is both consistently more accurate and consistently more fluent, regardless of whether the source or the target language is the original language, contradicting the corresponding BLEU scores.

To summarize the contributions of the paper:

- We introduce an APE model trained only on synthetic data generated with RTT for fixing typical translation errors from NMT output and investigate its scalability. To the best of our knowledge, this paper is the first to study the effect of an APE system trained at scale and only on synthetic data.
- We improve the BLEU of top submissions of the recent WMT evaluation campaigns.
- We show that the BLEU scores of the APE output only correlate well with human ratings when they are calculated with target-original references.
- We propose separately reporting scores on test sets whose source sentences are translated and whose target sentences are translated, and call for higher-quality test sets.

2 APE with RTT

2.1 Definition and Training

We formalize our APE model as a translation model from synthetic “translationese” (Gellerstam, 1986) text in one language to natural text in the same language. For a language pair (X, Y) and a monolingual corpus M_Y in language Y , the training procedure is as follows:

1. Train two translation models on bitext for $X \rightarrow Y$ and $Y \rightarrow X$
2. Use these models to generate round-trip translations for every target-language sentence y in M_Y , resulting in the synthetic dataset $\text{RTT}(M_Y)$.

3. Train a translation model on pairs of $(\text{RTT}(y), y)$, that translates from the round-tripped version of a sentence to its original form.

This procedure is illustrated in Figure 1.




Figure 1: Training procedure of our APE model with RTT in language Y .

2.2 Application

Given a trained translation model and a trained APE model, the procedure is simply to a) translate any source text from language X to language Y with the translation model, and b) post-edit the output of the translation by passing it through the APE model. In this sense, the APE model may also be viewed as a paraphrasing model to produce “naturalized” text. This procedure is illustrated in Figure 2.




Figure 2: Automatic Post-Editing (APE) as post-processor of NMT.

3 Experimental Setup

3.1 Architecture

For the translation models, we use the transformer implementation in *lingvo* (Shen et al., 2019), using the transformer-base model size for Romanian→English and transformer-big model size (Vaswani et al., 2017) for German→English and French→English. The reverse models, English→Romanian, English→German and English→French, are all transformer-big. All use a vocabulary of 32k subword units, and exponentially moving averaging of checkpoints (EMA decay) with the weight decrease parameter set to $\alpha = 0.999$ (Buduma and Locascio, 2017).

The APE models are also transformer models with 32k subword units and EMA decay trained

with lingvo. For the German and the French APE models, we use the transformer-big size, whereas for the Romanian APE model, we use the smaller transformer-base setup as we have less monolingual data.

3.2 Evaluation

We report BLEU (Papineni et al., 2002) and human evaluations. All BLEU scores are calculated with sacreBLEU (Post, 2018)¹.

Since 2014, the organizers of the WMT evaluation campaign (Bojar et al., 2017) have created test sets with the following method: first, they crawled monolingual data in both English and the target language from news stories from online sources. Thereafter they took about 1500 English sentences and translated them into the target language, and an additional 1500 sentences from the target language and translated them into English. This results in test sets of about 3000 sentences for each English-X language pair. The sgm files of each WMT test set include the original language for each sentence.

Therefore, in addition to reporting overall BLEU scores on the different test sets, we also report results on the two subsets (based on the original language) of each newstest20XX, which we call the {German,French,Romanian}-original and English-original halves of the test set. This is motivated by Koppel and Ordan (2011), who demonstrated that they can train a simple classifier to distinguish human-translated text from natural text with high accuracy. These text categorization experiments suggest that both the source language and the mere fact of being translated play a crucial role in the makeup of a translated text. One of the major goals of our APE model is to rephrase the NMT output in a more natural way, aiming to remove undesirable translation artifacts that have been introduced.

To collect human rankings, we present each output to crowd-workers, who were asked to score each sentence on a 5-point scale for:

- **fluency:** How do you judge the overall naturalness of the utterance in terms of its grammatical correctness and fluency?

Further, we included the source sentence and asked the raters to evaluate each sentence on a 2-point scale (binary decision) for:

¹sacreBLEU signatures: BLEU+case.mixed+lang.en-LANG+numrefs.1+smooth.exp+SET+tok.intl+version.1.2.20

- **accuracy:** Does the statement factually contradict anything in the reference information?

