

Human-Paraphrased References Improve Neural Machine Translation

Markus Freitag, George Foster, David Grangier, Colin Cherry

Google Research

{freitag, fosterg, grangier, colincherry}@google.com

Abstract

Automatic evaluation comparing candidate translations to human-generated paraphrases of reference translations has recently been proposed by Freitag et al. (2020). When used in place of original references, the paraphrased versions produce metric scores that correlate better with human judgment. This effect holds for a variety of different automatic metrics, and tends to favor natural formulations over more literal (*translationese*) ones. In this paper we compare the results of performing end-to-end system development using standard and paraphrased references. With state-of-the-art English-German NMT components, we show that tuning to paraphrased references produces a system that is significantly better according to human judgment, but 5 BLEU points worse when tested on standard references. Our work confirms the finding that paraphrased references yield metric scores that correlate better with human judgment, and demonstrates for the first time that using these scores for system development can lead to significant improvements.

1 Introduction

Machine Translation (MT) has shown impressive progress in recent years. Neural architectures (Bahdanau et al., 2015; Gehring et al., 2017; Vaswani et al., 2017) have greatly contributed to this improvement, especially for languages with abundant training data (Bojar et al., 2016, 2018; Barrault et al., 2019). This progress creates novel challenges for the evaluation of machine translation, both for human (Toral, 2020; Lübli et al., 2020) and automated evaluation protocols (Lo, 2019; Zhang et al., 2019).

Both types of evaluation play an important role in machine translation (Koehn, 2010). While human evaluations provide a gold standard evaluation, they involve a fair amount of careful and hence

expensive work by human assessors. Cost therefore limits the scale of their application. On the other hand, automated evaluations are much less expensive. They typically only involve human labor when collecting human reference translations and can hence be run at scale to compare a wide range of systems or validate design decisions. The value of automatic evaluations therefore resides in their capacity to be used as a proxy for human evaluations for large scale comparisons and system development.

The recent progress in MT has raised concerns about whether automated evaluation methodologies reliably reflect human ratings in high accuracy ranges. In particular, it has been observed that the best systems according to humans might fare less well with automated metrics (Barrault et al., 2019). Most metrics such as BLEU (Papineni et al., 2002) and TER (Snover et al., 2006) measure overlap between a system output and a human reference translation. More refined ways to compute such overlap have consequently been proposed (Banerjee and Lavie, 2005; Lo, 2019; Zhang et al., 2019).

Orthogonal to the work of building improved metrics, Freitag et al. (2020) hypothesized that human references are also an important factor in the reliability of automated evaluations. In particular, they observed that standard references exhibit simple, monotonic language due to human ‘translationese’ effects. These standard references might favor systems which excel at reproducing these effects, independent of the underlying translation quality. They showed that better correlation between human and automated evaluations could be obtained when replacing standard references with *paraphrased* references, even when still using surface overlap metrics such as BLEU (Papineni et al., 2002). The novel references, collected by asking linguists to paraphrase standard references, were shown to steer evaluation away from rewarding

translation artifacts. This improves the assessment of alternative, but equally good translations.

Our work builds on the success of paraphrased translations for evaluating existing systems, and asks if different design choices could have been made when designing a system with such an evaluation protocol in mind. This examination has several potential benefits: it can help identify choices which improve BLEU on standard references but have limited impact on final human evaluations; or those that result in better translations for the human reader, but worse in terms of standard reference BLEU. Conversely, it might turn out that paraphrased references are not robust enough to support system development due to the presence of ‘metric honeypots’: settings that produce poor translations, but which are nevertheless assigned high BLEU scores.

To address these points, we revisit the major design choices of the best English→German system from WMT2019 (Ng et al., 2019) step-by-step, and measure their impact on standard reference BLEU as well as on paraphrased BLEU. This allows us to measure the extent to which steps such as data cleaning, back-translation, fine-tuning, ensemble decoding and reranking benefit standard reference BLEU more than paraphrase BLEU. Revisiting these development choices with the two metrics results in two systems with quite different behaviors. We conduct a human evaluation for adequacy and fluency to assess the overall impact of designing a system using paraphrased BLEU.

Our main findings show that optimizing for paraphrased BLEU is advantageous for human evaluation when compared to an identical system optimized for standard BLEU. The system optimized for paraphrased BLEU significantly improves WMT newstest19 adequacy ratings (4.72 vs 4.27 on a six-point scale) and fluency ratings (63.8% vs 27.2% on side-by-side preference) despite scoring 5 BLEU points lower on standard references.

