

# Reproducibility Issues for BERT-based Evaluation Metrics

Yanran Chen<sup>1</sup>, Jonas Belouadi<sup>2</sup>, Steffen Eger<sup>2</sup>

<sup>1</sup>Computer Science Department, Technical University of Darmstadt, Germany

<sup>2</sup>NLLG, Faculty of Technology, Bielefeld University, Germany

<sup>1</sup>yanran.chen@stud.tu-darmstadt.de

<sup>2</sup>{jonas.belouadi, steffen.eger}@uni-bielefeld.de

## Abstract

Reproducibility is of utmost concern in machine learning and natural language processing (NLP). In the field of natural language generation (especially machine translation), the seminal paper of Post (2018) has pointed out problems of reproducibility of the dominant metric, BLEU, at the time of publication. Nowadays, BERT-based evaluation metrics considerably outperform BLEU. In this paper, we ask whether results and claims from four recent BERT-based metrics can be reproduced. We find that reproduction of claims and results often fails because of (i) heavy undocumented preprocessing involved in the metrics, (ii) missing code and (iii) reporting weaker results for the baseline metrics. (iv) In one case, the problem stems from correlating not to human scores but to a wrong column in the csv file, inflating scores by 5 points. Motivated by the impact of preprocessing, we then conduct a second study where we examine its effects more closely (for one of the metrics). We find that preprocessing can have large effects, especially for highly inflectional languages. In this case, the effect of preprocessing may be larger than the effect of the aggregation mechanism (e.g., greedy alignment vs. Word Mover Distance).

## 1 Introduction

Reproducibility is a core aspect in machine learning (ML) and natural language processing (NLP). It requires that claims and results of previous work can independently be reproduced and is a prerequisite to trustworthiness. The last few years have seen vivid interest in the topic and many issues of non-reproducibility have been pointed out, leading to claims of a “reproducibility crisis” in science (Baker, 2016). In the field of evaluation metrics for natural language generation (NLG), the seminal work of Post (2018) has demonstrated how different preprocessing schemes can lead to substantially different results when using the dominant metric at the time, BLEU (Papineni et al., 2002).

Thus, when researchers employ such different preprocessing steps (a seemingly innocuous decision), this can directly lead to reproducibility failures of (conclusions regarding) metric performances.

Even though BLEU and similar lexical-overlap metrics still appear to dominate the landscape of NLG (particular MT) evaluation (Marie et al., 2021), it is obvious that metrics which measure surface level overlap are unsuitable for evaluation, especially for modern text generation systems with better paraphrasing capabilities (Mathur et al., 2020). As a remedy, multiple higher-quality metrics based on BERT and its extensions have been proposed in the last few years (Zhang et al., 2019; Zhao et al., 2019). In this work, we investigate whether these more recent metrics have better reproducibility properties, thus filling a gap for the newer paradigm of metrics. We have good reason to suspect that reproducibility will be better: (i) as a response to the identified problems, recent years have seen many efforts in the NLP and ML communities to improve reproducibility, e.g., by requiring authors to fill out specific check lists.<sup>1</sup> (ii) Designers of novel evaluation metrics should be particularly aware of reproducibility issues, as reproducibility is a core concept of proper evaluation of NLP models (Gao et al., 2021).

Our results are disillusioning: out of four metrics we tested, three exhibit (severe) reproducibility issues. The problems relate to (i) heavy use of (undocumented) preprocessing, (ii) missing code, (iii) reporting lower results for competitors, and (iv) correlating with the wrong columns in the evaluation csv file. Motivated by the findings on the role of preprocessing and following Post (2018), we then study its impact more closely in the second part of the paper (for those metrics making use of it), finding that it can indeed lead to substantial performance differences also for BERT-based met-

<sup>1</sup>E.g., <https://aclrollingreview.org/responsibleNLPresearch/>.

rics. The code for this work is available at <https://github.com/cyr19/Reproducibility>.

## 2 Related Work

Relevant prior work to this work includes BERT-based evaluation metrics (Section 2.1) and reproducibility in NLP (Section 2.2).

### 2.1 BERT-based Evaluation Metrics

In recent years, many strong automatic evaluation metrics based on BERT (Devlin et al., 2018) or its variants have been proposed. It has been shown that those BERT-based evaluation metrics correlate much better with human judgements than traditional evaluation metrics such as BLEU (Papineni et al., 2002). Popular supervised BERT-based evaluation metrics include BLEURT (Selam et al., 2020) and COMET (Rei et al., 2020), which are trained on segment-level human judgements such as DA scores in WMT datasets. Unsupervised BERT-based evaluation metrics such as BERTScore (Zhang et al., 2019), MoverScore (Zhao et al., 2019), BaryScore (Colombo et al., 2021) and XMoverScore (Zhao et al., 2020) do not use training signals, thus potentially may generalize better to unseen language pairs (Belouadi and Eger, 2022). MoverScore, BaryScore, and BERTScore are *reference-based* evaluation metrics. In contrast, *reference-free* evaluation metrics directly compare system outputs to source texts. For MT, popular such metrics are Yisi-2 (Lo, 2019), XMoverScore, and SentSim (Song et al., 2021).

### 2.2 Reproducibility in NLP

Cohen et al. (2018) define **replicability** as the ability to repeat the process of experiments and **reproducibility** as the ability to obtain the same results. They further categorize reproducibility along 3 dimensions: (1) reproducibility of a *conclusion*, (2) reproducibility of a *finding*, and (3) reproducibility of a *value*. In a more recent study, Belz et al. (2021) categorize reproduction studies according to the *condition of the reproduction experiment*: (1) *reproduction under the same condition*, i.e., re-using as similar as possible resources and mimicking the authors' experimental procedure as closely as possible; (2) *reproduction under varied conditions*, aiming to test whether the proposed methods can obtain similar results with some changes in the settings; (3) *multi-test and multi-lab studies*, i.e., reproducing multiple papers using uniform meth-

ods and multiple teams attempting to reproduce the same paper, respectively.

In the first part of this work, our reproductions follow the first type described by Belz et al. (2021), i.e., we adhere to the original experimental setup and re-use the resources provided by the authors whenever possible, aiming at exact reproduction. The second part falls into the second category of reproduction study described by Belz et al. (2021), i.e., to change some settings on purpose to see if comparable results can be obtained.

According to Fokkens et al. (2013) and Wieling et al. (2018), the main challenge for reproducibility is the unavailability of the source code and data. Dakota and Kübler (2017) study reproducibility for text mining. They show that 80% of the failed reproduction attempts were due to the lack of information about the datasets. To investigate the availability of source data, Mieskes (2017) conducted quantitative analyses on the publications from five conferences. They found that though 40% of the papers reported having collected or changed existing data, only 65.2% of them provided the links to download the data; 18% of them were invalid. Similarly, Wieling et al. (2018) assessed the availability of both source code and data of papers from two ACL conferences (2011 and 2016). When comparing 2016 to 2011, the availability of both data and code increased, suggesting a growing trend of sharing resources for reproduction. However, even using the same code and data, they could only recreate identical values for one paper. More recently, Belz et al. (2021) analyzed 34 reproduction studies under the same condition (re-using the original resources when possible) for NLP papers. They found that only a small portion (14.03%) of values could be exactly reproduced and the majority (59.2%) of the reproduced values lead to worse results. Moreover, 1/4 deviations are >5%.

In NLG, Post (2018) attests to the non-comparability of BLEU (Papineni et al., 2002) scores across different papers. He argues that there are four causes. **First**, BLEU is a parameterized approach; he shows that on WMT17 (Bojar et al., 2017), the BLEU score for en-fi, increases by roughly 3% Pearson from changing parameters regarding multiple references. The **second** issue, which is regarded as the most critical, is the use of different preprocessing schemes. Among these, tokenization of the references plays a key role. The **third** problem is that preprocessing details are of-

ten omitted in papers. The **fourth** problem is different versions of datasets, in his case a particular problem with the en-de language pair in WMT14 (Macháček and Bojar, 2014). The reproducibility issue of BLEU has also been verified by Belz et al. (2022), using their novel approach, which is designed to quantify the degree of reproducibility.

### 3 Datasets & Metrics

In our reproduction experiments (Section 4), following Zhang et al. (2019), Zhao et al. (2019) and Colombo et al. (2021), we use WMT15-18 (Stanojević et al., 2015; Bojar et al., 2016, 2017; Ma et al., 2018) for MT evaluation. Besides, we follow Zhao et al. (2019) to use TAC2008<sup>2</sup> and TAC2009<sup>3</sup> for text summarization evaluation, MSCOCO (Guo et al., 2018) for image captioning (IC) evaluation, and BAGEL (Wen et al., 2015) and SFHOTEL (Mairesse et al., 2010) for data-to-text generation (D2T) evaluation. For the reference-free metric SentSim, we will mainly report results on the MLQE-PE dataset (Fomicheva et al., 2020b). In further experiments (Section 5), we consider WMT19 (Ma et al., 2019) for MT as well. The datasets for each NLG task are described in detail in the appendix (Section A.1). For our reproduction attempts, we consider MoverScore, BERTScore, BaryScore, and SentSim.

**Metrics** **MoverScore** measures semantic similarity between reference and hypothesis by aligning semantically similar words and computing the distance between these words using the Word Mover Distance (Kusner et al., 2015). **BERTScore** calculates the cosine similarity (of BERT representations) for each token in the reference with each token in the hypothesis, and uses greedy alignment to obtain the similarity scores between sentences. It has three variants: Recall, Precision, and F1. **BaryScore** computes the Wasserstein distance (i.e., Earth Mover Distance (Rubner et al., 2000)) between the barycentric distribution (Agueh and Carlier, 2011) of the contextual representations of reference and hypothesis to measure the dissimilarity between them. **SentSim** has both reference-free and -based versions; we experiment with its reference-free version in this work, which combines sentence- (based on Reimers and Gurevych (2020)) and word-level models (extending a.o. BERTScore to the

multilingual case) to score a pair of source text and hypothesis.

## 4 Reproduction Attempts

Our main focus will be to reproduce the results on machine translation (MT) reported in Zhang et al. (2019), Zhao et al. (2019), Colombo et al. (2021) and Song et al. (2021).

### 4.1 Reproduction on MT

At first, we examine the three reference-based metrics. MoverScore, BaryScore and BERTScore were all originally evaluated on MT but with different WMT datasets (Stanojević et al., 2015; Bojar et al., 2016, 2017; Ma et al., 2018). Zhang et al. (2019) used WMT18 (Ma et al., 2018) as the main evaluation dataset. Zhao et al. (2019) reported the results on WMT17 (Bojar et al., 2017) for both MoverScore and BERTScore-F1. Colombo et al. (2021) compared their metric BaryScore with MoverScore and BERTScore-F1 on WMT15 (Stanojević et al., 2015) and WMT16 (Bojar et al., 2016). MoverScore claims to outperform BERTScore (which was published earlier on Arxiv), and BaryScore claims to outperform the earlier two.

We evaluate the three metrics with the same BERT model (BERT-base-uncased) on all MT datasets mentioned above, using the reproduction resources provided by the authors of each metric. We also evaluate MoverScore and BaryScore on a BERT model finetuned on NLI (Wang et al., 2018) (as in the original papers). The code and data for reproduction were released on their respective githubs.<sup>4</sup> In our reproduction experiments, we use the metrics with the configurations found in their evaluation scripts or papers. Although Zhao et al. (2019) also reported the results for BERTScore-F1, they did not provide information about the used parameter settings. Similarly, Colombo et al. (2021) evaluated the other two metrics on WMT15-16, but except for the model choice, all other settings are unclear. Moreover, except for Zhang et al. (2019), who explicitly state which results were obtained using IDF-weighting, the authors of the other two approaches did not mention this in their papers.

<sup>4</sup>BERTScore (WMT18): [https://github.com/Tiiger/bert\\_score/tree/master/reproduce](https://github.com/Tiiger/bert_score/tree/master/reproduce); MoverScore (WMT17): <https://github.com/AIPHES/emnlp19-moverscore/tree/master/examples>; BaryScore (WMT15-16): [https://github.com/PierreColombo/nlg\\_eval\\_via\\_simi\\_measures/tree/main/raw\\_score](https://github.com/PierreColombo/nlg_eval_via_simi_measures/tree/main/raw_score).