Each task was given to three different raters. Consequently, each output has a separate score for each question that is the average of 3 different ratings.

3.3 Data

For the round-trip experiments we use the monolingual News Crawl data from the WMT evaluation campaign. We remove duplicates and apply a max-length filter on the source sentences and the round-trip translations, filtering to the minimum of 500 characters or 70 tokens. For German, we concatenate all News Crawl data from 2007 to 2017, comprising 216.5M sentences after filtering and removing duplicates. For Romanian, we use News Crawl '16, comprising 2.2M sentences after filtering and deduplication. For French, we concatenate News Crawl data from 2007 to 2014, comprising 34M sentences after filtering.

Our translation models are trained on WMT18 (~5M sentences for German after filtering), WMT16 (~0.5M sentences for Romanian after filtering) and WMT15 (~41M sentences for French) bitext. For Romanian and German we filter sentence pairs that have empty source or target, that have source or target longer than 250 tokens, or the ratio of whose length is greater than 2.0. For English→German and English→French, we also build a system based on noised back-translation, as in Edunov et al. (2018). We use the same monolingual sentences that we used for the APE model to generate the noisy back-translation data.

4 Experiments

4.1 English→German

The results of our English→German experiments are shown in Table 1. We trained the APE model on RTT produced by English→German and German→English NMT models that are only trained on bitext. Applying the APE model on the output of our NMT model also trained on only bitext improves the BLEU scores by up to 1.5 BLEU points for newstest2014 and 0.7 BLEU points for newstest2017. Nevertheless, the score drops by 1.4 points on newstest2016. To investigate the differing impact on the test sets, we split each test set by its original language (Table 2). The APE model consistently increases the BLEU on the German-original half of the test set, but decreases the BLEU

	newstest2014	newstest2015	newstest2016	newstest2017	average
Vaswani et al. (2017)	28.4	-	-	-	
Shaw et al. (2018)	29.2	-	-	-	
our bitext	29.2	31.4	35.0	29.4	31.2
+ RTT APE	30.7	31.2	33.6	30.1	31.4
+ RTT APE de-orig only	31.7	32.9	37.2	31.9	33.4
our NBT	33.5	34.4	38.3	32.5	34.7
+ RTT APE	32.5	32.7	35.2	31.3	32.9
+ RTT APE de-orig only	34.0	34.5	38.7	33.2	35.1

Table 1: BLEU scores for WMT18 English→German. We apply the same APE model (trained on RTT with bitext models) for both an NMT system based on pure bitext and an NMT system that uses noised back-translation (NBT) in addition to bitext.

	newstest2014		newstest2015		newstest2016		newstest2017		average	
	orig-de	orig-en	orig-de	orig-en	orig-de	orig-en	orig-de	orig-en	orig-de	orig-en
our bitext	28.4	29.4	26.5	33.3	29.9	38.2	25.9	31.6	27.7	33.1
+ RTT APE	34.1	27.6	31.3	30.9	35.7	32.2	32.1	28.5	33.3	29.8
our NBT	35.6	31.3	32.6	34.7	37.6	38.7	31.7	32.6	34.4	34.3
+ RTT APE	36.9	28.8	33.5	32.0	38.5	32.9	33.8	29.2	35.7	30.7

Table 2: BLEU scores for WMT18 English→German. Test sets are divided by their original source language (either German or English).

	newstest2016				newstest2017			
	fluency		accuracy		fluency		accuracy	
	orig-de	orig-en	orig-de	orig-en	orig-de	orig-en	orig-de	orig-en
baseline bitext	4.65	4.49	95.6%	94.4%	4.74	4.52	97.2%	94.6%
+ RTT APE	4.77	4.59	98.4%	95.0%	4.84	4.58	98.0%	94.8%
our NBT	4.79	4.64	98.2%	95.8%	4.79	4.65	98.2%	95.4%
+ RTT APE	4.82	4.63	98.0%	96.2%	4.86	4.68	98.0%	96.4%
reference	4.85	4.67	98.6%	98.6%	4.83	4.70	98.0%	99.2%

Table 3: English→German human evaluation results split by original language of the test set.

on the English-original half. Consequentially, we applied our APE model only on the sentences with original language in German (+RTT APE de-orig only in Table 1) and see consistent improvements over all test sets with an average BLEU improvement of 2.2 points.