2 Related Work

Collecting human paraphrases of existing references has recently been shown to be useful for system evaluation (Freitag et al., 2020). Our work considers applying the same methodology for system tuning. There is some earlier work relying on *automated* paraphrases for system tuning, especially for Statistical Machine Translation (SMT). Madnani

et al. (2007) introduced an automatic paraphrasing technique based on English-to-English translation of full sentences using a statistical MT system, and showed that this permitted reliable system tuning using half as much data. Similar automatic paraphrasing has also been used to augment training data, e.g. (Marton et al., 2009), but relying on standard references for evaluation. In contrast to human paraphrases, the quality of current machine generated paraphrases degrades significantly as overlap with the input decreases (Mallinson et al., 2017; Roy and Grangier, 2019). This makes their use difficult for evaluation since (Freitag et al., 2020) suggests that substantial paraphrasing – ‘paraphrase as much as possible’ – is necessary for evaluation.

Our work can be seen as replacing the regular BLEU metric with a new paraphrase BLEU metric for system tuning. Different alternative automatic evaluation metrics have also been considered for system tuning (He and Way, 2010; Servan and Schwenk, 2011) with Minimum Error Rate Training, MERT (Och, 2003). This work showed some specific cases where Translation Error Rate (TER) was superior to BLEU.

Our work is also related to the bias that the human translation process introduces in the references, including source language artifacts—*Translationese* (Koppel and Ornan, 2011)—as well as source-independent artifacts—*Translation Universals* (Mauranen and Kujamäki, 2004). The professional translation community studies 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). For MT, Freitag et al. (2019) point at Translationese as a source of mismatch between BLEU and human evaluation, raising concerns that overlap-based metrics might reward hypotheses with translationese language more than hypotheses using more natural language. The impact of Translationese on human evaluation of MT has recently received attention as well (Toral et al., 2018; Zhang and Toral, 2019; Graham et al., 2019). More generally, the question of bias to a specific reference has also been raised, in the case of monolingual manual evaluation (Fomicheva and Specia, 2016; Ma et al., 2017). Different from the impact of Translationese on evaluation, the impact of Translationese in the training data has also been studied (Kurokawa et al., 2009; Lembersky et al., 2012a; Bogoychev and Sennrich, 2019; Riley et al., 2020).

Finally, our work is also related to studies measuring the importance of the test data quality, looking specifically at the test set translation direction. For SMT evaluation, [Lembersky et al. \(2012b\)](#) and [Stymne \(2017\)](#) explored how the translation direction affects translation 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. Recently, a similar study was conducted for neural MT ([Bogoychev and Sennrich, 2019](#)).

3 Experimental Setup

We first describe data and models, then present our human evaluation protocol.

3.1 Data

We ran all experiments on the WMT 2019 English→German news translation task ([Barrault et al., 2019](#)). The task provides $\sim 38\text{M}$ parallel sentences. As German monolingual data, we concatenate all News Crawl data from 2007 to 2018, comprising $\sim 264\text{M}$ sentences after removing duplicates.

In addition to the training data, we use newstest2018 for development and newstest2019 for evaluation only. There is an important difference between these two test sets. Newstest2018 was created from monolingual news data from both English and German online sources. Half of the data consists of English text translated into German, while the other half consists of German text translated into English. This results in a joint test set of 2,998 sentences. Newstest2019, on the other hand, consists only of 1,997 sentences translated from English into German (see Figure 1). To provide a joint test set similar to newstest2018, we took newstest2019 from the reverse translation direction German→English, swapped source and target, and concatenated it with the original test sets. This results in a new joint newstest2019 test set of 3,997 sentences.

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 *orig-en* and the *orig-de* halves of the test set.

[Freitag et al. \(2020\)](#) provided an alternative reference translation for the *orig-en* half of new-




Figure 1: Sentences in a test set are either natural in the source and forward-translated into the target language, or vice-versa. If a test set consists of both kinds of sentences, we call it a joint test set. WMT English→German newstest2018 is a joint test set with half of the sentences being forward-translated. WMT English→German newstest2019 is a forward-translated test set.

stest2019. For both standard and alternative references, they provided an additional paraphrased ‘as much as possible’ version (four different references in all). In order to enable our parameter tuning experiments, we created a paraphrased version of the reference for the *orig-en* half of *newstest2018* (1,500 sentences) following the instructions from [Freitag et al. \(2020\)](#). We will release this new paraphrased reference, *newstest2018.orig-en.p*, as part of our work.