<sup>2</sup><https://tac.nist.gov/2008/>

<sup>3</sup><https://tac.nist.gov/2009/>

	metric	cs-en	de-en	et-en	fi-en	ru-en	tr-en	zh-en	avg
Reproduced	BaryScore-W	0.360	0.525	0.379	0.280	0.322	0.254	<b>0.252</b>	0.339
	MoverScore-1	0.362	0.529	0.391	<b>0.297</b>	0.338	0.288	0.244	0.350
	BERTScore-F1	<b>0.376</b>	<b>0.538</b>	<b>0.393</b>	0.295	<b>0.341</b>	<b>0.290</b>	0.244	<b>0.354</b>
Reported	BERTScore-F1	0.375	0.535	0.393	0.294	0.339	0.289	0.243	0.353

Table 1: Reproduction: Segment-level Kendall’s  $\tau$  on WMT18 to-English language pairs using the evaluation script provided by Zhang et al. (2019). Reported values are taken from Zhang et al. (2019). Values in green denote reproduced results that are better than the reported. Bold values refer to the best reproduced results with the BERT-base-uncased model.

	metric	cs-en	de-en	fi-en	lv-en	ru-en	tr-en	zh-en	avg
Reproduced	BaryScore-W	0.646	0.652	0.819	0.689	0.697	<b>0.737</b>	0.719	0.709
	MoverScore-1	<b>0.660</b>	<b>0.690</b>	0.806	0.685	<b>0.736</b>	0.732	<b>0.720</b>	<b>0.718</b>
	BERTScore-F1	0.655	0.682	<b>0.823*</b>	<b>0.713</b>	0.725	0.718	0.712	0.718
	MoverScore-1 <sup>+</sup>	0.670*	0.708*	0.821	0.717*	0.738*	0.762*	0.744*	0.737*
Reported	MoverScore-1 <sup>+</sup>	0.670	0.708*	0.835*	0.746*	0.738*	0.762*	0.744*	0.743*
	BERTScore-F1	0.670	0.686	0.820	0.710	0.729	0.714	0.704	0.719

Table 2: Reproduction: Segment-level Pearson’s  $r$  on WMT17 to-English language pairs using evaluation script provided by Zhao et al. (2019). Reported results are cited from Zhao et al. (2019). <sup>+</sup> refers to using the finetuned BERT-based-uncased model on MNLI. Values in green/red denote the reproduced results are better/worse than the reported. Bold values refer to the best results with BERT-base-uncased model. Values with \* denote the best reproduced/reported results.

For unclear metric configurations, we keep them at default. The configurations used here are:

- **BERTScore** We report the reproduced results for BERTScore-F1 that uses BERT-base-uncased, with the default layer 9 of the BERT representation for this model, and IDF-weighting.
- **MoverScore** We report the reproduced results for unigram MoverScore (MoverScore-1) using BERT-base-uncased or its finetuned version on MNLI, the last five layers from BERT aggregated by power means (Rücklé et al., 2018), IDF-weighting, punctuation removal and subwords removal (only keep the first subword in each word).
- **BaryScore** We report the reproduced results for BaryScore<sup>5</sup> that makes use of BERT-base-uncased or its finetuned version on MNLI<sup>6</sup>, the last five layers aggregated using Wasserstein Barycenter, and IDF-weighting.

The metrics with finetuned models are marked with <sup>+</sup> in the following.

**Results** As Table 1 shows, we do not obtain identical results for BERTScore-F1 with Zhang et al.

<sup>5</sup>BaryScore outputs scores relying on Wasserstein distance and those relying on Sinkhorn distance (Cuturi, 2013) together. We report the results for Wasserstein distance, which are denoted as BaryScore-W or Bary-W.

<sup>6</sup>As the authors of BaryScore did not release their finetuned model, we use the NLI model released by the authors of MoverScore for BaryScore.

(2019) on WMT18 to-English language pairs. The maximal deviation between the reported and reproduced results can be seen on the evaluated data for de-en – around 0.003 absolute Pearson’s  $r$ . Most of the deviations are about 0.001. This might be because of tiny differences in rounding strategies and random seeds<sup>7</sup> etc. Further, among the three evaluation metrics, BERTScore-F1 performs best, whereas BaryScore is worst.

Table 2 displays the reproduction results on WMT17 to-English language pairs, leveraging the resources from Zhao et al. (2019). As for MoverScore-1<sup>+</sup>, 5 out of 7 values can be perfectly reproduced (excluding the average value). The unreproducible results on fi-en and lv-en are 0.012 and 0.031 lower than the reported, respectively. On personal communication, the authors told us that they changed the preprocessing for these settings, which is impossible to identify from the released paper or code. We obtain comparable average value for BERTScore-F1 with Zhao et al. (2019) (0.718 vs. 0.719), but the results on individual language pairs differ. Except for fi-en, MoverScore-1<sup>+</sup> correlates better with humans than BERTScore-F1, which is in line with the observation from Zhao et al. (2019). When applying the same BERT model, BaryScore performs slightly worse than

<sup>7</sup>We noted that this evaluation script produces different results every time, but the discrepancy across different runs in the averaged correlations is tiny ( $\sim 0.001$ - $0.002$ ).



	metric	WMT15					WMT16						
		cs-en	de-en	fi-en	ru-en	avg	cs-en	de-en	ru-en	fi-en	ro-en	tr-en	avg
Reproduced	BERT-F	0.750	<b>0.733</b>	0.752	<b>0.745</b>	0.745	<b>0.747</b>	<b>0.640</b>	0.672	0.661	<b>0.723</b>	<b>0.688</b>	0.689
	Mover-1	0.734	0.731	0.743	0.731	0.735	0.740	0.633	<b>0.676</b>	0.655	0.714	0.693	0.685
	Bary-W	<b>0.751</b>	<b>0.731</b>	<b>0.769</b>	0.740	<b>0.748</b>	<b>0.735</b>	<b>0.672</b>	<b>0.659</b>	<b>0.673</b>	<b>0.715</b>	<b>0.709</b>	<b>0.694</b>
	Mover-1 <sup>+</sup>	0.745	0.755*	0.774	0.765*	0.760	0.765*	0.676	0.696*	0.707*	0.742*	0.736	0.720*
	Bary-W <sup>+</sup>	<b>0.753*</b>	<b>0.755</b>	<b>0.787*</b>	<b>0.763</b>	<b>0.764*</b>	<b>0.758</b>	<b>0.700*</b>	<b>0.677</b>	0.706	<b>0.732</b>	<b>0.744*</b>	<b>0.720</b>
Reported	BERT-F	<b>0.743</b>	0.722	0.747	<b>0.740</b>	0.738	0.741	<b>0.653</b>	0.651	0.654	0.702	<b>0.707</b>	0.685
	Mover	0.688	0.718	0.700	0.686	0.698	0.674	0.609	0.644	0.631	0.642	0.661	0.644
	Bary	0.742	<b>0.741</b>	<b>0.766</b>	0.737	<b>0.747</b>	<b>0.742</b>	0.646	<b>0.675</b>	<b>0.671</b>	<b>0.725</b>	0.693	<b>0.692</b>
	Mover <sup>+</sup>	0.710	0.711	0.722	0.673	0.704	0.707	0.624	0.640	0.645	0.664	0.663	0.657
	Bary <sup>+</sup>	0.759*	0.758*	0.799*	0.776*	0.773*	0.766*	0.685*	0.694*	0.702*	0.743*	0.738*	0.721*

Table 3: Reproduction: Segment-level Pearson’s r on WMT15-16 using evaluation script provided by Colombo et al. (2021). Reported values are cited from Colombo et al. (2021). <sup>+</sup> represents using the fine-tuned BERT-base-uncased model on MNLI. Values in green/red denote the reproduced results are better/worse than the reported. Bold values refer to the best results with BERT-base-uncased model. Values with \* denote the best reproduced/reported results.

the other two metrics, except for tr-en.

Table 3 shows the results of the reproduction attempts on **WMT15-16** based on the code and data provided by Colombo et al. (2021). Colombo et al. (2021) reported Pearson, Spearman and Kendall correlation with human ratings; we relegate the reproduction results for Kendall and Spearman correlation, which are similar to those for Pearson correlation, to Section A.2. We are not able to reproduce identical values for any evaluation metric, even for BaryScore. However, the reproduced results for BaryScore and BaryScore<sup>+</sup> are comparable with the reported – around 0.001 Pearson off the reported average values in 3 out of 4 cases. For BERTScore-F1, the reproduced average values are around 0.005 Pearson better than the reported, while for MoverScore/MoverScore<sup>+</sup>, they are about 0.05 Pearson better. Colombo et al. (2021) observed that BaryScore<sup>+</sup> performs best on all language pairs in WMT15-16, which is inconsistent with our observation: MoverScore-1<sup>+</sup> outperforms BaryScore<sup>+</sup> on half the language pairs in these two datasets. With BERT-base-uncased, BaryScore performs best among the three evaluation metrics on these two datasets, however — it achieves the highest correlation on 6 out of 10 language pairs.

**Summary** We can rarely reconstruct identical values but obtained comparable results for the three discussed metrics, even when some of the metric configurations are missing. However, we can overall not reproduce the *conclusions* for three main reasons: (i) authors report lower scores for competitor metrics; (ii) authors selectively evaluate on specific datasets (maybe omitting those for which their metrics do not perform well?); (iii) unlike the au-

thors of BERTScore, the authors of BaryScore and MoverScore do not provide a unique hash, making reproduction of the original values more difficult; (iv) undocumented preprocessing involved.

Following the three reproduction attempts, we cannot conclude that the newer approaches are better than the prior ones (BertScore), as Zhao et al. (2019) and Colombo et al. (2021) claim. We also point out that the three metrics perform very similar when using the same underlying BERT model; using a BERT model fine-tuned on NLI seems to have a bigger impact. This casts some doubt on whether the more complicated word alignments (as used in BaryScore and MoverScore) really have a critical effect.

**SentSim** For reference-free evaluation, Song et al. (2021) use MLQE-PE as their primary evaluation dataset. They compare SentSim to so-called glass-box metrics which actively incorporate the MT system under test into the scoring process (Fomicheva et al., 2020a).

Using the original model configuration, we were able to exactly reproduce the reported scores for all SentSim variants on MLQE-PE. However, we noticed that the provided code for loading the dataset does not retrieve human judgments but averaged log-likelihoods of the NMT model used to generate the hypotheses. Since computing correlations with model log-likelihoods is not meaningful and the z-standardized means of the human judgments that should have been used instead are in an adjacent column of the dataset, we assume that this is an off-by-one error.

Table 4 shows how much fixing this error affects the achieved correlations of BERTScore- and WMD-based SentSim. The baselines were not af-

	metric	en-de	en-zh	ru-en	ro-en	et-en	ne-en	si-en	avg
Fixed	SentSim(BERTScore-based)	6.15	22.23	47.30	78.55	55.09	57.09	51.14	45.36
	SentSim(WMD-based)	3.86	22.62	47.46	77.72	54.60	57.00	49.79	44.72
Reported	SentSim(BERTScore-based)	48.40	42.70	47.50	72.70	55.30	39.20	42.60	49.80
	SentSim(WMD-based)	47.20	42.70	47.60	72.40	55.30	39.00	42.60	49.50
Baselines	D-TP	25.90	32.10	—	69.30	64.20	55.80	46.00	48.90
	D-Lex-Sim	17.20	31.30	—	66.30	61.20	60.00	51.30	47.90

Table 4: Correlations of SentSim on MLQE-PE with model log-likelihoods (Reported), as erroneously done in the official paper, and with human judgments (Fixed). The green and red highlighted results on human judgments indicate that they are better or worse than the corresponding results computed with log-likelihoods. We cite baseline scores from Fomicheva et al. (2020a).

ected by this, as Song et al. (2021) copied their scores from their original papers. Evaluation on human judgments leads to vast score differences on many language pairs. This is especially noticeable for English-German and English-Chinese language pairs, where the correlations achieved with our fixed implementation are substantially worse. This result is much more in line with the findings of related research, which also notes very poor performance for these languages on this dataset (Fomicheva et al., 2020b; Specia et al., 2020). We note that after fixing the error, SentSim falls below the baselines, which it had otherwise outperformed.