To verify that the drop in BLEU score is because of the unnatural reference translations, we run a human evaluation (see Section 3.2) for both fluency/grammatical correctness and accuracy. Based on the human ratings (Table 3), our APE model also improves on the English-original half of the test set (which is a more realistic use case).

Without re-training, we use the APE model that is trained on the bitext RTT and apply it to a stronger NMT system that also includes all the available monolingual data in the form of noised back-translation. We see a very similar pattern to the previous experiments. Regarding automatic scores, our APE model only improves on the German-original part of the test sets, with an average improvement of 1.3 BLEU points. The human evaluations show the same inconsistency with the automatic scores for the English-original half. As with the weaker baseline, humans rate the output of our APE model at least as fluent and accurate as the original output of the NMT model (Table 3). Further, we also run a human evaluation on

the reference sentences and found that the scores for both fluency and accuracy are only minimally higher than for our APE NBT output.

Comparing only the BLEU scores from our bitext and NBT models in Table 2 reveals that augmenting the parallel data with back-translated data also mostly improves the BLEU scores on the German-original half of the test set. This is in line with the results of our APE model and opens the question of how much of the original bitext data is natural on the target side.

As our APE model seems agnostic to the model which produced the RTT, we applied it to the best submissions of the recent WMT18 evaluation campaign, applying to German-original half of the test set only. Table 4 shows the results for the 2 top submissions of Microsoft (Junczys-Dowmunt, 2018) and Cambridge (Stahlberg et al., 2018). Both systems improved by up to 0.8 points in BLEU.

	Microsoft	Cambridge
WMT18 submission	48.7	47.2
+ APE only de-orig	49.5	47.7

Table 4: BLEU scores for WMT18 English→German newstest2018. We apply our APE model only on the German-original half of the test set. BLEU scores are calculated on the full newstest2018 set and the English-original half is just copied from the submission.

Finally, we train our APE model on different random subsets of the available 216.5M monolingual data (see Figure 3). The average BLEU scores on newstest2014-newstest2017 show that we can achieve similar performance by using 24 million training examples only, and that large improvements are seen using as few as 4M training examples.

4.2 English→Romanian

Experimental results for the WMT16 English→Romanian task are summarized in Table 5. By applying our APE model on top of a baseline that is only trained on bitext, we see improvements of 3.0 BLEU (dev) and 0.3 BLEU (test) over our baseline system when we automatically post edit only to the Romanian-original half of the test set. Similar to English→German, we apply our APE model on the top 2 submissions of the WMT16 evaluation campaign (Table 6). Both the QT21 submission (Peter et al., 2016), which is a system combination of several NMT systems,




Figure 3: English→German - Average BLEU scores for newstest2014-newstest2017: Our APE model is trained on different subsets of the monolingual data.

and the ensemble of the University of Edinburgh (Sennrich et al., 2016a) improve, by 0.3 BLEU and 0.2 BLEU on test, respectively.

	dev	test
Sennrich et al. (2016a)	-	28.8
our bitext	27.0	28.9
+ RTT APE	27.3	29.0
+ RTT APE only ro-orig	30.0	29.2

Table 5: BLEU scores for our models on WMT16 English→Romanian.

	QT21	Edinburgh
WMT16 submission	29.4	28.8
+ RTT APE only ro-orig	29.7	29.0

Table 6: BLEU scores for WMT16 English→Romanian test set. Our APE model was applied on top of the best WMT16 submissions.

4.3 English→French

Experimental results for English→French are summarized in Table 7. We see the same tendency as we saw for German and Romanian. When applying our APE system on the output of the bitext baseline, we get a small improvement of 0.1 BLEU. By only post-editing the French-original half, we get an improvement of 1.0 BLEU points. The same effect can be seen on the English→French system that is trained with Noised BT. We yield quality improvements of 0.8 BLEU by applying our APE model on the French-original half of the test set only.

	newstest2014
our bitext	43.2
+ RTT APE	43.3
+ RTT APE only fr-orig	44.2
our NBT	45.3
+ RTT APE	44.6
+ RTT APE only fr-orig	46.1

Table 7: BLEU scores for WMT15 English→French.

