

# Sentence-level Aggregation of Lexical Metrics Correlates Stronger with Human Judgements than Corpus-level Aggregation

Paulo Cavalin<sup>1</sup>, Pedro H. Domingues<sup>2</sup>, Claudio Pinhanez<sup>1</sup>

<sup>1</sup>IBM Research Brazil

<sup>2</sup>PUC-Rio

pcavalin@br.ibm.com

## Abstract

In this paper we show that corpus-level aggregation hinders considerably the capability of lexical metrics to accurately evaluate machine translation (MT) systems. With empirical experiments, we demonstrate that averaging individual segment-level scores can make metrics such as BLEU and chrF correlate much stronger with human judgements and make them behave considerably more similarly to neural metrics such as COMET and BLEURT. We show that this difference exists because corpus- and segment-level aggregation differs considerably owing to the classical average of ratio versus ratio of averages mathematical problem. Moreover, as we also show, this difference considerably affects the statistical robustness of corpus-level aggregation. Considering that neural metrics currently only cover a small set of sufficiently-resourced languages, the results in this paper can help make the evaluation of MT systems for low-resource languages more trustworthy.

## 1 Introduction

Currently, we can group machine translation (MT) metrics into two main groups: lexical and neural metrics (Freitag et al. 2023). Although lexical metrics such as the BLEU score (Papineni et al. 2002a) have contributed considerably to the progress in MT in the past 20 years, neural metrics have emerged in recent years as viable alternatives not only to overcome the shortcomings of *n-gram* lexical matches, but also to leverage corpus-based training to improve the evaluation of MT (Rei et al. 2020a; Sellam, Das, and Parikh 2020b; Zhang\* et al. 2020).

Although it is quite clear that neural metrics tend to be more robust and will eventually replace lexical ones (Mathur, Baldwin, and Cohn 2020; Freitag et al. 2021, 2022), we believe that not only lexical metrics are still quite needed but also that there is room to improve the robustness of such metrics. We say that they are needed because most progress with neural metrics is observed on the one hundred or so most-resourced languages in the world, but almost 7,000 languages in the world still lack the minimum amount of data to train a MT model (Lorandi and Belz 2024). Furthermore, neural metrics can struggle in unseen domains

(Zouhar et al. 2024), which is an issue that can be considerably amplified with the lack of sufficient data. It is thus unrealistic to think that neural metrics will be applied for such scenarios in the near future. With this perspective in mind, we argue that there is a very under-explored and important component of lexical metrics which is the aggregation method.

Lexical metrics such as *BLEU* and *chrF* usually rely on *corpus-level aggregation (CLA)* (Papineni et al. 2002a; Popović 2015), but one can easily rely on *segment-level aggregation (SLA)*. The main difference between CLA and SLA is that, while in the former we compute *n-gram* matching statistics for all samples in a first step and then we compute a global score for the entire test set, in the latter we compute the statistics and the score for each sample individually and then use the mean of these scores for evaluating a test set. We notice that SLA is a very underexplored method for lexical metrics and, to the best of our knowledge, there is no previous work demonstrating why corpus-level should be preferred besides generic theoretical statistical assumptions.

At a glance, it surely looks that in the worst-case scenario, both CLA and SLA present comparable results, so there would be no good reason to question the aggregation method choice. But in this paper we show that, counter the common belief that CLA is better and should be the aggregation method of choice for lexical metrics, SLA is far more correlated to human judgements and to the more robust neural metrics. To support this claim, we first show that there is a conceptual difference between these two aggregation methods, mathematically related to the *average of ratios vs. ratios of averages* problem. Then, based on empirical experiments, we show not only that the choice between the aggregation method significantly impacts the resulting system-level scores but also that CLA is not statistically robust.

In greater details, in a first set of experiments we investigated whether the aforementioned mathematical differences between CLA and SLA statistically impact the scores provided by the metrics. For this, we considered 492 system outputs from the WMT23 metrics shared task (Freitag et al. 2023) and computed system-level scores considering not only these two aggregation methods with both BLEU and chrF as base metrics but also scores computing the mean of *bootstrap-resampled scores (BRS)*, as a reference point of a more statistically-robust approach. For this evaluation, we

computed the pairwise Pearson correlation of the scores of the two aggregation methods and the results show not only that the scores from CLA and SLA differ considerably but also that the scores from SLA correlate stronger with those from BRS.