## 4.2 Reproduction for other tasks

In Section A.3, we reproduce results for other tasks, especially summarization, image captioning and data-to-text generation, with a focus on MoverScore. We find that we can only reproduce the reported results for summarization, and our results are on average 0.1 Pearson’s  $r$  (-12.8%) down for IC and 0.06 Spearman’s  $\rho$  (-27.8%) down for D2T generation. A reason is that the authors of MoverScore did not release their evaluation scripts and we can only speculate as to their employed preprocessing steps. As long as these are not reported in the original papers or released code, claims regarding performance of the metrics are hard to verify.<sup>8</sup>

## 5 Sensitivity Analysis

In the previous section, we have seen that preprocessing may play a vital role for obtaining state-of-the-art results (at least for some of the metrics). Similar to the case of BLEU (Post, 2018), we now

<sup>8</sup>We cannot rule out the possibility that we made mistakes in our reproduction attempts (e.g., incorrect evaluation scripts or use of datasets), but the unavailability of the resources makes the detection of potential errors difficult.

examine this aspect in more detail.

According to the papers and evaluation scripts, MoverScore uses the following main preprocessing steps (besides those handled by BERT): (i) **Subwords Removal**: discard BERT representations of all subwords except the first. (ii) **Punctuation Removal**: discard BERT representations of punctuations. (iii) **Stopwords Removal**: discard BERT representations of stopwords (only for summarization).<sup>9</sup> The preprocessing steps for BERTScore and BaryScore are only related to lowercasing and tokenization, both of which are handled by BERT. We observe that (i) MoverScore uses much more preprocessing than BERTScore and BaryScore on WMT datasets; (ii) authors may take different preprocessing steps for different tasks, e.g., Zhao et al. (2019) remove stopwords for summarization but not for MT.

Besides preprocessing in a narrower sense, all three considered evaluation metrics use *parameters*. This makes them more flexible, but also complicates reproduction: the difference in one parameter setting can lead to reproduction failure. We study the impact of the parameters related to **IDF-weighting**. IDF-weighting measures how critical a word is to a corpus; thus, it is corpus-dependent. The choice of corpus may lead to deviations of metric scores.

MoverScore is the main experiment object in the remainder. Compared to the other metrics, its authors took more preprocessing steps to achieve the results in their paper, suggesting that it is more likely to obtain uncomparable scores across different users when using MoverScore. We will also investigate the sensitivity of BERTScore to the factors discussed above; we omit BaryScore and SentSim from further consideration. Impor-

<sup>9</sup>In Section A.4, we describe subword removal and stopword and punctuation removal used in MoverScore.

tantly, we move beyond English-only evaluation, as reported in the original MoverScore paper. This will estimate how much uncertainty there is from preprocessing when a user applies MoverScore to a non-English language pair, which requires new IDF corpora, new stopwords lists and may have higher morphological complexity (which is related to choice of subwords).

We use two statistics to quantify the sensitivity of the evaluation metrics. When there are only **two compared values**  $a, b$ , we compute Relative Difference (RD) to reflect the relative performance variation regarding a certain parameter. When there are **more than two compared values**, we compute Coefficient of Variation (CV) to reflect the extent of variability of the metric performance:

$$\text{RD}(a, b) = \frac{a - b}{b} * 100\%,$$

$$\text{CV}(x) = \frac{\sigma}{\mu} * 100\%$$

where  $\sigma$  is the standard deviation and  $\mu$  is the mean of a set of values  $x$ . Larger absolute values of the statistics indicate higher sensitivity of the evaluation metrics.

We only consider MT and summarization evaluation in this part. In each experiment, we only adjust the settings of the tested factors and keep the others default (given in Section A.5). In addition to English (“to-English”), we consider MT evaluation for other 6 languages (“from-English”), for which we use multilingual BERT: Chinese (zh), Turkish (tr), Finnish (fi), Czech (cs), German (de), and Russian (ru). Note that in these cases, we compare a Chinese reference to a Chinese hypothesis and analogously for the other languages.

## 5.1 Stopwords Removal

In this experiment, we consider 4 stopword settings including disabling stopwords removal and applying 3 different stopword lists for the examined languages. We obtain the stopword lists from the resources listed in Section A.6. We inspect the sensitivity of MoverScore-1, MoverScore-2 (MoverScore using bigrams) and BERTScore-F1 to stopword settings, despite that BERTScore does originally not employ stopwords.

For English MT, we calculate CV of the correlations with humans over the 4 stopword settings for each language pair in the datasets, then average CVs over the language pairs in each dataset

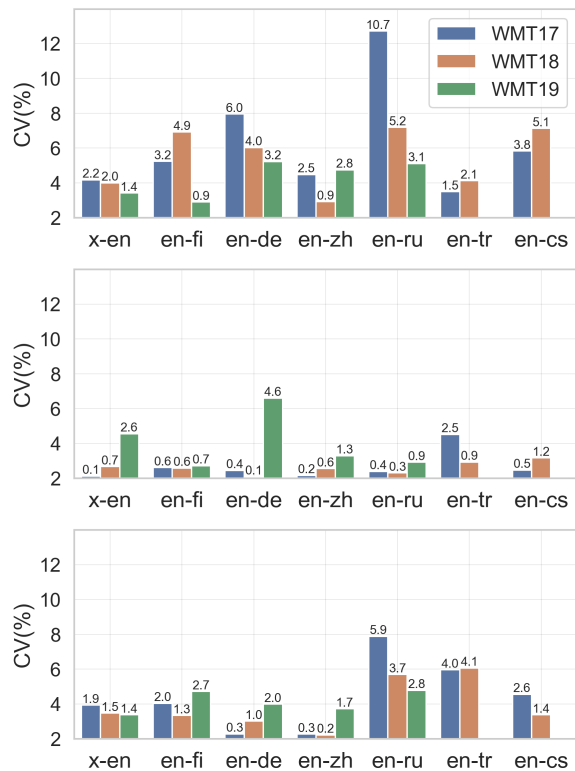


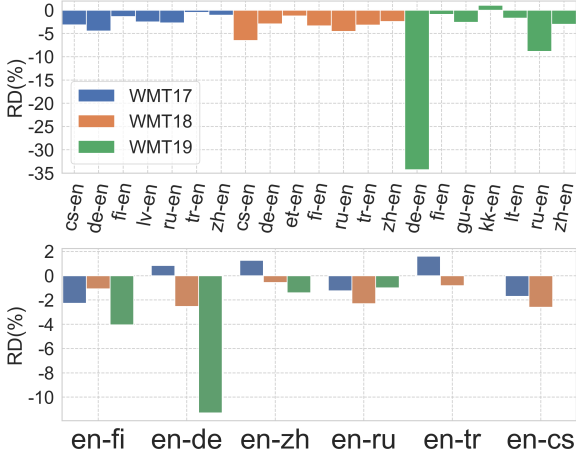
Figure 1: From top to bottom: CV<sub>STOP</sub>, CV<sub>IDF</sub>, CV<sub>SUB</sub>. WMT17-19, segment-level evaluation, MoverScore-1. x-en denotes the average results on all to-English language pairs (where metrics operate on English texts).

to obtain the average CV per dataset. For summarization, we calculate CV of the correlations over the 4 stopword settings for each criterion on each dataset.<sup>10</sup>

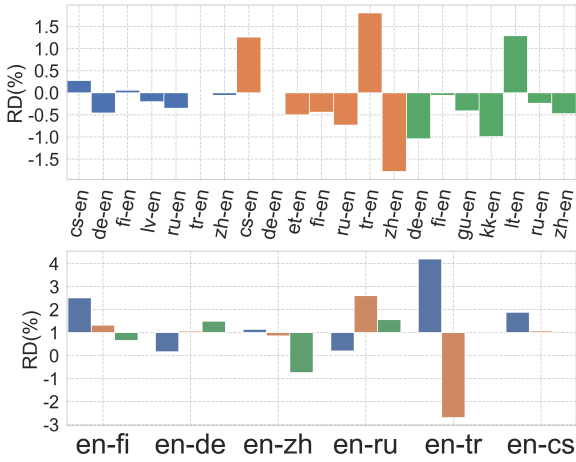
**Results** On segment-level MT, as Figure 1 (top) shows, the sensitivity varies across datasets and languages. Most of the CV<sub>STOP</sub> are in range of 2-4%. This leads to 6-11% absolute variation of the metric performance when the average correlation is, for example, 0.7 (95% confidence interval). For some datasets and languages, the variation is even more pronounced: for example, for Russian on WMT17, the CV<sub>STOP</sub> is above 10%.

Among the examined metrics, MoverScore-2 behaves slightly more sensitively than MoverScore-1, whereas BERTScore-F1 is much more sensitive than MoverScore-1 on Chinese and English. Compared to other tasks, stopwords removal has the largest (but negative) impact in segment-level MT evaluation (cf. Section A.7).

<sup>10</sup>TAC datasets have human judgements according to two criteria: Responsiveness and Pyramid; details are given in Section A.1.



(a) RD(dis,ori)



(b) RD(dis,pr)

Figure 2: RD(dis,ori) (for IDF weighting), RD(dis,pr) (for punctuation removal). WMT17-19, segment-level evaluation, MoverScore-1. The top graphs in (a) and (b) are the results on to-English language pairs (where metrics operate on English texts), whereas the bottom ones are those on from-English language pairs (where metrics operate on texts in other languages).

## 5.2 IDF-weighting

In this test, we first disable IDF-weighting for the evaluation metrics ( $\text{idf}_{\text{dis}}$ ), and compare the metric performance to that when applying original IDF-weighting<sup>11</sup> ( $\text{idf}_{\text{ori}}$ ) by calculating the RD between them. We denote this statistic as RD(dis,ori); negative values indicate  $\text{idf}_{\text{ori}}$  works better and vice versa. Next, to inspect the sensitivity to varying IDF-weighting corpora, we apply IDF-weighting from four randomly generated corpora to the evaluation metrics additionally ( $\text{idf}_{\text{rand}}$ ): each corpus

<sup>11</sup>The original IDF weights for MoverScore are extracted from the reference and hypothesis corpus; those for BERTScore are computed using the reference corpus.

consists of 2k English segments randomly selected from the concatenated corpus of all tested datasets. The corresponding variability of the metric performance is quantified by the CV of the correlations with humans over the 5 IDF-weighting corpus selections ( $\text{idf}_{\text{ori}} + 4 \text{idf}_{\text{rand}}$ ), marked with  $\text{CV}_{\text{IDF}}$ . We examine the sensitivity regarding IDF-weighting of MoverScore-1, MoverScore-2, and BERTScore-F1. Subsequently, we test the IDF-weighting from large-scale corpora ( $\text{idf}_{\text{large}}$ ). These corpora are obtained from Hugging Face Datasets.<sup>12</sup>

**Results** As seen in Figure 2(a), RD(dis,ori) is positive on only one to-English language pair (WMT19 kk-en), but on three from-English language pairs (WMT17 en-de, en-zh, and en-tr). Overall, IDF-weighting is thus beneficial. The maximal performance drops are on WMT19 de-en ( $>35\%$ ) and en-de ( $>10\%$ ), respectively. Most RD(dis,ori) have absolute values  $<5\%$ . This means, suppose the correlation is 0.7, the performance can fall by around 0.035 because of disabling IDF-weighting.

Next,  $\text{CV}_{\text{IDF}}$  for segment-level MT is presented in Figure 1 (middle). In English evaluation, the maximal variation is also caused by the result for de-en in WMT19, where  $\text{idf}_{\text{ori}}$  yields considerably better result than  $\text{idf}_{\text{rand}}$  (0.22 vs. 0.17 Kendall’s  $\tau$ ). While en-de has CV values above 4.5%, most  $\text{CV}_{\text{IDF}}$  are smaller than 1%.

BERTScore-F1 is less sensitive to IDF-weighting than both MoverScore variants. Among the evaluation tasks, the metrics are again most sensitive on segment-level MT, where for English,  $\text{idf}_{\text{ori}}$  works best for MoverScore (even  $\text{idf}_{\text{large}}$  cannot improve its performance), while  $\text{idf}_{\text{rand}}$  and  $\text{idf}_{\text{ori}}$  are almost equally effective for BERTScore-F1 (cf. Section A.9).