5 Example Output

We would like to highlight a few short examples where our APE model improves the NMT translation in German. Although our APE model is also quite helpful for long sentences, we will focus on short examples for the sake of simplicity. In Table 8 there are examples from the English→German noised back-translated (NBT) setup (see Table 1), with and without automatic post editing. In the first example, NMT translates *club* (i.e. cudgel) incorrectly into *Club* (i.e. organization). Based on the context of the sentence, our APE model learned that *club* has to be translated into *Schlagstock* (i.e. cudgel). The next two examples are very similar as our APE model improves the word choice of the translations by taking the context of the sentence into account. The NMT translations of the last two examples make little sense and our APE model rephrases the output into a fluent, meaningful sentence.

6 Discussion

In this section, we focus on the results on target-language-original test sets, like the English-original subset of newstest2016 (Table 2 and Table 3), where the APE model lowered the score by 6 BLEU, yet improved human evaluations. A naïve take-away from this result would be that evaluation sets whose target side is natural text are inherently superior. However, translating *from* translationese also has its own problems, including 1) it does not represent any real-world translation task, and 2) translationese sources may be much easier to translate “correctly”, and reward MT biases like word-for-word translation. The take-away, therefore, must be to report scores both on the source-language-original and the target-language-original test sets, rather than lumping two test-sets together into one as has heretofore been done. This gives a higher-precision glimpse into the strengths and

weaknesses of different modeling techniques, and may prevent some effects (like improvements in naturalness of output) from being hidden.

Going forward, our results should also be seen as a call for higher-quality test sets. Multi-reference BLEU is one option and less likely to suffer these biases as acutely, and has previously been used in the NIST projects. Another option could be to align sentence pairs from monolingual data sets in two languages and run human evaluation to exclude bad sentence pairs.

7 Related Work

Automatic Post-Editing

Probably most similar to our work, [Junczys-Dowmunt and Grundkiewicz \(2016, 2018\)](#) uses RTT as additional training data for the automatic post-editing (APE) task of the WMT evaluation campaign ([Chatterjee et al., 2018](#)). They claimed that the provided post-editing data is orders of magnitude too small to train neural models, and combined the training data with artificial training data generated with RTT. They found that the additional artificial data helps against early overfitting and makes it possible to overcome the problem of too little training data. In contrast to our work, they do not report results for models only trained on the artificial RTT data. Further, their RTT data is much smaller (10M sentences) compared to ours (up to 200M sentences) and they only report results for the APE subtask.

There have been several earlier approaches using RTT for APE. [Hermet and Alain \(2009\)](#) used RTT to improve a standard preposition error detection system. Although their evaluation corpus was limited to 133 prepositions, the hybrid system outperformed their standard method by roughly 13%. [Madnani et al. \(2012\)](#) combined RTT obtained from Google Translate via 8 different pivot languages into a lattice for grammatical error correction. Similar to system combination, their final output is extracted by the shortest path scored by different features. They claimed that their preliminary experiments yield fairly satisfactory results but leave significant room for improvement.

Back-translation

Back-translation ([Sennrich et al., 2016b](#); [Poncelas et al., 2018](#)) augments relatively scarce parallel data with plentiful monolingual data, allowing one to train source-to-target (S2T) models with the help of target-to-source (T2S) models. Specif-

source	Using a club , they beat the victim in the face and upper leg.
NBT	Mit einem Club schlagen sie das Opfer in Gesicht und Oberschenkel.
+ RTT APE	Mit einem Schlagstock schlagen sie dem Opfer ins Gesicht und in den Oberschenkel.
source	Müller put another one in with a with a penalty .
NBT	Müller setzte einen weiteren mit einer Strafe ein.
+ RTT APE	Müller netzte einen weiteren per Elfmeter ein.
source	Obama receives Netanyahu
NBT	Obama erhält Netanjahu
+ RTT APE	Obama empfängt Netanjahu
source	At least one Bayern fan was taken injured from the stadium .
NBT	Mindestens ein Bayern-Fan wurde vom Stadion verletzt .
+ RTT APE	Mindestens ein Bayern-Fan wurde verletzt aus dem Stadion gebracht .
source	The archaeologists made a find in the third construction phase of the Rhein Boulevard.
NBT	Die Archäologen haben in der dritten Bauphase des Rheinboulevards gefunden .
+ RTT APE	Die Archäologen sind im dritten Bauabschnitt des Rheinboulevards fündig geworden .