We then conducted a deeper investigation of the statistical robustness of the two aggregation methods where we focused on evaluating the impact of the size of the test set on the correlation of scores. For that, we relied on downsampled test sets and computed the correlations of scores from downsampled versions of both CLA and SLA against each other, and against the statistically-robust BRS. The results corroborate our previous findings that SLA is not only more statistically robust than CLA but also show that this method is as statistically robust as BRS and can replace it as a much computationally-cheaper alternative.

However, the most surprising and relevant result, in our opinion, is that CLA is not only statistically weaker than SLA but it actually lacks any statistical robustness. In other words, when compared to BRS, the correlations of CLA scores computed on larger test sets are quite close to those computed on a test set with only a single sample. That means that the corpus-level evaluation could be simply replaced by single-sample evaluations.

Finally, in order to materialize what actually means the previous mathematical and statistical differences between the aggregation methods, in terms of the impact on the resulting quality of system-level scores, we computed correlations between the metrics and human judgements. For this, we considered human annotations from the WMT23 Metrics Shared Task (Freitag et al. 2023), and we also included additional neural metrics, i.e. *COMET*, *BLEURT*, and *BERTScore*, to provide a better view on the impact of the aggregation method. Our results provide strong evidence that SLA correlates much stronger with human judgements and are much more comparable to the outcomes of *BERTScore*. Considering that *BERTScore* is the only neural metric among these three which does not take into account the input sentence to compute the score, which is compatible to the way lexical metrics work, we believe that our results show that the use of segment-level aggregation reduces considerably the gap between lexical and neural metrics.

## 2 Related Work

The most well-known *lexical metric* for machine translation evaluation is the BLEU score, introduced more than two decades ago as a solution to make the development of MT system more scalable (Papineni et al. 2002b). The idea was to take advantage of a set of translations created by humans and to somehow measure the discrepancy between the outputs generated by a MT system and the reference translations. With that approach, one could develop different systems and select the one which produced the highest BLEU score in a completely automated fashion.

The way BLEU works is based on computing a Precision-like metric on overlaps of *n-grams* between the MT outputs and the references. That is done by counting up the number of *n-grams* generated in the MT outputs that also appear in the references. This computation is then heuristically refined

to address some issues such as wrongly-generated repetitions and very short texts and to combine different *n-gram* levels. BLEU also inspired other popular lexical metrics such as *chrF* (Popović 2015) and *chrF++* (Popović 2017), which play important roles to expand current NLP efforts into low-resource languages.

Despite the wide adoption of BLEU for about two decades, several works have focused on exploiting and overcoming its limitations (Graham and Baldwin 2014; Mathur, Baldwin, and Cohn 2020; Freitag et al. 2022, 2023). One issue already tackled by the community is the reliance in a large set of parameters and lack of standard and transparency in reporting results (Post 2018a). But another limitation, which is the reliance on lexical matches, resulted in the proposal of different alternative metrics, notably *neural metrics* such as *COMET* (Rei et al. 2020b), *BLEURT* (Sellam, Das, and Parikh 2020b), and *UniTE* (Wan et al. 2022).

As outlined in the WMT22 shared task results (Freitag et al. 2022), across diverse domains and tasks, neural-based metrics like *MetricX XXL* (Juraska et al. 2023), *COMET-22* (Rei et al. 2022), *UniTE* (Wan et al. 2022) and *BLEURT-20* (Sellam, Das, and Parikh 2020a) consistently outperformed BLEU and other non-neural counterparts in capturing evaluation nuances. In the subsequent WMT23 shared task (Freitag et al. 2023), the evaluation framework has been enhanced, expanding the metrics set and relying on a global score calculated through a weighted average across tasks. The results underscored the better alignment of neural-based metrics with human judgments than with non-neural ones.

Nevertheless, it is worth highlighting that neural metrics come with additional cost. Some metrics such as *UniTE* and *COMET* compute scores by relying also on the input provided to generate the translation, which obviously limits the application to cases where both the source and the target languages were used to train the underlying model. Even the neural metrics that consider only the MT outputs and the reference to compute the scores are quite limited, since they are usually trained with just dozens of languages. That limits neural metrics to, at best, hundreds of languages, the high- and mid-resourced ones.

Considering the currently-limited application of neural metrics and the vast number of under-studied languages in the world, there is still a vast application field for metrics based on overlaps of *n-grams* such as BLEU and *chrF*. At the same time, we observe here that there is a gap on a better understanding on the shortcomings of corpus-level aggregation, considering that most of the recent metrics rely on averages of segment-level scores.