## 5.3 Subwords & Punctuation

In this experiment, we evaluate the sensitivity to (i) subword selection and (ii) punctuation removal (PR). (i) In addition to the original two selections of subwords (keeping the first subword and keeping all subwords), we also average the embeddings of the subwords in a word to get the word-level BERT representations. To quantify the sensitivity to subword selection, we calculate CV of the correlations with humans over the 3 subword selections, denoted as  $\text{CV}_{\text{SUB}}$ . (ii) We measure the

<sup>12</sup><https://huggingface.co/datasets>. The corpora used here are listed in Section A.8.



performance change from using to disabling PR by calculating the RD between them, which we denote as  $RD(\text{dis},\text{pr})$ ; negative values indicate MoverScore with PR performs better and vice versa. We inspect the corresponding sensitivity of MoverScore-1.

**Results** Figure 2(b) shows that most  $RD(\text{dis},\text{pr})$  have absolute values  $<1\%$ , while both values for en-tr are  $>3\%$ . Further, the  $CV_{\text{SUB}}$  for segment-level MT is presented in Figure 1 (bottom). The average  $CV_{\text{SUB}}$  over all datasets for most languages are  $<2\%$ , whereas highly inflectional languages such as Turkish and Russian are considerably more sensitive, with average values  $>4\%$ .

Similar as for stopwords and IDF weighting, MoverScore-1 behaves most sensitively on segment-level MT, where the default configuration of PR and subwords, which uses the first subword and removes punctuations, works best for English. However, for other languages, only in 2 out of 16 cases is it best to select the default configuration (cf. Section A.10). As the authors of MoverScore only reported the results on English data, they may thus select an optimal preprocessing strategy only for that case.

## 5.4 Discussion

We summarize the findings from the previous experiments along 4 dimensions.

**Evaluation Tasks:** Among the considered NLG tasks, BERT-based evaluation metrics are more likely to generate inconsistent scores in segment-level MT evaluation. Their sensitivity is less pronounced in system-level MT and summarization. In the latter two cases, average scores are considered, over the translations within one system or over the multiple references. Thus, some of the variation in metric scores will cancel out, leading to a less fluctuating metric performance from varying preprocessing schemes. **Evaluation metrics:** Among the two variants of MoverScore, MoverScore-2 are more sensitive to parameter settings. BERTScore-F1 behaves less sensitively to IDF-weighting than MoverScore while it behaves much more sensitively to stopwords in the evaluation of Chinese and English compared with MoverScore-1. **Languages:** Overall, the considered evaluation metrics have different sensitivities in different languages. Furthermore, highly inflectional languages such as Turkish and Russian as well as German often become “outliers” or obtain extrema in our experiments. **Importance of fac-**

**tors:** Stopwords removal has the largest but mostly negative impact. IDF-weighting positively impacts evaluation metrics in English evaluation but its contribution is much less stable in the evaluation of other languages. MoverScore benefits from subwords and punctuation removal in segment-level MT evaluation for English, but on other tasks or for other languages, no configuration of PR and subword selection consistently performs best.

## 6 Conclusion

We investigated reproducibility for BERT-based evaluation metrics, finding several problematic aspects, including using heavy undocumented preprocessing, reporting lower scores for competitors, selective evaluation on datasets, and copying correlation scores from wrong indices. Our findings cast some doubts on previously reported results and findings, i.e., whether more the complex alignment schemes are really more effective than the greedy alignment of BERTScore. In terms of preprocessing, we found that it can have a large effect depending (a.o.) on the languages and tasks involved. For a fairer comparison between metrics, we recommend to (1) additionally report the results on the datasets that the competitors used, (2) check whether the used versions of the competitor metrics can obtain comparable results as in the original papers, and (3) minimize the role of preprocessing (ideally employing uniform preprocessing across metrics). On the positive side, as authors are nowadays much more willing to publish their resources, it is considerably easier to spot such problems, which may also be one reason why critique papers such as ours have become more popular in the last few years (Beese et al., 2022). In a wider context, our paper contributes to addressing the “cracked foundations” of evaluation for text generation (Gehrmann et al., 2022) and to better understanding their limitations (Leiter et al., 2022).

In the future, we would like to reproduce more recent BERT-based metrics — e.g., with other aggregation mechanisms (Chen et al., 2020), normalization schemes (Zhao et al., 2021), different design choices (Yuan et al., 2021; Chen and Eger, 2022), or metrics that use supervision (Rei et al., 2020; Sellam et al., 2020; Rony et al., 2022) — to obtain a broader assessment of reproducibility issues in this context. We would also like to quantify, at a larger scale, the bias in research induced from overestimating one’s own model vis-à-vis competitor models.

## 7 Limitations

Limitations of our work include (1) a limited number of explored evaluation metrics, (2) a restricted focus on MT only and (3) reliance on author-provided reproduction resources.

(1) Although we did point out very important issues, we only reproduced four metrics. Further, the sensitivity analysis only concerned two evaluation metrics. In the future, we would like to include more reproducibility studies on recent BERT-based evaluation metrics for a broader analysis. It is possible that our particular sample is representative of more severe underlying problems in the community or that it is particularly affected by reproducibility issues.

(2) Our reproduction attempts, with the exception of MoverScore, focused only on MT. For example, the authors of BaryScore also reported results on summarization, IC, and D2T generation, which (for computational costs) we did not consider in this work. While we believe that our findings generalize from MT to other tasks, we did not confirm this expectation experimentally.

(3) Our reproduction attempts were mainly based on the author-provided resources, such as the code and datasets they released, with which we could obtain comparable results in most instances. Nevertheless, we did not investigate their legitimacy, e.g., whether the implementation of the approach is in accordance with the description in its paper or whether the datasets uploaded by the authors are the official ones, etc.

## 8 Acknowledgement

We thank the Ubiquitous Knowledge Processing (UKP) Lab at TU Darmstadt for providing the computing resources. We thank the reviewers for their thoughtful feedback, which improved the final version of the paper. We also thank the authors of MoverScore and BaryScore for their feedback on the initial version of this paper.

## References

Martial Agueh and Guillaume Carlier. 2011. Barycenters in the wasserstein space. *SIAM Journal on Mathematical Analysis*, 43(2):904–924.

Peter Anderson, Basura Fernando, Mark Johnson, and Stephen Gould. 2016. Spice: Semantic propositional image caption evaluation. In *European conference on computer vision*, pages 382–398. Springer.

Monya Baker. 2016. Reproducibility crisis. *Nature*, 533(26):353–66.

Dominik Beese, Begum Altunbacs, Gorkem Guzeler, and Steffen Eger. 2022. Detecting stance in scientific papers: Did we get more negative recently? *ArXiv*, abs/2202.13610.

Jonas Belouadi and Steffen Eger. 2022. Uscore: An effective approach to fully unsupervised evaluation metrics for machine translation. *ArXiv*, abs/2202.10062.

Anya Belz, Shubham Agarwal, Anastasia Shimorina, and Ehud Reiter. 2021. [A systematic review of reproducibility research in natural language processing](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 381–393, Online. Association for Computational Linguistics.

Anya Belz, Maja Popović, and Simon Mille. 2022. Quantified reproducibility assessment of nlp results. *arXiv preprint arXiv:2204.05961*.

Steven Bird, Ewan Klein, and Edward Loper. 2009. *Natural language processing with Python: analyzing text with the natural language toolkit*. " O'Reilly Media, Inc."

Ondřej Bojar, Yvette Graham, Amir Kamran, and Miloš Stanojević. 2016. Results of the wmt16 metrics shared task. In *Proceedings of the First Conference on Machine Translation: Volume 2, Shared Task Papers*, pages 199–231.

Ondřej Bojar, Yvette Graham, and Amir Kamran. 2017. [Results of the wmt17 metrics shared task](#). pages 489–513.

Xi Chen, Nan Ding, Tomer Levinboim, and Radu Soricut. 2020. [Improving text generation evaluation with batch centering and tempered word mover distance](#). In *Proceedings of the First Workshop on Evaluation and Comparison of NLP Systems*, pages 51–59, Online. Association for Computational Linguistics.

Yanran Chen and Steffen Eger. 2022. Menli: Robust evaluation metrics from natural language inference. *arXiv preprint arXiv:2208.07316*.

Kevin Bretonnel Cohen, Jingbo Xia, Pierre Zweigenbaum, Tiffany J. Callahan, Orin Hargraves, Foster R. Goss, Nancy Ide, Aurélie Névél, Cyril Grouin, and Lawrence E. Hunter. 2018. Three dimensions of reproducibility in natural language processing. *LREC ... International Conference on Language Resources & Evaluation : [proceedings]. International Conference on Language Resources and Evaluation*, 2018:156–165.

Pierre Colombo, Guillaume Staerman, Chloe Clavel, and Pablo Piantanida. 2021. Automatic text evaluation through the lens of wasserstein barycenters. *arXiv preprint arXiv:2108.12463*.

- Yin Cui, Guandao Yang, Andreas Veit, Xun Huang, and Serge Belongie. 2018. Learning to evaluate image captioning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5804–5812.
- Marco Cuturi. 2013. Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in neural information processing systems*, 26.
- Daniel Dakota and Sandra Kübler. 2017. Towards replicability in parsing. In *RANLP*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Antske Fokkens, Marieke van Erp, Marten Postma, Ted Pedersen, P. Vossen, and Nuno Freire. 2013. Offspring from reproduction problems: What replication failure teaches us. In *ACL*.
- M. Fomicheva, Shuo Sun, Lisa Yankovskaya, F. Blain, Francisco Guzmán, M. Fishel, Nikolaos Aletras, Vishrav Chaudhary, and Lucia Specia. 2020a. Unsupervised quality estimation for neural machine translation. *Transactions of the Association for Computational Linguistics*, 8:539–555.
- Marina Fomicheva, Shuo Sun, Erick Fonseca, Chrysoula Zerva, Frédéric Blain, Vishrav Chaudhary, Francisco Guzmán, Nina Lopatina, Lucia Specia, and André FT Martins. 2020b. Mlqe-pe: A multilingual quality estimation and post-editing dataset. *arXiv preprint arXiv:2010.04480*.
- Wikimedia Foundation. [Wikimedia downloads](#).
- Yang Gao, Steffen Eger, Wei Zhao, Piyawat Lertvittayakumjorn, and Marina Fomicheva, editors. 2021. *Proceedings of the 2nd Workshop on Evaluation and Comparison of NLP Systems*. Association for Computational Linguistics, Punta Cana, Dominican Republic.
- Sebastian Gehrmann, Elizabeth Clark, and Thibault Selam. 2022. [Repairing the cracked foundation: A survey of obstacles in evaluation practices for generated text](#).
- Mandy Guo, Zihang Dai, Denny Vrandečić, and Rami Al-Rfou. 2020. [Wiki-40B: Multilingual language model dataset](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 2440–2452, Marseille, France. European Language Resources Association.
- Yinuo Guo, Chong Ruan, and Junfeng Hu. 2018. [Meteor++: Incorporating copy knowledge into machine translation evaluation](#). In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pages 740–745, Belgium, Brussels. Association for Computational Linguistics.
- Matthew Honnibal and Ines Montani. 2017. spaCy 2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing. To appear.
- Matt Kusner, Yu Sun, Nicholas Kolkin, and Kilian Weinberger. 2015. From word embeddings to document distances. In *International conference on machine learning*, pages 957–966. PMLR.
- Christoph Leiter, Piyawat Lertvittayakumjorn, M. Fomicheva, Wei Zhao, Yang Gao, and Steffen Eger. 2022. Towards explainable evaluation metrics for natural language generation.
- 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, Florence, Italy. Association for Computational Linguistics.
- Qingsong Ma, Ondřej Bojar, and Yvette Graham. 2018. [Results of the wmt18 metrics shared task: Both characters and embeddings achieve good performance](#). pages 671–688.
- Qingsong Ma, Johnny Wei, Ondřej Bojar, and Yvette Graham. 2019. Results of the wmt19 metrics shared task: Segment-level and strong mt systems pose big challenges. In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 62–90.
- Andrew L. Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, Andrew Y. Ng, and Christopher Potts. 2011. [Learning word vectors for sentiment analysis](#). In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 142–150, Portland, Oregon, USA. Association for Computational Linguistics.
- Matouš Macháček and Ondřej Bojar. 2014. [Results of the WMT14 metrics shared task](#). In *Proceedings of the Ninth Workshop on Statistical Machine Translation*, pages 293–301, Baltimore, Maryland, USA. Association for Computational Linguistics.
- François Mairesse, Milica Gasic, Filip Jurcicek, Simon Keizer, Blaise Thomson, Kai Yu, and Steve Young. 2010. Phrase-based statistical language generation using graphical models and active learning. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, pages 1552–1561.
- Benjamin Marie, Atsushi Fujita, and Raphael Rubino. 2021. [Scientific credibility of machine translation research: A meta-evaluation of 769 papers](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 7297–7306, Online. Association for Computational Linguistics.