Table 8: Example output for English→German.

ically, given a set of sentences in the target language, a pre-constructed T2S translation system is used to generate translations to the source language. These synthetic sentence pairs are combined with the original bilingual data when training the S2T NMT model.

Iterative Back-translation

Iterative back-translation (Zhang et al., 2018; Cotterell and Kreutzer, 2018; Hoang et al., 2018) is a joint training algorithm to enhance the effect of monolingual source and target data by iteratively boosting the source-to-target and target-to-source translation models. The joint training method uses the monolingual data and updates NMT models through several iterations. A variety of flavors of iterative back-translation have been proposed, including Niu et al. (2018), who simultaneously perform iterative S2T and T2S back-translation in a multilingual model, and He et al. (2016); Xia et al. (2017), who combine dual learning with phases of back- and forward-translation.

Artetxe et al. (2018a,b) and Lample et al. (2018a,b) used iterative back-translation to train two unsupervised translation systems in both directions ($X \rightarrow Y$ and $Y \rightarrow X$) in parallel. Further, they used back-translation to generate a synthetic source to construct a dev set for tuning the parameters of their unsupervised statistical machine translation system. In a similar formulation, Cheng et al. (2016) jointly learn a translation system with a round-trip autoencoder.

Round-tripping and Paraphrasing

Round-trip translation has seen success as a

method to generate paraphrases. Bannard and Callison-Burch (2005) extracted paraphrases by using alternative phrase translations from bilingual phrase tables from Statistical Machine Translation. Mallinson et al. (2017) presented PARANET, a neural paraphrasing model based on round-trip translations with NMT. They showed that their paraphrase model outperforms all traditional paraphrase models.

Wu et al. (2018) train a paraphrasing model on $(X, RTT(X))$ pairs, translating from natural text into a simplified version. They apply this sentence-simplifier on the source sentences of the training data and report quality gains for IWSLT.

Translationese and Artifacts from NMT

The difference between translated sentence pairs based on whether the source or the target is the original sentence has long been recognized by the human translation community, but only partially investigated by the machine translation community. An introduction to the latter is presented in Koppel and Ordan (2011), who train a high-accuracy classifier to distinguish human-translated text from natural text in the Europarl corpus. This is in line with research from the professional translation community, which has seen various works investigating both systematic biases inherent to translated texts (Baker, 1993; Selinker, 1972), as well as biases resulting specifically from interference from the source text (Toury, 1995). There has similarly long been a focus on the conflict between *Fidelity* (the extent to which the translation is faithful to the source) and *Transparency* (the ex-

tent to which the translation appears to be a natural sentence in the target language) (Warner, 2018; Schleiermacher, 1816; Dryden, 1685). To frame our hypotheses in these terms, the present work hypothesizes that outputs from NMT systems often err on the side of disfluent fidelity, or word-by-word translation.

There are a few papers that discuss the effect of translationese on MT models. Lembersky et al. (2012); Stymne (2017) explored how the translation direction for statistical machine translation affects the translation result. They found that using training and tuning data translated in the same direction as the translation systems tends to give the best results. Holmqvist et al. (2009) noted that the original language of the test sentences influences the BLEU score of translations. They showed that the BLEU scores for target-original sentences are on average higher than sentences that have their original source in a different language. Popel (2018) split the WMT Czech-English test set based on the original language. They found that when training on synthetic data, the model performs much better on the Czech-original half than on the non Czech-original half. When trained on authentic data, it is the other way round. Fomicheva et al. (2017) found that both the average score and Pearson correlation with human judgments is substantially higher when both the MT output and human translation were generated from the same source language.

8 Ablation

8.1 Iterative APE

We can apply our APE model in an iterative fashion several times on the same output. In Table 9, we applied our APE model on the already post-edited output to see if we can further naturalize the sentences. As a result, 75% of the sentences did not change. The remaining sentences lowered the BLEU scores on average by 0.1 points for German-original half and by 0.7 points for the English-original half of the test sets.