3.2 Models

For our translation models, we adopt the transformer implementation from *Lingvo* ([Shen et al., 2019](#)), using the transformer-big model size ([Vaswani et al., 2017](#)). We 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](#)). We used a batch size of around 32k sentences in all our experiments.

We report BLEU ([Papineni et al., 2002](#)) in addition to human evaluation. All BLEU scores are calculated with sacreBLEU ([Post, 2018](#))¹.

3.3 Human Evaluation

To collect human rankings, we ran side-by-side evaluation for overall quality and fluency. We hired 20 linguists and divided them equally between the two evaluations. Each evaluation included 1,000 items with each item being rated exactly once. We acquired only a single rating per sentence from the professional linguists as we found that they were

¹BLEU+case.mixed+lang.ende+numrefs.1+smooth.exp+SET+tok.13a+version.1.4.12 SET $\in\{\text{wmt18, wmt19, wmt19/google/ar, wmt19/google/arp, wmt19/google/wmtp}\}$

more reliable than crowd workers (Toral, 2020). We evaluated the orig-en sentences corresponding to the official WMT-19 English→German test set (Barrault et al., 2019). Results in this natural translation direction are more meaningful as pointed out by Zhang and Toral (2019), who show that translating a ‘translationese’ source is simpler and should not be used for human evaluation. Our human evaluation followed the protocol:

- **Fluency:** We present two translations of the same source sentence to professional linguists without showing the actual source sentence. We then ask the rater whether they prefer one of the outputs or rate them equally based on fluency.
- **Overall Quality:** We present two translations along with the source and ask the raters to evaluate each translation on a 6-point scale. A score of 6 will be assigned to translations with ‘perfect meaning and grammar’, while a score of 0 will be assigned to ‘nonsense/ no meaning preserved’ translations. The average over all ratings yields the system’s final quality score.

4 Experimental Results

This section first presents our main result comparing the same system tuned with BLEU on standard versus paraphrased references. We then break down how system design choices impact each metric differently. Throughout, we refer to scores computed with standard references as BLEU, and those computed with paraphrased references as BLEUP.

4.1 Overall Performance

We compare the performance of a system optimized on newstest2018 with standard references (opt-on-BLEU) with one optimized on newstest2018.orig-en with paraphrased references (opt-on-BLEUP). Both systems were developed using only newstest2018 data, keeping newstest2019 as a blind test set. Table 1 summarizes the results on newstest2019. Details of how these two systems were developed and how they differ are given in Section 4.2.

The opt-on-BLEU system outperforms opt-on-BLEUP by 5.2 BLEU points. Normally this would lead us to discard opt-on-BLEUP. However, the BLEUP scores tell a different story: opt-on-BLEUP outperforms by 0.3 points, a potentially large improvement given the smaller natural range of this

metric. Under a significance test with random approximation (Riezler and Maxwell III, 2005), both the BLEU and BLEUP differences are significant at $p < 5e-18$.

	opt-on-BLEU	opt-on-BLEUP
BLEU	45.0	39.8
BLEUP	13.4	13.7
human quality	4.27	4.72
human fluency	27.2%	63.8%

Table 1: BLEU scores and human ratings for WMT newstest2019 English→German (original English sources). We optimized the system to perform best on either newstest2018 with standard reference translations (opt-on-BLEU) or newstest2018.orig-en with paraphrased reference translations (opt-on-BLEUP). BLEU differences are significant according to random approximation (Riezler and Maxwell III, 2005) with $p < 5e-18$. Human score differences are significant according to a Wilcoxon rank-sum test with $p < 5e-18$.

Freitag et al. (2020) showed that BLEU scores calculated on paraphrased references have higher correlation with human judgment than those calculated on standard references. To verify their findings, we ran a human evaluation for the two different outputs on 1,000 sentences randomly drawn from newstest2019 (orig-en), as described above. As shown in Table 1, opt-on-BLEUP is consistently evaluated as better for both quality and fluency. To measure the significance between the two ratings, we ran a Wilcoxon rank sum test on the human ratings and found that both improvements are significant with $p < e-18$.