- Nitika Mathur, Timothy Baldwin, and Trevor Cohn. 2020. [Tangled up in BLEU: Reevaluating the evaluation of automatic machine translation evaluation metrics](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4984–4997, Online. Association for Computational Linguistics.
- Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. 2016. [Pointer sentinel mixture models](#).
- Margot Mieskes. 2017. A quantitative study of data in the nlp community. In *EthNLP@EACL*.
- Jekaterina Novikova, Ondřej Dušek, Amanda Cercas Curry, and Verena Rieser. 2017. [Why we need new evaluation metrics for NLG](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2241–2252, Copenhagen, Denmark. Association for Computational Linguistics.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *ACL*.
- 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, Brussels, Belgium. Association for Computational Linguistics.
- Ricardo Rei, Craig Stewart, Ana C Farinha, and Alon Lavie. 2020. [Comet: A neural framework for mt evaluation](#).
- Nils Reimers and Iryna Gurevych. 2020. [Making monolingual sentence embeddings multilingual using knowledge distillation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4512–4525, Online. Association for Computational Linguistics.
- Md Rashad Al Hasan Rony, Liubov Kovriguina, Debanjan Chaudhuri, Ricardo Usbeck, and Jens Lehmann. 2022. [RoMe: A robust metric for evaluating natural language generation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5645–5657, Dublin, Ireland. Association for Computational Linguistics.
- Yossi Rubner, Carlo Tomasi, and Leonidas J Guibas. 2000. The earth mover’s distance as a metric for image retrieval. *International journal of computer vision*, 40(2):99–121.
- Andreas Rücklé, Steffen Eger, Maxime Peyrard, and Iryna Gurevych. 2018. [Concatenated power mean word embeddings as universal cross-lingual sentence representations](#). *arXiv*.
- Thibault Sellam, Dipanjan Das, and Ankur Parikh. 2020. [BLEURT: Learning robust metrics for text generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7881–7892, Online. Association for Computational Linguistics.
- Yurun Song, Junchen Zhao, and Lucia Specia. 2021. [SentSim: Crosslingual semantic evaluation of machine translation](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3143–3156, Online. Association for Computational Linguistics.
- Lucia Specia, F. Blain, M. Fomicheva, E. Fonseca, Vishrav Chaudhary, Francisco Guzmán, and André F. T. Martins. 2020. Findings of the wmt 2020 shared task on quality estimation. In *WMT@EMNLP*.
- Miloš Stanojević, Amir Kamran, Philipp Koehn, and Ondřej Bojar. 2015. Results of the wmt15 metrics shared task. In *Proceedings of the Tenth Workshop on Statistical Machine Translation*, pages 256–273.
- Martin Thoma. 2018. [WiLI-2018 - Wikipedia Language Identification database](#).
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2018. [GLUE: A multi-task benchmark and analysis platform for natural language understanding](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355, Brussels, Belgium. Association for Computational Linguistics.
- Tsung-Hsien Wen, Milica Gasic, Nikola Mrksic, Pei-Hao Su, David Vandyke, and Steve Young. 2015. [Semantically conditioned lstm-based natural language generation for spoken dialogue systems](#).
- Martijn Wieling, Josine Rawee, and Gertjan van Noord. 2018. [Squib: Reproducibility in computational linguistics: Are we willing to share?](#) *Computational Linguistics*, 44(4):641–649.
- Weizhe Yuan, Graham Neubig, and Pengfei Liu. 2021. [Bartscore: Evaluating generated text as text generation](#). *Advances in Neural Information Processing Systems*, 34.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2019. [Bertscore: Evaluating text generation with bert](#). *arXiv preprint arXiv:1904.09675*.
- Wei Zhao, Steffen Eger, Johannes Bjerva, and Isabelle Augenstein. 2021. [Inducing language-agnostic multilingual representations](#). In *Proceedings of \*SEM 2021: The Tenth Joint Conference on Lexical and Computational Semantics*, pages 229–240, Online. Association for Computational Linguistics.
- Wei Zhao, Goran Glavaš, Maxime Peyrard, Yang Gao, Robert West, and Steffen Eger. 2020. [On the limitations of cross-lingual encoders as exposed by reference-free machine translation evaluation](#). *arXiv preprint arXiv:2005.01196*.



Wei Zhao, Maxime Peyrard, Fei Liu, Yang Gao, Christian M Meyer, and Steffen Eger. 2019. Moverscore: Text generation evaluating with contextualized embeddings and earth mover distance. *arXiv preprint arXiv:1909.02622*.

## A Appendix

### A.1 Datasets

#### A.1.1 Machine Translation

Each WMT dataset contains evaluation data for different pairs of source and translation languages. Two types of human judgments serve as the golden standard. The first one is Direct Assessment (DA), which contains human scores for each translation. The second one is DArr, which consists of conclusions about one translation being better than another drawn from DA scores. According to [Bojar et al. \(2017\)](#), [Ma et al. \(2018\)](#) and [Ma et al. \(2019\)](#), when there is insufficient amount of DA scores for each individual translation (smaller than 15), DArr is then considered. In this work, We follow the official instructions to calculate correlations of evaluation metrics with DA judgments using absolute Pearson’s  $r$ , and with DArr judgments using Kendall’s  $\tau$ -like formulation proposed by [Bojar et al. \(2017\)](#). On those datasets, DA always serve as the golden truth for system-level evaluation. For segment-level evaluation, WMT18 and WMT19 use DArr, WMT15 and WMT16 rely on DA, and WMT17 uses DA for all to-English and 2 from-English languages pairs (en-ru and en-zh) and DArr for the remaining from-English language pairs.

#### A.1.2 Text Summarization

Each TAC dataset contains several clusters each with 10 news articles. There are more than 50 system and 4 reference summaries with fewer than 100 words for each article. Each system summary receives 4 human judgements according to two criteria: 1) Pyramid, which reflects the level of content coverage of the summaries; and 2) Responsiveness, which measures the response level to the overall quality of linguistic and content of the summaries. The difference between these two datasets is the fact that TAC2008 contains 48 clusters and summaries from 57 systems, while TAC2009 contains 44 clusters and summaries from 55 systems. [Zhao et al. \(2019\)](#) calculated Pearson and Spearman correlation with summary-level human judgments when evaluating MoverScore. In addition, We compute Kendall correlation as well, allowing for a comparison among the three correlations.

#### A.1.3 Image Captioning

Following [Cui et al. \(2018\)](#), [Zhao et al. \(2019\)](#) evaluated MoverScore on the validation set of

MSCOCO, which contains roughly 40k images each with 5 reference and 12 system captions. Besides, there are system-level human judgements about 5 criteria: M1-M5 ([Anderson et al., 2016](#)). In the reproduction experiment, following the experiment setup of [Zhao et al. \(2019\)](#), We calculate Pearson correlation with M1 and M2 scores, which refer to the ratio of captions better or equal to human captions and the ratio of captions indistinguishable from human captions, respectively.

#### A.1.4 Data-to-Text Generation

There are 202 Meaning Representation (MR) instances in BAGEL and 398 MR instances in SFHOTEL datasets. Multiple references and about two system utterances exist for each MR instance. The datasets provide utterance-level human judgments according to 3 criteria: 1) informativeness, which measures how informative the utterance is; 2) naturalness, which refers to the similarity extent between a system utterance and an native speaker-generated utterance; 3) quality, which reflects the fluency and grammar level of a system utterance ([Novikova et al., 2017](#)). In the reproduction experiment, We follow [Zhao et al. \(2019\)](#) to calculate Spearman correlation with utterance-level human judgements about these 3 criteria.

### A.2 Reproduction on WMT15-16

Table 5 and 6 display the reproduced Spearman and Kendall correlations on WMT15 and WMT16.

### A.3 Reproduction of other tasks

[Zhao et al. \(2019\)](#) released the evaluation scripts for WMT17 and TAC2008/2009 and the corresponding datasets on a github<sup>13</sup>. We take them as the resources for reproduction. As for IC and D2T generation evaluation, we write our own evaluation scripts and download those datasets on our own. We obtained MSCOCO, BAGEL, and FSHOTEL datasets from an open question<sup>14</sup> on its Github page, where [Zhao et al. \(2019\)](#) provided the links to download them. Since [Zhao et al. \(2019\)](#) did not provide much information about how they evaluated on MSCOCO, we also inspect the BERTScore paper ([Zhang et al., 2019](#)), where the authors gave details of the evaluation process. As each system caption in MSCOCO has multiple references, it

<sup>13</sup><https://github.com/AIPHES/emnlp19-moverscore>

<sup>14</sup><https://github.com/AIPHES/emnlp19-moverscore/issues/16>

	metric	WMT15					WMT16						
		cs-en	de-en	fi-en	ru-en	avg	cs-en	de-en	ru-en	fi-en	ro-en	tr-en	avg
Reproduced	BERT-F1	<b>0.750*</b>	0.720	0.731	<b>0.712</b>	0.728	<b>0.749</b>	<b>0.643</b>	0.661	<b>0.660</b>	<b>0.706</b>	<b>0.666</b>	<b>0.681</b>
	Mover-1	0.728	<b>0.721</b>	0.722	0.696	0.717	0.737	0.626	<b>0.671</b>	0.650	0.696	0.669	0.675
	Bary-W	0.747	<b>0.705</b>	<b>0.757</b>	0.709	<b>0.730</b>	<b>0.730</b>	<b>0.650</b>	<b>0.649</b>	<b>0.660</b>	<b>0.699</b>	<b>0.671</b>	<b>0.676</b>
	Mover-1 <sup>+</sup>	0.737	0.741*	0.753	0.733	0.741	0.755*	0.662	0.685*	0.700	0.722*	0.702*	0.704*
	Bary-W <sup>+</sup>	<b>0.744</b>	<b>0.734</b>	<b>0.771*</b>	<b>0.734*</b>	<b>0.746*</b>	<b>0.751</b>	<b>0.680*</b>	<b>0.667</b>	<b>0.701*</b>	<b>0.719</b>	<b>0.702</b>	<b>0.703</b>
Reported	BERT-F1	0.735	0.707	0.725	0.705	0.718	0.736	<b>0.646</b>	0.646	0.641	0.676	<b>0.671</b>	0.669
	Mover	0.701	0.694	0.700	0.655	0.688	0.695	0.591	0.628	0.622	0.654	0.640	0.638
	Bary	<b>0.738</b>	<b>0.722</b>	<b>0.745</b>	<b>0.706</b>	<b>0.728</b>	<b>0.743</b>	0.642	<b>0.664</b>	<b>0.664</b>	<b>0.714</b>	<b>0.671</b>	<b>0.683</b>
	Mover <sup>+</sup>	0.711	0.682	0.720	0.647	0.690	0.704	0.607	0.622	0.626	0.660	0.607	0.638
	Bary <sup>+</sup>	0.752*	0.745*	0.787*	0.750*	0.759*	0.762*	0.677*	0.683*	0.695*	0.730	0.705	0.709*

Table 5: Reproduction: Segment-level Spearman correlation on WMT15-16 using evaluation script provided by Colombo et al. (2021). Reported values are cited from Colombo et al. (2021). <sup>+</sup> represents using the fine-tuned bert-base-uncased model on MNLI. Values in green/red denote the reproduced results are better/worse than the reported. Bold values refer to the best results with bert-base-uncased model. Values with \* denote the best reproduced/reported results.