8.2 Reverse APE

Instead of training an APE model on $(RTT(y), y)$ sentence pairs (see Section 2), we train in this section a `reverse APE` model that flips source and target and is trained on $(y, RTT(y))$ sentence pairs. Experimental results can be seen in Table 10. Overall, the performance decreases on

	average	
	orig-de	orig-en
our bitext	27.7	33.1
+ APE	33.3	29.8
+ 2xAPE	33.2	29.1

Table 9: Average BLEU scores for WMT18 English→German newstest2014-2017. We run our APE model a second time on the output of the already post-edited output.

both the German-original half and the English-original half. Interestingly, the BLEU scores of the reverse APE model on the English-original half are higher than those of the normal APE model. This demonstrates again that sentences evaluated on the English-original half prefer output that is biased by the translation process.

	average	
	orig-de	orig-en
our bitext	27.7	33.1
+ RTT APE	33.3	29.8
+ Reverse APE	25.1	30.6
our NBT	34.4	34.3
+ RTT APE	35.7	30.7
+ Reverse APE	27.0	31.3

Table 10: Average BLEU scores for WMT18 English→German newstest2014-2017.

8.3 Inside the black box of RTT

In this subsection we are interested in how much RTT changes translation outputs. We calculate the BLEU scores of all English→German test sets (11,175 sentences in total) by taking the original German sentences as references and their RTT as hypotheses. Although the RTT hypotheses are a less clean (paraphrased) version of the references, having been forward-translated from an already noisy back-translated source, the BLEU score is 40.9, with unigram precision of 72.3%, bigram precision of 48.9%, trigram precision of 35.6% and 4gram precision of 26.6%.

We further found that the original sentences use a larger vocabulary than the artificial RTT data. While the output of the RTT has only 29,635 unique tokens, the original sentences contain 33,814 unique tokens. Even more interesting, the NMT output (from the model trained only

on bitext) of the same test sets has a vocabulary size of 30,540, but after running our APE on the same test sets the vocabulary size increases to 31,471. The NMT output from the NBT model has a vocabulary size of 32,170 and its post-edited version increases the number of unique words to 32,283. Overall, we see that both the RTT and the NMT output have a smaller vocabulary size than the original or post-edited versions, and that BLEU score grows directly with increased number of unique tokens in the target side.

9 Conclusion

We propose an APE model that is only trained on RTT and increases the quality of NMT translations, measured both by BLEU and human evaluation. We see improvements both when automatically post editing our model translations and when automatically post editing outputs from the winning submissions to the WMT competition. Our APE has the advantage that it is agnostic to the model which produced the translations, and so can be used on top of any new advance in the field, without need for re-training. Further, we demonstrate that we need only a subset of 24M training examples to train our APE model. We furthermore use this model as a tool to reveal systematic problems with reference translations, and propose finer-grained BLEU reporting on both source-language-original test sets and target-language-original test sets, as well as calling for higher-quality and multi-reference test sets.

References

- Mikel Artetxe, Gorika Labaka, and Eneko Agirre. 2018a. [Unsupervised Statistical Machine Translation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3632–3642.
- Mikel Artetxe, Gorika Labaka, Eneko Agirre, and Kyunghyun Cho. 2018b. [Unsupervised Neural Machine Translation](#). In *International Conference on Learning Representations*.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. [Neural Machine Translation by Jointly Learning to Align and Translate](#). In *3rd International Conference on Learning Representations, ICLR 2015*.
- Mona Baker. 1993. *Corpus Linguistics and Translation Studies: Implications and Applications*. *Text and technology: in honour of John Sinclair*, pages 233–252.
- Colin Bannard and Chris Callison-Burch. 2005. [Paraphrasing with Bilingual Parallel Corpora](#). In *Proceedings of the 43rd Annual Meeting on Association for Computational Linguistics*, pages 597–604.
- Ondřej Bojar, Rajen Chatterjee, Christian Federmann, Yvette Graham, Barry Haddow, Shujian Huang, Matthias Huck, Philipp Koehn, Qun Liu, Varvara Logacheva, et al. 2017. [Findings of the 2017 Conference on Machine Translation \(WMT17\)](#). In *Proceedings of the Second Conference on Machine Translation*, pages 169–214.
- Thorsten Brants, Ashok C Popat, Peng Xu, Franz J Och, and Jeffrey Dean. 2007. [Large Language Models in Machine Translation](#). In *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*.
- Nikhil Buduma and Nicholas Locascio. 2017. [Fundamentals of Deep Learning: Designing Next-Generation Machine Intelligence Algorithms](#). “O’Reilly Media, Inc.”.
- 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*, pages 710–725.
- Yong Cheng, Wei Xu, Zhongjun He, Wei He, Hua Wu, Maosong Sun, and Yang Liu. 2016. [Semi-Supervised Learning for Neural Machine Translation](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, volume 1, pages 1965–1974.
- Ryan Cotterell and Julia Kreutzer. 2018. [Explaining and Generalizing Back-Translation through Wakesleep](#). *arXiv preprint arXiv:1806.04402*.
- Tobias Domhan and Felix Hieber. 2017. [Using Target-side Monolingual Data for Neural Machine Translation through Multi-task Learning](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1500–1505.
- John Dryden. 1685. *Sylvae [Translator’s preface]*. A Scholar Press facsimile. Scholar Press.
- Sergey Edunov, Myle Ott, Michael Auli, and David Grangier. 2018. [Understanding Back-Translation at Scale](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 489–500.
- Marina Fomicheva et al. 2017. [The Role of Human Reference Translation in Machine Translation Evaluation](#). Ph.D. thesis, Universitat Pompeu Fabra.
- Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N. Dauphin. 2017. [Convolutional Sequence to Sequence Learning](#). In *Proceedings of the 34th International Conference on Machine Learning - Volume 70*, pages 1243–1252.