This experiment demonstrates that we can actually tune our MT system on paraphrased references to yield higher translation quality when compared to a typical system tuned on standard BLEU. Interestingly, the BLEU score for the better system is much lower, supporting our contention that BLEU rewards spurious translation features (e.g. monotonicity and common translations) that are filtered out by BLEUP.

4.2 Analysing Performance

We now describe the individual model decisions that went into the two final systems of Section 4.1. To build a classical system optimized on BLEU with standard references, we replicate the WMT 2019 winning submission (Ng et al., 2019) and examine

the effect of each of its major design decisions.² In particular, we are looking into the effect of data cleaning, back-translation, fine tuning, ensembling and noisy channel reranking. We examine the impact of each method on BLEU and BLEUP. For our experiments, we used newstest2018 as our development set and newstest2019 as our held-out test set. All model decisions (checkpoint, variants) are solely made on newstest2018.

Experimental results are presented in Table 2. As described in Section 3.1, we report 4 different BLEU scores for newstest2018 (dev) and newstest2019 (test). In addition to reporting BLEU score on the joint or the orig-de/orig-en halves of the test sets, we also report BLEU scores that are calculated on paraphrased references (BLEUP).

4.2.1 Data Cleaning

For data cleaning, we used CDS (Wang et al., 2018). We trained a CDS model for English→German taking news-commentary as the in-domain/clean data set. We scored all parallel sentences with our trained CDS model and kept the 70% highest scoring sentences. Our experimental results suggest that data cleaning is useful for all four types of test sets and consistently improves over a baseline system that is trained on raw parallel data. We conclude that data cleaning is useful for all systems independently of which test set it will be optimized for.

4.2.2 Back-Translation

We trained a strong German→English model on the same parallel data (with flipped source/target) and used that model to (back-)translate (BT) all deduped German monolingual sentences from NewsCrawl 2007-2018 into English. We filtered sentences with a source-target ratio lower than 0.5 or higher than 1.5. We further run language identification and filtered out all backtranslations going into the wrong language. We then oversample our bitext data to match the size of the backtranslation data and train a NMT model on the concatenation of both datasets.

As previously reported by (Freitag et al., 2019; Bogoychev and Sennrich, 2019), the original language of the sentences within a test is crucial and can lead to very different conclusions, in particular for back-translation systems. This difference is visible when looking at the BLEU scores on the

²Our replication achieves 45.0 BLEU on newstest19, competitive with the reference system at 42.7 BLEU.

standard references. While the BLEU score on orig-de does improve by 7.5 points, the BLEU score drops by 2.9 points on the orig-en half. Due to the big gain on the orig-de half, BT also improves the BLEU score on the joint set. The paraphrased references were designed to overcome these kinds of mismatches and they show a gain of 0.5 BLEU points. We can conclude that back-translation helps improve BLEU and BLEUP and we include BT for systems that are optimized for both standard or paraphrased BLEU scores.

4.2.3 Fine-Tuning

Similar to (Ng et al., 2019), we fine-tuned our back-translated model on a concatenation of previous WMT testsets (newstest{2013,2015,2016,2017}) and the clean in-domain news-commentary corpus. In total, we fine-tuned the model on 330k sentences. We kept all model parameters the same (batch size, learning rate) and continued training on the fine-tuned data for one epoch. The BLEU scores on the standard references suggest a small improvement of 0.3 BLEU on the joint test set. Interestingly, the improvement is visible on the orig-en half by 0.7 points while the BLEU scores on orig-de actually drop by 1.7 points. Nevertheless, BLEUP does improve by 0.5 points, suggesting that fine-tuning is especially helpful when measuring scores with paraphrased references. Despite the small gain on standard references, we include fine-tuning in both our optimized systems.

4.2.4 Ensemble

Combining different predictions is a standard approach in MT to boost BLEU scores. We run ensemble decoding with 4 previously built models. In addition to using the 3 models described in Section 4.2.1, 4.2.2, and 4.2.3, we build a second fine-tuned model with the same approach, but different initialization.