	metric	WMT15					WMT16						
		cs-en	de-en	fi-en	ru-en	avg	cs-en	de-en	ru-en	fi-en	ro-en	tr-en	avg
Reproduced	BERT-F1	<b>0.559*</b>	<b>0.541</b>	0.547	<b>0.532</b>	<b>0.545</b>	<b>0.564</b>	0.474	0.483	<b>0.484</b>	<b>0.520</b>	<b>0.485</b>	<b>0.502</b>
	Mover-1	0.537	0.538	0.540	0.515	0.532	0.552	0.458	<b>0.488</b>	0.475	0.510	0.487	0.495
	Bary-W	0.551	<b>0.528</b>	<b>0.566</b>	<b>0.525</b>	<b>0.543</b>	<b>0.543</b>	<b>0.477</b>	<b>0.471</b>	<b>0.484</b>	<b>0.513</b>	<b>0.491</b>	<b>0.496</b>
	Mover-1 <sup>+</sup>	0.544	0.556*	0.569	0.546*	0.554	0.566*	0.486	0.499*	0.513	0.533*	0.518*	0.519*
	Bary-W <sup>+</sup>	<b>0.549</b>	<b>0.553</b>	<b>0.580*</b>	<b>0.546*</b>	<b>0.557*</b>	<b>0.561</b>	<b>0.499*</b>	<b>0.485</b>	<b>0.516*</b>	<b>0.529</b>	<b>0.518</b>	<b>0.518</b>
Reported	BERT-F1	0.543	0.529	0.541	0.525	0.535	0.555	0.463	0.469	0.470	0.495	0.490	0.490
	Mover	0.520	0.503	0.523	0.469	0.504	0.526	0.442	0.448	0.451	0.482	0.437	0.464
	Bary	<b>0.549</b>	<b>0.531</b>	<b>0.563</b>	<b>0.532</b>	<b>0.544</b>	<b>0.563</b>	<b>0.479</b>	<b>0.481</b>	<b>0.483</b>	<b>0.529</b>	<b>0.514</b>	<b>0.508</b>
	Mover <sup>+</sup>	0.520	0.503	0.529	0.473	0.506	0.534	0.448	0.452	0.458	0.486	0.449	0.471
	Bary <sup>+</sup>	0.569*	0.562*	0.597*	0.567*	0.574*	0.575*	0.500*	0.513*	0.509*	0.545*	0.524*	0.528*

Table 6: Reproduction: Segment-level Kendall correlation on WMT15-16 using evaluation script provided by Colombo et al. (2021). Reported values are cited from Colombo et al. (2021). <sup>+</sup> represents using the fine-tuned bert-base-uncased model on MNLI. Values in green/red denote the reproduced results are better/worse than the reported. Bold values refer to the best results with bert-base-uncased model. Values with \* denote the best reproduced/reported results.

is critical to know how to obtain the caption-level scores. Zhang et al. (2019) clearly state that they use the maximal score for each caption as its final score. According to the evaluation scripts for TAC2008/2009 from Zhao et al. (2019), they averaged the scores for each summary to obtain the summary-level scores, so we assume they might apply the same strategy on MSCOCO. Therefore, we test these two strategies in the reproduction experiment for IC. To check the reliability of our evaluation script, we use it to reproduce the results reported in the BERTScore paper as well. If one can get comparable correlations and the other can not, it may suggest that the authors did extra processing to achieve the results, such as more pre-processing steps on the dataset. The configurations of the evaluation metrics used here are the same as in reproduction attempts on MT.

dataset criteria	BAGEL			SFHOTEL		
	Inf	Nat	Qual	Inf	Nat	Qual
original	0.285	0.195	0.158	0.207	0.270	0.183
reproduced	0.244	0.145	0.092	<b>0.223</b>	0.167	0.065
reproduced (stopwords)	0.230	0.135	0.078	<b>0.208</b>	0.145	0.042

Table 7: Reproduction: Utterance-level Spearman correlations of MoverScore-1 on BAGEL and SFHOTEL datasets. Original results are cited from Zhao et al. (2019). Bold values refer to the reproduced results that are better than the original.

**Results** Overall, we could only reconstruct the identical values for summarization in these 3 reproduction attempts.

Table 7 displays the reproduction results for D2T generation. The reproduced scores with/without stopwords removal go down 0.1/0.08 on average. The maximum deviation is reached in the evaluation of quality on SFHOTEL, down up to 0.14 absolute Spearman correlation. Only two reproduced

metric criteria	MoverScore-1		BERTScore-R	
	M1	M2	M1	M2
original	0.813	0.810	0.834	0.783
reproduced (mean)	0.687	0.674	-	-
reproduced (max)	0.690	0.714	<b>0.851</b>	<b>0.793</b>
reproduced (mean+stopwords)	0.707	0.709	-	-
reproduced (max+stopwords)	0.686	0.718	-	-

Table 8: *Reproduction*: System-level Pearson correlations of MoverScore-1 and BERTScore-R on MSCOCO dataset. Original results are cited from Zhao et al. (2019) and Zhang et al. (2019). Bold values refer to the reproduced results that are better than the original.

values are higher than the original, which are the results for informativeness on SFHOTEL dataset. Besides, the reproduced values also deviate least in the assessment of this criterion on both datasets. As for IC, as Table 8 shows, the correlations for MoverScore are down by over 0.1 across all evaluation setups. Nevertheless, BERTScore-Recall performs even on average 0.03 better in our evaluation. This kind of inconsistency between the reproduction results for these two evaluation metrics may suggest that Zhao et al. (2019) did more preprocessing in the evaluation of IC, which is impossible for others to identify if the authors neither document them nor share the relevant code. In contrast, although different preprocessing schemes were applied to MT and summarization evaluation, it is possible to reproduce most of the values because Zhao et al. (2019) released the evaluation scripts. All of the facts mentioned imply the importance of sharing code and data for reproducibility. However, even with the author-provided code and datasets, there is no guarantee that the results can be perfectly reproduced. The authors may ignore some details of the evaluation setup or metric configurations.

#### A.4 Subwords, Stopwords, Punctuation

**Subword Removal** BERT leverages a subword-level tokenizer, which breaks a word into subwords when the full word is excluded from its built-in vocabulary (e.g., *smarter* → *smart*, *##er*). BERT automatically tags all subwords except the first one with *##*, so we can easily remove them. There are two advantages to doing so. Firstly, it can speed up the system due to the smaller number of embeddings to process. Secondly, it is sometimes equally effective to lemmatization or stemming. E.g., the suffix *er* of the word *smarter* can be removed with this. In some cases, it may keep a less informative part; e.g., the prefix *un* in the word *unhappy*.

#### Stopwords Removal & Punctuation Removal

Both of these two common preprocessing techniques aim to remove less relevant parts of the text data. A typical stopword list consists of function words such as prepositions articles and conjunctions. As an example, MoverScore achieves a higher correlation with human judgments when removing stopwords on text summarization.

#### A.5 Default configuration of evaluation metrics

- **MoverScore** For English evaluation, we use the released version of MoverScore, which makes use of 1) BERT base uncased model finetuned on MNLI dataset, 2) the embeddings of the last five layers aggregated by power means, 3) punctuation removal and the first subword, and 4) IDF-weighting from references and hypotheses separately. We disable stopwords removal in the whole experiment except stopwords tests. For other languages, we replace the model with multilingual BERT base uncased, to keep in line with English evaluation.
- **BERTScore** For English evaluation, we use BERTScore incorporating with BERT base uncased model, the default layer 9, and IDF-weighting from the references. For other languages, similar to MoverScore, we replace the model with multilingual BERT base uncased model.

#### A.6 Stopword lists

For English, the first stopword list is obtained from the Github repository of MoverScore<sup>15</sup>, which contains 153 words. Since users may first choose existing stopword lists from popular libraries, we consider the stopword lists from NLTK (Bird et al., 2009) and SpaCy (Honnibal and Montani, 2017), which consist of 179 and 326 words, respectively.

We obtain the stopword lists for other languages from:

- I. NLTK (Bird et al., 2009);
- II. SpaCy (Honnibal and Montani, 2017);
- III. a Github repository containing stopword lists for many languages;<sup>16</sup>

<sup>15</sup><https://github.com/AIPHES/emnlp19-moverscore/blob/master/examples/stopwords.txt>

<sup>16</sup><https://github.com/orgs/stopwords-iso/repositories?type=all>



Metric	Segment-level				System-level			
	WMT17- $r$	WMT18- $\tau$	WMT19- $\tau$	AVG	WMT17- $r$	WMT18- $r$	WMT19- $r$	AVG
Mover-1	2.18%	2.00%	1.42%	1.87%	0.44%	0.20%	0.12%	0.25%
Mover-2	2.09%	2.04%	1.99%	2.04%	0.14%	0.16%	0.20%	0.17%
BERT-F1	8.74%	8.18%	6.24%	7.72%	0.16%	0.48%	0.25%	0.30%

Table 9:  $CV_{STOP}$  on WMT17-19 to-English language pairs.

IV. a dataset on Kaggle containing stopword lists for many languages <sup>17</sup>;

V. a Github repository containing Chinese stopword lists <sup>18</sup>;

VI. a web containing stopword lists for many languages <sup>19</sup>.

Below are the size of each stopword list and its resource:

- tr: 551(II); 53(I); 504(III);
- de: 543(II); 231(I) 620(III);
- ru: 264(II); 151(I) 556(III);
- cs: 423(III); 405(VI) 256(IV);
- fi: 747(VI); 847(III) 229(IV);
- zh: 747(V); 1891(II) 794(III);

### A.7 Other results for stopwords

Table 9 and Figure 3 display the  $CV_{STOP}$  in English evaluation. We can observe that: (i) Among the three evaluation metrics, MoverScore-1 is least sensitive to stopwords removal, while BERTScore-F1 behaves most sensitively. (ii) The metrics are most sensitive in segment-level MT evaluation among the examined evaluation tasks. (iii) Kendall’s  $\tau$  varies most with changing stopword settings, while Pearson is least sensitive. In other language environment, we can also observe that the metrics are more sensitive at segment-level than at system-level (Figure 8, 10 and 12 (top)). Further, except for Chinese and English, where BERTScore-F1 behaves much more sensitively than MoverScore-F1, the difference between their sensitivity is less pronounced (see Figure 1 and 10 (top)).

<sup>17</sup><https://www.kaggle.com/heeraldedhia/stop-words-in-28-languages>

<sup>18</sup><https://github.com/goto456/stopwords>

<sup>19</sup><https://countwordsfree.com/stopwords/>

metric dataset	BERTScore-F1			MoverScore-1		
	WMT17	WMT18	WMT19	WMT17	WMT18	WMT19
en	0	0	0	0	0	0
zh	0	0	0	0	0	0
de	0	0	0	0	0	0
ru	0	0	0	0	0	0
fi	0	0	0	0	0	229
cs	0	0	-	0	0	-
tr	504	551	-	504	551	-

Table 10: Distribution of the best stopword settings for **all tested languages** in **segment-level MT** evaluation. Values indicate the size of the stopword lists.

Figure 4 illustrates the distribution of the best stopword settings for **English**. In **segment-level MT** evaluation (Figure 4(a)), there is only one case that the best result is achieved by removing stopwords, which takes place on MoverScore-1. In contrast, the best stopword lists for **system-level MT** evaluation can be any of the settings for all evaluation metrics (Figure 4(b)). However, in about 50% of the test cases, MoverScore still performs best when disabling stopwords removal. In **Pyramid** evaluation (Figure 4(c)), MoverScore-1 achieves the best results using the original stopword list for all test cases, whereas disabling stopwords removal is still the best choice for MoverScore-2 and BERTScore-F1. In the evaluation of **Responsiveness** (Figure 4(d)), two cases (33.3%) can be seen that MoverScore-1 applying the original stopword list performs best; this happens only once on MoverScore-2 (16.7%). BERTScore-F1 never benefit from stopwords removal on all evaluation tasks.