- Martin Gellerstam. 1986. Translationese in Swedish novels translated from English. *Translation Studies in Scandinavia*, pages 88–95.
- Çağlar Gülçehre, Orhan Firat, Kelvin Xu, Kyunghyun Cho, Loic Barrault, Huei-Chi Lin, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2015. [On using Monolingual Corpora in Neural Machine Translation](#). *arXiv preprint arXiv:1503.03535*.
- Çağlar Gülçehre, Orhan Firat, Kelvin Xu, Kyunghyun Cho, and Yoshua Bengio. 2017. [On Integrating a Language Model into Neural Machine Translation](#). *Comput. Speech Lang.*, pages 137–148.
- Di He, Yingce Xia, Tao Qin, Liwei Wang, Nenghai Yu, Tiejun Liu, and Wei-Ying Ma. 2016. [Dual Learning for Machine Translation](#). In *Conference on Advances in Neural Information Processing Systems (NeurIPS)*.
- Matthieu Hermet and Désilets Alain. 2009. [Using First and Second Language Models to Correct Preposition Errors in Second Language Authoring](#). In *Proceedings of the Fourth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 64–72.
- Vu Cong Duy Hoang, Philipp Koehn, Gholamreza Haffari, and Trevor Cohn. 2018. [Iterative Backtranslation for Neural Machine Translation](#). In *Proceedings of the 2nd Workshop on Neural Machine Translation and Generation*, volume 1, pages 18–24.
- Maria Holmqvist, Sara Stymne, Jody Foo, and Lars Ahrenberg. 2009. [Improving Alignment for SMT by Reordering and Augmenting the Training Corpus](#). In *Proceedings of the Fourth Workshop on Statistical Machine Translation*, pages 120–124. Association for Computational Linguistics.
- Marcin Junczys-Dowmunt. 2018. [Microsoft’s Submission to the WMT2018 News Translation Task: How I Learned to Stop Worrying and Love the Data](#). In *Proceedings of the Third Conference on Machine Translation, Volume 2: Shared Task Papers*, pages 429–434. 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.
- 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*, pages 822–826.
- Moshe Koppel and Noam Ordan. 2011. [Translationese and Its Dialects](#). In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies - Volume 1*, pages 1318–1326.
- Guillaume Lample, Alexis Conneau, Ludovic Denoyer, and Marc’Aurelio Ranzato. 2018a. [Unsupervised Machine Translation Using Monolingual Corpora Only](#). In *International Conference on Learning Representations*.
- Guillaume Lample, Myle Ott, Alexis Conneau, Ludovic Denoyer, and Marc’Aurelio Ranzato. 2018b. [Phrase-Based & Neural Unsupervised Machine Translation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Gennadi Lembersky, Noam Ordan, and Shuly Wintner. 2012. [Language Models for Machine Translation: Original vs. Translated Texts](#). *Computational Linguistics*, 38(4).
- Nitin Madnani, Joel Tetreault, and Martin Chodorow. 2012. [Exploring grammatical error correction with not-so-crummy machine translation](#). In *Proceedings of the Seventh Workshop on Building Educational Applications Using NLP*, pages 44–53. Association for Computational Linguistics.
- Jonathan Mallinson, Rico Sennrich, and Mirella Lapata. 2017. [Paraphrasing Revisited with Neural Machine Translation](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, volume 1, pages 881–893.
- Xing Niu, Michael Denkowski, and Marine Carpuat. 2018. [Bi-Directional Neural Machine Translation with Synthetic Parallel Data](#). In *Proceedings of the 2nd Workshop on Neural Machine Translation and Generation*, pages 84–91.
- 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 on association for computational linguistics*, pages 311–318. Association for Computational Linguistics.
- Jan-Thorsten Peter, Tamer Alkhouli, Hermann Ney, Matthias Huck, Fabienne Braune, Alexander Fraser, Aleš Tamchyna, Ondřej Bojar, Barry Haddow, Rico Sennrich, Frédéric Blain, Lucia Specia, Jan Niehues, Alex Waibel, Alexandre Allauzen, Lauriane Aufrant, Franck Burlot, elena knyazeva, Thomas Lavergne, François Yvon, Mărcis Pinnis, and Stella Frank. 2016. [The QT21/HimL Combined Machine Translation System](#). In *Proceedings of the First Conference on Machine Translation (WMT16)*, pages 344–355.
- Alberto Poncelas, Dimitar Shterionov, Andy Way, Gideon Maillette de Buy Wenniger, and Peyman Passban. 2018. [Investigating Backtranslation in Neural Machine Translation](#). In *Proceedings of the*