Although ensemble decoding improves the performance on our standard references by up to 1.9 BLEU points, the quality is rated as lower by 0.3 BLEU points on the paraphrased references. We suspect that using an ensemble for decoding favors common, average language by promoting target spans where all systems agree. Paraphrase translations actually downweight the importance of this language, which seems important for agreeing with human judgments (Freitag et al., 2020). This promotion of average language and monotonic translation may explain the effectiveness of ensembling

	newstest2018 (dev)				newstest2019 (test)			
	joint	orig-de	orig-en	orig-en.p	joint	orig-de	orig-en	orig-en.p
(1) bitext	46.0	38.8	50.6	12.8	38.5	34.9	40.9	12.1
(2) + CDS	46.1	39.4	50.5	13.4	39.6	35.6	42.3	12.6
(3) + BT	47.2	45.3	47.7	13.6	40.9	43.1	39.4	13.1
(4) + Fine tuning	47.7	43.6	49.2	13.8	41.2	41.3	41.1	13.6
(5) + Ensemble of 4	49.8	45.4	52.1	13.7	43.1	42.1	43.6	13.3
+ reranking of (5) (opt on BLEU)	50.7	44.8	53.9	13.8	43.4	41.2	45.0	13.4
+ reranking of (4) (opt on BLEUP)	47.1	45.9	47.1	14.7	41.6	44.0	39.8	13.7

Table 2: BLEU scores for WMT 2019 English→German. The *joint* sets combine *orig-en* and *orig-de* subsets. The *orig-en.p* sets use paraphrased references instead of standard references. Our experiments compared *newstest2018.joint* and *newstest2018.orig-en.p* for system tuning. The standard newstest2018 and newstest2019 sets are *newstest2018.joint* and *newstest2019.orig-en*, respectively.

only for standard reference BLEU. Similar to the WMT 2019 winning submission, we include the ensemble approach in our system that is optimized on the joint BLEU scores. However, we do not include it in our system optimized on BLEUP.

4.3 Reranking

Finally, we extend the noisy-channel approach (Yee et al., 2019) which consists of re-ranking the top-50 beam search output of either the ensemble model (when tuned for BLEU) or the fine-tuned model (when tuned for BLEUP). Instead of using 4 features—forward probability, backward probability, language model and word penalty—we use 11 forward probabilities, 10 backward probabilities and 2 language model scores. Different to (Ng et al., 2019), we did not pick the re-ranking weights through random search, but used MERT (Och, 2003) for efficient tuning.

The 11 different forward translation scores come from different English→German NMT models that are replicas of the previous described models (Section 4.2.1, 4.2.2, and 4.2.3). The 10 backward translation scores come from the same approaches, but trained in the reverse direction. These 21 NMT model scores are combined with 2 language model (LM) scores. The first LM is trained on the German monolingual NewsCrawl data, while the second LM is trained on forward-translated English NewsCrawl data. The first LM should assign high scores to genuine German text, while the second LM should assign high scores to translationese German originating from English.

We first reranked the 50-best list generated by the ensemble model with MERT on newstest2018. Similar to the original WMT 2019 submission, the

BLEU scores on the joint and orig-en set increase. This reranked output corresponds to our opt-on-BLEU model. Next, we reranked the 50-best list generated by the fine-tuned model with MERT on newstest2018.orig-en with paraphrased references. This led to further small increases in BLEUP, and corresponds to our opt-on-BLEUP model.

In summary, optimizing on BLEUP leads us to keep back-translation, even though evaluation with standard English-original references would have us drop it, and also leads us to drop the ensembling step. Rescoring using MERT weights learned with BLEU or BLEUP further separates the systems according to these metrics.

5 Analysis

This section confirms the results from the previous section with additional references for newstest2019 and illustrates the behaviour of our systems on individual sentences.

5.1 Alternative Reference Translations

Freitag et al. (2020) released an additional standard reference translation (AR) and two ‘paraphrase as-much-as-possible’ reference translations for newstest2019 (WMT.p and AR.p). We used WMT.p in all our above experiments; here we report BLEU scores for all four available reference translations in table 3. The BLEU improvements between the two standard reference translations agree perfectly. Similarly, the BLEUP improvements between the two paraphrased references also coincide. This indicates that by optimizing on BLEU or BLEUP we have not somehow overfit to a specific set of reference translations or their paraphrases, but instead

have molded our model to better match a style of reference translation.