Further, in Table 10, we present the best stopword setting for **all examined languages** in **segment-level MT** evaluations. Except Finnish and Turkish, disabling stopwords removal is always the best choice for all other languages. For Finnish, only on one dataset, MoverScore-1 performs better using stopwords removal, whereas, for Turkish, both evaluation metrics achieve the best performance applying the same stopword lists. The reason might be that both Turkish and Finnish belong to agglutinative languages, and those lan-

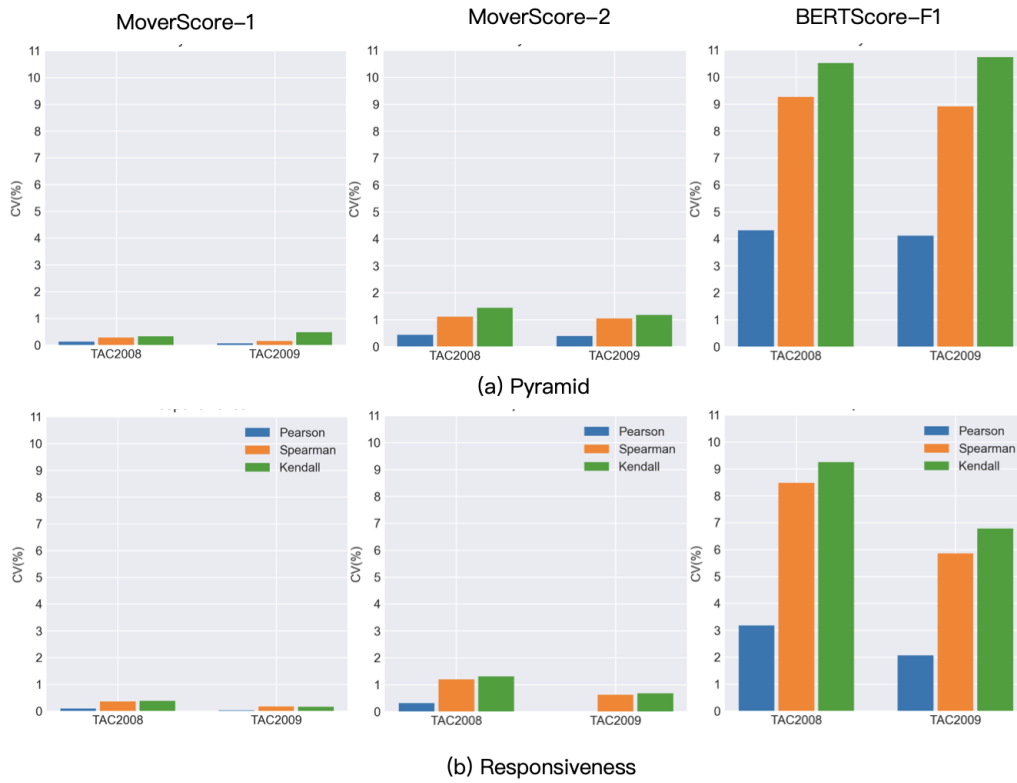


Figure 3:  $CV_{STOP}$  on TAC2008-2009.

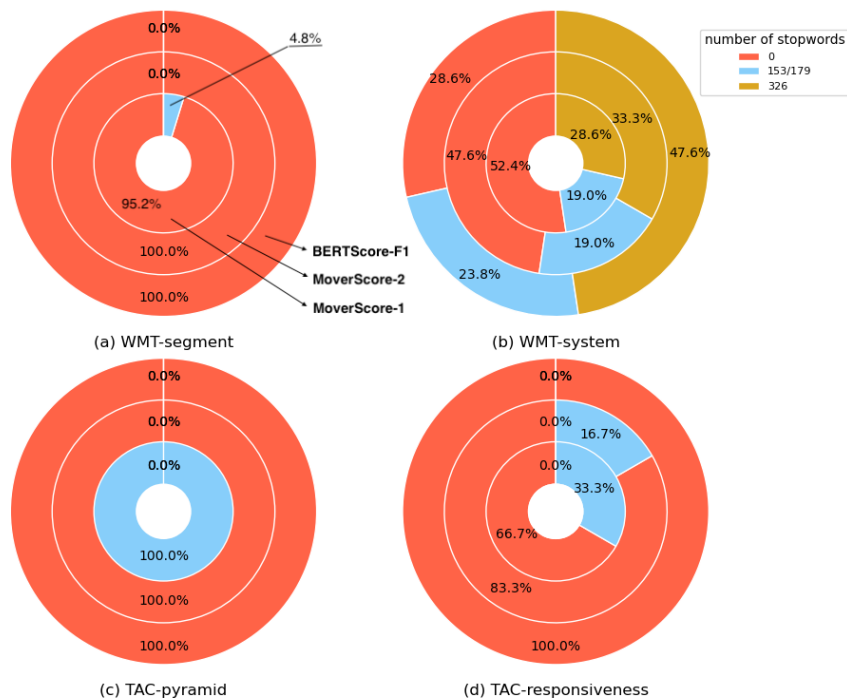


Figure 4: Distribution of the best stopword setting of each evaluation metric on each evaluation task for **English**. The rings from the inside to the outside represent MoverScore-1, MoverScore-2 and BERTScore-F1. For MT, each language pair in WMT datasets is regarded as a test case, resulting in 21 test cases (3 datasets times 7 language pairs). For summarization tasks, each type of correlation is regarded as a test case for each criterion, resulting in 6 test cases (3 correlations times 2 datasets). The MoverScore (153) and SpaCy (179) stopword lists yield exactly the same results.

Metric	Segment-level				System-level			
	WMT17- $r$	WMT18- $\tau$	WMT19- $\tau$	AVG	WMT17- $r$	WMT18- $r$	WMT19- $r$	AVG
Mover-1	0.13%	0.67%	2.56%	1.12%	0.05%	0.06%	0.21%	0.11%
Mover-2	0.78%	1.19%	3.84%	1.94%	0.25%	0.17%	0.49%	0.30%
BERT-F1	0.20%	0.32%	0.41%	0.31%	0.05%	0.02%	0.08%	0.05%

Table 11:  $CV_{IDF}$  for WMT17-19 to-English language pairs.

Metric	Pyramid						Responsiveness					
	TAC2008			TAC2009			TAC2008			TAC2009		
	$r$	$\rho$	$\tau$	$r$	$\rho$	$\tau$	$r$	$\rho$	$\tau$	$r$	$\rho$	$\tau$
Mover-1	0.11%	0.15%	0.40%	0.08%	0.04%	0.43%	0.13%	0.27%	0.31%	0.14%	0.33%	0.37%
Mover-2	0.13%	0.23%	0.27%	0.11%	0.31%	0.35%	0.19%	0.37%	0.39%	0.13%	0.42%	0.46%
BERT-F1	0.19%	0.42%	0.51%	0.06%	0.24%	0.34%	0.21%	0.32%	0.31%	0.17%	0.21%	0.19%

Table 12:  $CV_{IDF}$  for TAC2008-2009.

guages tend to have a high rate of affixes or morphemes per word, which means there may exist more noise in word embeddings.

Corpora	WMT18-AVG	WMT19-AVG
ORI	<b>0.355</b>	<b>0.333</b>
Wili_2008(117500)	0.349	0.323
Wikipedia(100000)	0.350	0.320
Wikipedia(1000000)	0.351	0.320
Wikipedia(2500000)	0.351	0.320
Wikipedia(5000000)	0.350	0.320
Wikipedia(7500000)	0.351	0.320
Wikipedia(10000000)	0.351	0.320
IMDB_train(25000)	0.347	0.323
Wikitext(23767)	0.350	0.324
Wiki40b(2926536)	0.347	0.324

Table 13: Average segment-level Kendall correlation of MoverScore-1 using  $idf_{large}$  with human judgements in WMT18-19 to-English language pairs. Bold values refer to the best results. Number in bracket represents the number of documents in this corpus.

## A.8 IDF Corpora

- **Wikipedia**<sup>20</sup> (Foundation) Wikipedia dataset contains clean full articles of Wikipedia pages but with many non-content segments such as citations, links and so on. Due to memory limit, we can only test a few segments in this dataset.
- **Wiki40b**<sup>21</sup> (Guo et al., 2020) This dataset

<sup>20</sup><https://huggingface.co/datasets/wikipedia>

<sup>21</sup><https://huggingface.co/datasets/wiki40b>

aims at entity identification task, and is cleaned up by excluding ambiguity and non-entity pages from Wikipedia, and non-content and structured part from each page.

- **WikiText**<sup>22</sup> (Merity et al., 2016) This is a language modelling dataset, containing texts extracted from the set of verified good and featured articles on English Wikipedia.
- **Wili\_2008**<sup>23</sup> (Thoma, 2018) The goal of this dataset is to train and test language identification models. It contains short paragraphs of many languages from Wikipedia.
- **IMDB**<sup>24</sup> (Maas et al., 2011) This dataset contains movie reviews and their sentiment label, aiming at binary sentiment classification for English data.

## A.9 Other results for IDF-weighting

As shown in Figure 5 and 6, in English evaluation, the metric performance of the three evaluation metrics drop most from disabling IDF-weighting in segment-level MT evaluation, where the varying IDF corpora also have the largest impact among the examined evaluation tasks (see Table 11 and 12). Among the three metrics, BERTScore-F1 is least sensitive to IDF-weighting, to which  $idf_{ori}$  and  $idf_{rand}$  are almost equally effective, whereas  $idf_{ori}$  yields considerably better results than  $idf_{rand}$  for

<sup>22</sup><https://huggingface.co/datasets/wikitext#wikitext-2-raw-v1-1>

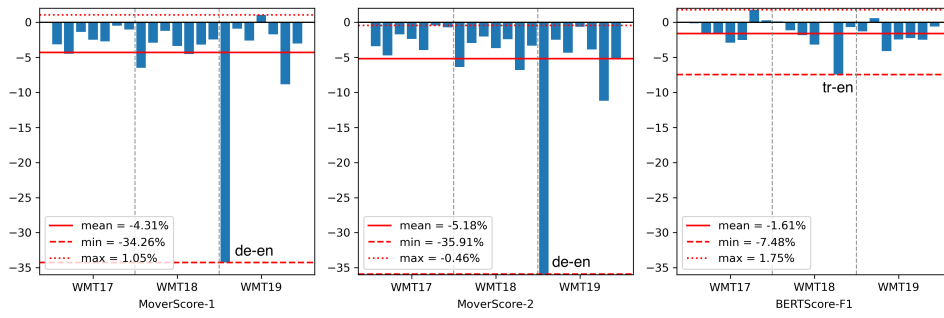
<sup>23</sup>[https://huggingface.co/datasets/wili\\_2018](https://huggingface.co/datasets/wili_2018)

<sup>24</sup><https://huggingface.co/datasets/imdb>

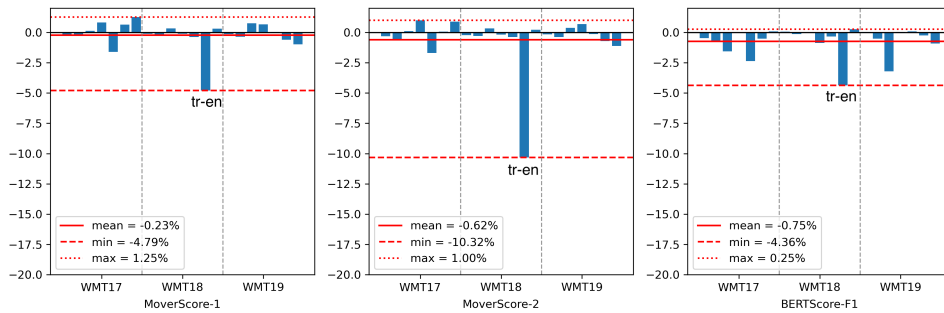
MoverScore-1/2; MoverScore-2 behaves slightly more sensitively than MoverScore-1 (see Figure 7). Moreover, unlike in English evaluation, the contribution of IDF-weighting seems less stable for other languages (see Figure 2(a) and 11).

Further, Table 13 presents the results for  $\text{idf}_{\text{large}}$  in English evaluation. First, the size of those corpora is much larger than the original corpora, but MoverScore still performs better with original IDF-weighting. Secondly, the results for Wikipedia shows that the metric performance does not enhance with the increasing size of IDF corpora. Thirdly, although those corpora contain articles in many domains, they do not provide more applicable IDF-weighting neither. In conclusion, no IDF-weighting from large-domain and large-scale corpora works as well as the original IDF-weighting in segment-level MT evaluation for English, where MoverScore-1 behaves most sensitively to IDF.