- 21st Annual Conference of the European Association for Machine Translation, pages 249–258.
- Martin Popel. 2018. [CUNI Transformer Neural MT System for WMT18](#). In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pages 482–487.
- Matt Post. 2018. [A Call for Clarity in Reporting Bleu Scores](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191.
- Friedrich Schleiermacher. 1816. *Über die verschiedenen Methoden des Übersetzens*. Abhandlungen der Königlichen Akademie der Wissenschaften in Berlin. Walter de Gruyter GmbH.
- Larry Selinker. 1972. Interlanguage. *International Review of Applied Linguistics*, pages 209–241.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016a. [Edinburgh Neural Machine Translation Systems for WMT 16](#). In *Proceedings of the First Conference on Machine Translation*, pages 371–376.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016b. [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)*, pages 86–96.
- Peter Shaw, Jakob Uszkoreit, and Ashish Vaswani. 2018. [Self-Attention with Relative Position Representations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 464–468.
- Jonathan Shen, Patrick Nguyen, Yonghui Wu, Zhifeng Chen, Mia X. Chen, Ye Jia, Anjali Kannan, Tara N. Sainath, and Yuan Cao et al. 2019. [Lingvo: a Modular and Scalable Framework for Sequence-to-Sequence Modeling](#). *CoRR*, abs/1902.08295.
- Felix Stahlberg, Adri de Gispert, and Bill Byrne. 2018. [The University of Cambridge’s Machine Translation Systems for WMT18](#). In *Proceedings of the Third Conference on Machine Translation, Volume 2: Shared Task Papers*, pages 508–516.
- Sara Stymne. 2017. [The Effect of Translationese on Tuning for Statistical Machine Translation](#). In *The 21st Nordic Conference on Computational Linguistics*, pages 241–246.
- Gideon Toury. 1995. *Descriptive Translation Studies and Beyond*. Benjamins translation library. John Benjamins Publishing Company.
- 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*, pages 5998–6008.
- Marina Warner. 2018. The Politics of Translation. *London Review of Books*, 40 no. 19:22.
- Tian Wu, Zhongjun He, Enhong Chen, and Haifeng Wang. 2018. [Improving Neural Machine Translation with Neural Sentence Rewriting](#). In *2018 International Conference on Asian Language Processing (IALP)*.
- Yingce Xia, Tao Qin, Wei Chen, Jiang Bian, Nenghai Yu, and Tie-Yan Liu. 2017. [Dual Supervised Learning](#). In *International Conference on Machine Learning (ICML)*.
- Zhirui Zhang, Shujie Liu, Mu Li, Ming Zhou, and Enhong Chen. 2018. [Joint Training for Neural Machine Translation Models with Monolingual Data](#). In *Thirty-Second AAAI Conference on Artificial Intelligence*.