5.2 Translation Examples

This section presents translation examples from our two differently optimized systems in Table 4. The first 3 examples show sentences where opt-on-BLEUP has higher translation quality than opt-on-BLEU. One observation of (Freitag et al., 2020) was that BLEU scores calculated on standard references prefer monotonic translations. This is visible in our first translation example, where opt-on-BLEU incorrectly translates the saying *Tomorrow’s a different beast* into *Morgen ist ein anderes Biest*, using an inappropriately monotonic strategy. On the other hand, the opt-on-BLEUP system captures the meaning of the source sentence and generates a valid translation.

Another drawback of standard reference BLEU is the preference for literal translation. This is visible in our second example where the word *cap* is translated into *Kappe* and *tip* into *kippen*. Both are valid word-by-word translations, but do not make much sense in this context. The third example is another example of the monotonic translation style of a regular tuned system. The opt-on-BLEU translation is an incorrect word-by-word translation. The opt-on-BLEUP system is able to introduce a German natural sentence structure and generate a flawless translation.

The last translation example is a loss for the paraphrased-tuned system and demonstrates that sometimes a more literal translation can be better. Even though the word *run* can be translated into *Ansturm*, it is not appropriate in this context and the simpler translation *Lauf* is correct.

5.3 Matched n-grams

The BLEU scores calculated on the two different references yield different conclusions. BLEU on standard references evaluated opt-on-BLEU higher by more than 5 BLEU points. BLEUP came to a different conclusion and gave a higher score to opt-on-BLEUP. In this section, we look at the n-grams that contributed most to these different outcomes. Those that contribute most to the difference in BLEU across the two systems are:

- **Er sagte, dass** (*He said that*)
- **, sagte er der** (*, he said the*)
- **stellte fest, dass** (*noted that*)

These are all generic, high-frequency n-grams. They are crucial for attaining high BLEU scores, and tend to appear in translations that employ the same structure as the source sentence. In contrast, the n-grams that contribute most to the difference in BLEUP are:

- **Menschen ums Leben kamen** (*humans died*)
- **Grossbritannien keine Steuern zahlen** (*Great Britain pay no tax*)
- **von BBC Scotland** (*from BBC Scotland*)

These are much less frequent sequences with more semantic content.

6 Conclusions

Prior work has shown that BLEU measured on paraphrased references (BLEUP) has better correlation with human evaluation than BLEU measured on regular references (BLEU) for the comparison of existing systems (Freitag et al., 2019). Motivated by this finding, we collected a development set of paraphrased references and assessed BLEUP for system development. This allowed us to evaluate if the design choices of a modern neural MT system impact BLEU and BLEUP differently, including tuning a re-ranking noisy channel model to these metrics. Our experiments followed the setup from the winning newstest19 English→German entry at WMT19 (Ng et al., 2019).

For design choices, we observe that BLEUP seems to emphasize the importance of back-translation even when test sets are source original. On the other end, BLEUP seems to de-emphasize the importance of ensembles, as the reliable prediction of common language by ensembles is less rewarded by this metric.

Our tuning experiments led to positive results. In human evaluation, the system tuned on BLEUP showed significant improvements in terms of adequacy and even greater gains in terms of fluency compared to the system tuned on BLEU. Example translations indicate that the model tuned on BLEUP produces noticeably less literal translations. Our experiments also highlight a disconnect between regular BLEU and human evaluation: the system tuned on BLEUP degrades standard BLEU scores by over 5 points, while faring significantly better in human evaluation. Paraphrased automatic evaluation therefore seems to be a promising proxy

	newstest2019			
	WMT (orig-en)	AR (orig-en)	WMT.p (orig-en.p)	AR.p (orig-en.p)
(1) bitext	40.9	32.2	12.1	12.0
(2) + CDS	42.3	34.2	12.6	12.3
(3) + BT	39.4	33.6	13.1	13.0
(4) + Fine tuning	41.1	35.5	13.6	13.4
(5) + Ensemble of 4	43.6	36.0	13.3	13.0
+ reranking of (5) (opt-on-BLEU)	45.0	36.7	13.4	13.1
+ reranking of (4) (opt-on-BLEUP)	39.8	34.4	13.7	13.5

Table 3: BLEU scores for English→German newstest2019 for the additional references from (Freitag et al., 2020).

source	Tomorrow’s a different beast.
opt on BLEU	Morgen ist ein anderes Biest.
opt on BLEUP	Morgen ist alles anders.
source	You have to tip your cap.
opt on BLEU	Sie müssen Ihre Kappe kippen.
opt on BLEUP	Man muss den Hut ziehen.
source	He averaged 5.6 points and 2.6 rebounds a game last season.
opt on BLEU	Er durchschnittlich 5,6 Punkte und 2,6 Rebounds ein Spiel in der vergangenen Saison.
opt on BLEUP	In der vergangenen Saison erzielte er im Schnitt 5,6 Punkte und 2,6 Rebounds pro Spiel.
source	Thirty-two percent supported such a run.
opt on BLEU	32 Prozent unterstützten einen solchen Lauf.
opt on BLEUP	32 Prozent sprachen sich für einen solchen Ansturm aus.