(a) Segment-level

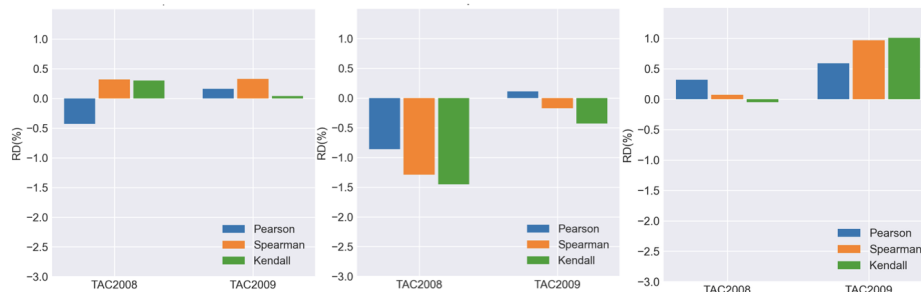


(b) System-level

Figure 5: RD(dis,ori), WMT17-19, MT evaluation, MoverScore-1/2 and BERTScore-F1. Negative values indicate  $idf_{ori}$  works better.



(a) Pyramid



(b) Responsiveness

Figure 6: RD(dis,ori), TAC2008-2009, summary-level summarization evaluation, MoverScore-1/2 and BERTScore-F1. Negative values indicate  $idf_{ori}$  works better.

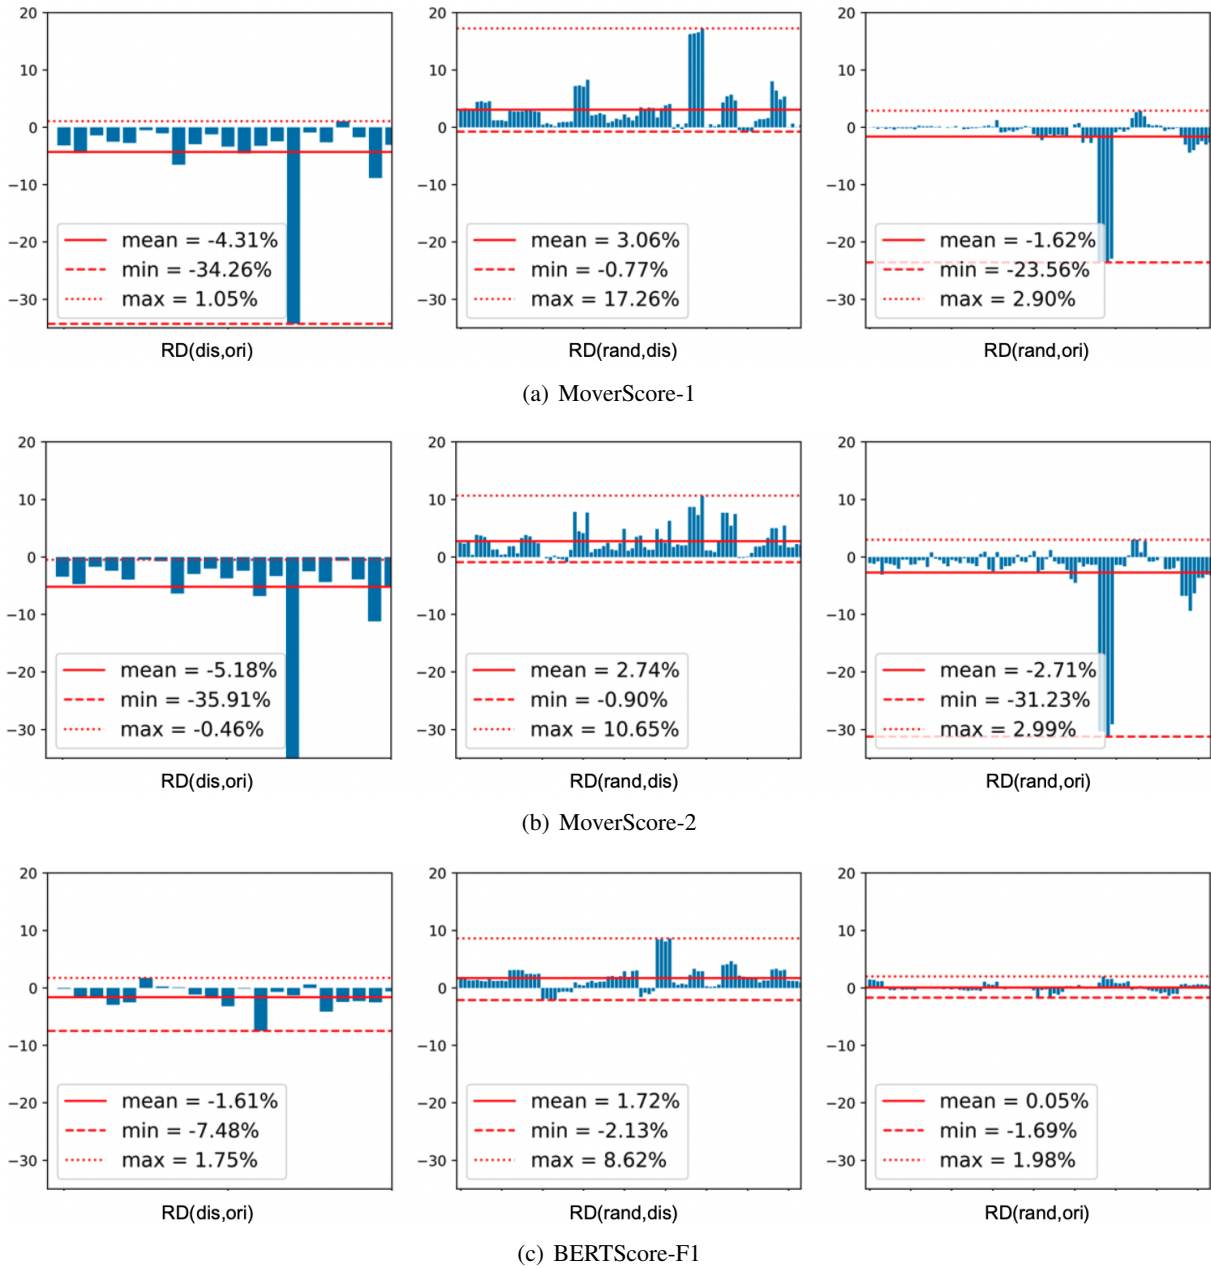


Figure 7: RD(dis,ori), RD(rand,dis), and RD(rand,ori); WMT17-19, segment-level MT evaluation, MoverScore-1/2 and BERTScore-F1. Negative values indicate the latter  $idf_x$  works better.

## A.10 Best settings of subword selection + PR

	WMT17						WMT18						WMT19			
	de	zh	ru	fi	tr	cs	de	zh	ru	fi	tr	cs	de	zh	ru	fi
first							✓									
all		✓		✓	✓	✓			✓							
ave-all	✓							✓		✓						✓
first+PR*											✓	✓	✓	✓	✓	✓
all+PR			✓							✓	✓	✓				
ave-all+PR																

Table 14: Best configuration of MoverScore-1 regarding subwords and punctuations for **other languages**. WMT17-19, **segment-level MT** evaluation. We mark the default configuration of MoverScore with \*.

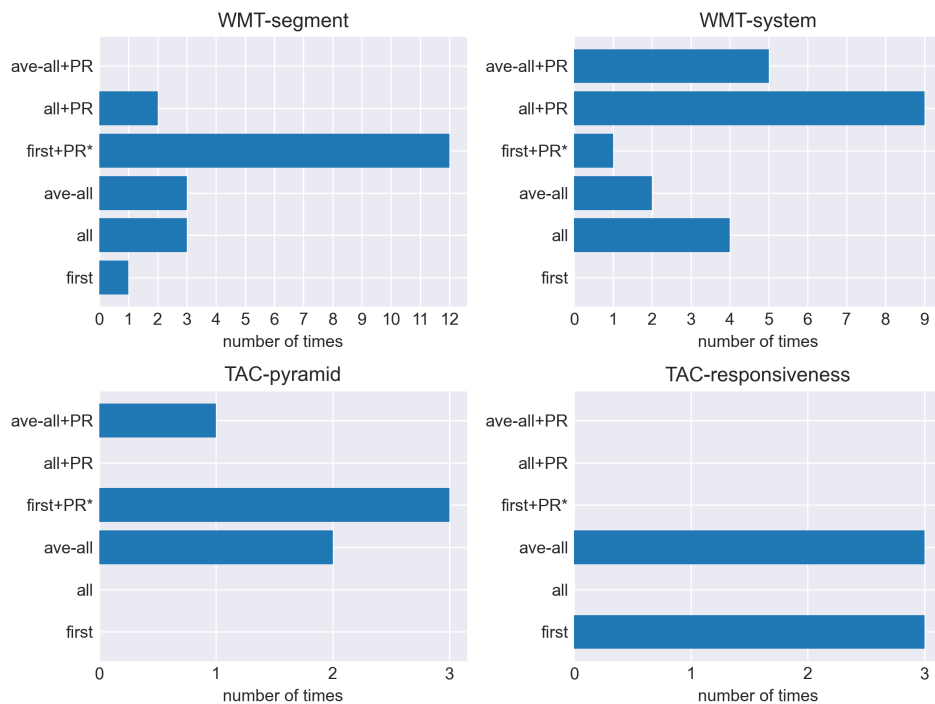


Table 15: Best configuration of MoverScore-1 regarding subwords and punctuations for **English**. WMT17-19 and TAC2008-2009. We mark the default configuration of MoverScore with \*.

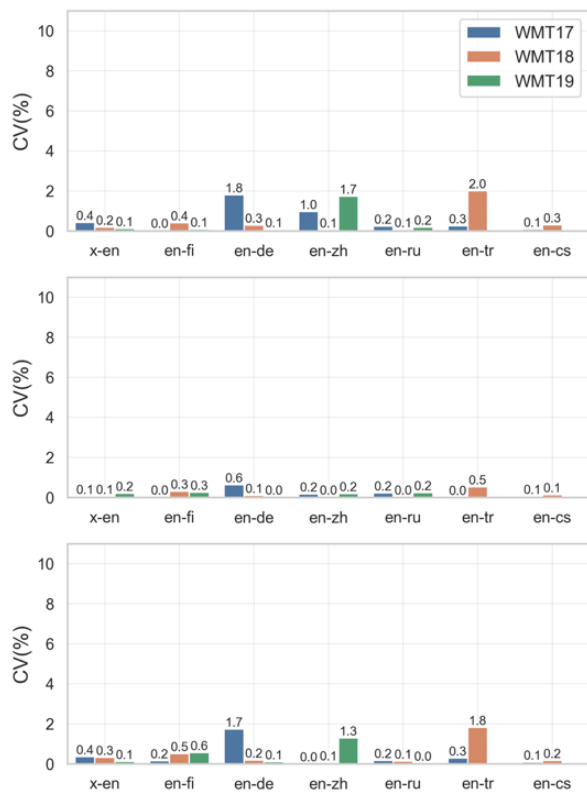


Figure 8: From top to bottom:  $CV_{STOP}$ ,  $CV_{IDF}$ ,  $CV_{SUB}$ . WMT17-19, system-level evaluation, MoverScore-1.

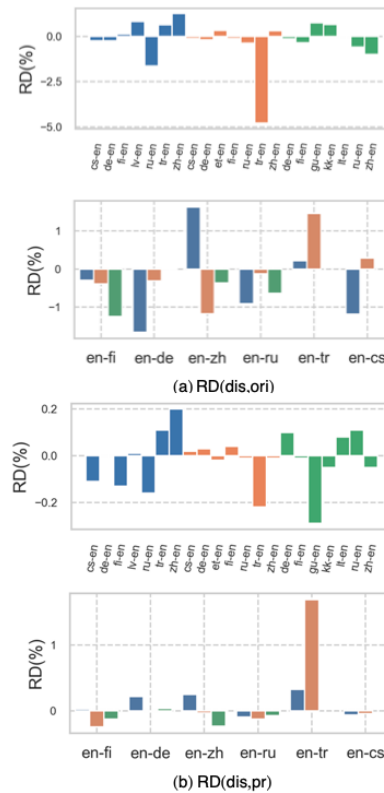


Figure 9:  $RD(dis,ori)$ ,  $RD(dis,pr)$ . WMT17-19, system-level evaluation, MoverScore-1.