Table 4: Example output for English→German for systems optimized on standard BLEU or BLEUP. Translations for opt-on-BLEU tend to be more literal, and adhere closely to the source sentence structure.

for human evaluation when making design choices for MT systems.

This research opens the question of whether these results can be confirmed over a wide range of language pairs. We also hope to achieve further improvements by refining the paraphrased evaluation protocol.

References

- 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, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Mona Baker. 1993. Corpus Linguistics and Translation Studies: Implications and Applications. *Text and technology: in honour of John Sinclair*, pages 233–252.
- Satanjeev Banerjee and Alon Lavie. 2005. [METEOR: An Automatic Metric for MT Evaluation with Improved Correlation with Human Judgments](#). In *Proceedings of the ACL workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, pages 65–72.
- Loïc Barrault, Ondřej Bojar, Marta R. Costa-jussà, Christian Federmann, Mark Fishel, Yvette Graham, Barry Haddow, Matthias Huck, Philipp Koehn, Shervin Malmasi, Christof Monz, Mathias Müller, Santanu Pal, Matt Post, and Marcos Zampieri. 2019. [Findings of the 2019 Conference on Machine Translation \(WMT19\)](#). In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 1–61, Florence, Italy. Association for Computational Linguistics.
- Nikolay Bogoychev and Rico Sennrich. 2019. [Domain, Translationese and Noise in Synthetic Data for Neural Machine Translation](#). *arXiv preprint arXiv:1911.03362*.
- Ondřej Bojar, Rajen Chatterjee, Christian Federmann, Yvette Graham, Barry Haddow, Matthias Huck, Antonio Jimeno Yepes, Philipp Koehn, Varvara Logacheva, Christof Monz, Matteo Negri, Aurélie Névéal, 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*, pages 131–198, Berlin, Germany. Association for Computational Linguistics.
- Ondřej Bojar, Christian Federmann, Mark Fishel, Yvette Graham, Barry Haddow, Philipp Koehn, and Christof Monz. 2018. *Findings of the 2018 Conference on Machine Translation (WMT18)*. In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pages 272–303, Belgium, Brussels. Association for Computational Linguistics.
- Nikhil Buduma and Nicholas Locascio. 2017. *Fundamentals of deep learning: Designing next-generation machine intelligence algorithms*. “O’Reilly Media, Inc.”.
- Marina Fomicheva and Lucia Specia. 2016. *Reference bias in monolingual machine translation evaluation*. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 77–82, Berlin, Germany. Association for Computational Linguistics.
- 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*, pages 34–44, Florence, Italy. Association for Computational Linguistics.
- Markus Freitag, David Grangier, and Isaac Caswell. 2020. *BLEU might be Guilty but References are not Innocent*. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*, Online. Association for Computational Linguistics.
- Jonas Gehring, Michael Auli, David Grangier, and Yann Dauphin. 2017. *A convolutional encoder model for neural machine translation*. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 123–135, Vancouver, Canada. Association for Computational Linguistics.
- Yvette Graham, Barry Haddow, and Philipp Koehn. 2019. *Translationese in machine translation evaluation*.
- Yifan He and Andy Way. 2010. *Metric and reference factors in minimum error rate training*. *Machine Translation*, 24(1):27–38.
- 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.
- Philipp Koehn. 2010. *Statistical Machine Translation*. Cambridge University Press.
- 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.
- David Kurokawa, Cyril Goutte, and Pierre Isabelle. 2009. *Automatic detection of translated text and its impact on machine translation*. In *Proceedings of MT-Summit XII*, pages 81–88.
- Samuel Lüubli, Sheila Castilho, Graham Neubig, Rico Sennrich, Qinlan Shen, and Antonio Toral. 2020. *A set of recommendations for assessing human-machine parity in language translation*. *Journal of Artificial Intelligence Research*, 67:653–672.
- Gennadi Lembersky, Noam Ordan, and Shuly Wintner. 2012a. *Adapting Translation Models to Translationese Improves SMT*. In *Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics, EACL ’12*, pages 255–265, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Gennadi Lembersky, Noam Ordan, and Shuly Wintner. 2012b. *Language models for machine translation: Original vs. translated texts*. *Computational Linguistics*, 38(4):799–825.
- Chi-kiu Lo. 2019. *Yisi-a unified semantic mt quality evaluation and estimation metric for languages with different levels of available resources*. In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 507–513.
- Qingsong Ma, Yvette Graham, Timothy Baldwin, and Qun Liu. 2017. *Further investigation into reference bias in monolingual evaluation of machine translation*. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2476–2485, Copenhagen, Denmark. Association for Computational Linguistics.
- Nitin Madnani, Necip Fazil Ayan, Philip Resnik, and Bonnie J Dorr. 2007. *Using paraphrases for parameter tuning in statistical machine translation*. In *Proceedings of the Second Workshop on Statistical Machine Translation*, pages 120–127. 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*, Valencia, Spain.
- Yuval Marton, Chris Callison-Burch, and Philip Resnik. 2009. *Improved statistical machine translation using monolingually-derived paraphrases*. In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing: Volume 1, EMNLP ’09*, pages 381–390, Stroudsburg, PA, USA. Association for Computational Linguistics.

- Anna Mauranen and Pekka Kujamäki. 2004. *Translation universals: Do they exist?*, volume 48. John Benjamins Publishing.
- 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.
- Franz Josef Och. 2003. [Minimum error rate training in statistical machine translation](#). In *Proceedings of the 41st annual meeting of the Association for Computational Linguistics*, pages 160–167.
- 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.
- Matt Post. 2018. [A Call for Clarity in Reporting Bleu Scores](#). *arXiv preprint arXiv:1804.08771*.
- Stefan Riezler and John T Maxwell III. 2005. [On some pitfalls in automatic evaluation and significance testing for mt](#). In *Proceedings of the ACL workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, pages 57–64.
- Parker Riley, Isaac Caswell, Markus Freitag, and David Grangier. 2020. [Translationese as a language in “multilingual” NMT](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7737–7746, Online. Association for Computational Linguistics.
- Aurko Roy and David Grangier. 2019. [Unsupervised Paraphrasing without Translation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6033–6039. Association for Computational Linguistics.
- Larry Selinker. 1972. Interlanguage. *International Review of Applied Linguistics*, pages 209–241.
- Christophe Servan and Holger Schwenk. 2011. [Optimising multiple metrics with mert](#). *The Prague Bulletin of Mathematical Linguistics*, 96(1):109–117.
- Jonathan Shen, Patrick Nguyen, Yonghui Wu, Zhifeng Chen, Mia X. Chen, Ye Jia, Anjuli 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.
- 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*. Cambridge, MA.
- Sara Stymne. 2017. [The Effect of Translationese on Tuning for Statistical Machine Translation](#). In *The 21st Nordic Conference on Computational Linguistics*, pages 241–246.
- Antonio Toral. 2020. [Reassessing Claims of Human Parity and Super-Human Performance in Machine Translation at WMT 2019](#). *arXiv preprint arXiv:2005.05738*.
- Antonio Toral, Sheila Castilho, Ke Hu, and Andy Way. 2018. [Attaining the Unattainable? Reassessing Claims of Human Parity in Neural Machine Translation](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 113–123, Belgium, Brussels. Association for Computational Linguistics.
- 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.
- Wei Wang, Taro Watanabe, Macduff Hughes, Tetsuji Nakagawa, and Ciprian Chelba. 2018. [Denoising neural machine translation training with trusted data and online data selection](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 133–143, Belgium, Brussels. Association for Computational Linguistics.
- Kyra Yee, Nathan Ng, Yann N Dauphin, and Michael Auli. 2019. [Simple and effective noisy channel modeling for neural machine translation](#). *arXiv preprint arXiv:1908.05731*.
- Mike Zhang and Antonio Toral. 2019. [The effect of translationese in machine translation test sets](#). *CoRR*, abs/1906.08069.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2019. [Bertscore: Evaluating text generation with BERT](#). *Arxiv*, 1904.09675.