We broadened this literature review by conducting an inspection of a total of 345 papers for a more comprehensive overview of the use of aggregation method. For that, we analysed MT-related papers published in recent editions<sup>1</sup> of *ACL* and *EMNLP*, two of the major venues for MT research, and catalogued the ways the evaluation of MT systems are reported. We observe that the BLEU metric dominates the field, being reported by 341 out of 345 papers, 98.8%. But details about the aggregation method are usually

---

<sup>1</sup>From 2020 to 2023

omitted, given that only 10 papers explicitly mention the aggregation method used for BLEU, where 6 of them explicitly mention that they are using corpus-level aggregation and 3 state they are using averages of bootstrapped resamplings, and only a single paper mentions the reliance on averages of segment-level BLEU scores. For the remaining 335 papers, we needed to infer the aggregation method by looking for references for specific tools, such as SacreBLEU (Post 2018b) and Moses. We were able to infer that 210 papers, 60.1%, contained at least some minimal references to such tools, being 183 to SacreBLEU and 36 to Moses. Considering that both tools implement the corpus-level aggregation as default, we considered that most of these papers relied on corpus-level BLEU, so there is likely 216 (63%) papers which implemented corpus-level aggregation for BLEU.

The main findings of the previously-mentioned study is that it looks like very few recently-published research relies on segment-level aggregation, that is, one paper using segment-level aggregation for the BLEU metric (Chen et al. 2023); and some papers using segment-level aggregation for chrF, such as in the reports for the WMT Metrics shared task (Freitag et al. 2021, 2022, 2023). And even when the aggregation method is mentioned, we see a clear lack of comprehension of the impact of the aggregation method. This paper aims to bridge this gap.

### 3 Corpus- vs. Segment-level Aggregation

In this section we focus on describing in details both corpus-level aggregation (CLA) and segment-level aggregation (SLA) methods and on explaining why the choice of one over another is mathematically different and how it might affect the resulting scores. For the sake of simplicity we will focus on the BLEU score but in our empirical evaluations presented afterwards we demonstrate that our hypotheses are not limited to BLEU and are at least also applicable to chrF.

#### 3.1 A Case Study with BLEU

In this section we discuss why we expect differences in the results according to the the aggregation method. For that we will rely on a simplified abstraction of BLEU, considering that this metric consists of computing a Precision-like score for the generated translation or for a set of generated translations (Papineni et al. 2002b).

Before explaining the aggregation method, it is worth explaining how BLEU is computed for a single sentence, i.e. the so-called *sentence-level* BLEU score. For this case, we basically compute the total of *n-gram matches* between the *candidate* sentence, i.e. the sentence generated by the ML model, and the *references*, representing the ground-truth translations generated by a person that is fluent enough in the target language. The matching is computed by evaluating the number of n-grams present in the candidate sentence that also appear in the references. After we computed that number of matches, we divide this number by the total of n-grams contained in the candidate output, to compute a Precision-ish score that represents the sentence-level evaluation score. Of course, since this is a simplified abstraction, we are not taking into account modified n-gram precision,

clipping, combination of different n-gram levels, and brevity penalty, but this does not affect our rationale.

When we expand sentence-level BLEU to evaluate the MT system on a corpus of text or on an entire test set, for instance, the main approach is to rely on the so-called *corpus-level* BLEU (check Section 2.2.1 in (Papineni et al. 2002b) for further details). The corpus-level aggregation (CLA) of BLEU consists of a global computation of n-gram matches and a single scoring for all samples in the test set, done at once. That is, all n-gram matches are counted and summed up, and this number is divided by the sum of the lengths of all candidate sentences.

Although corpus-level BLEU is the usual choice and the default options in tools such as SacreBLEU (Post 2018a), it is quite easy to adopt averages of segment-level scores, or simply *segment-level aggregation* (SLA), to compute system-level scores (Bugliarello et al. 2020; Niu et al. 2020). The implementation is straightforward, where one simply need to compute sentence-level BLEU scores for each individual test sample, than using the mean of such scores as the final system-level scores. This aggregation method presents some advantages, such as allowing to compute statistical metrics such as standard deviations, which is not possible with corpus-level aggregation. Notice that bootstrap resampling is a somewhat popular method to compute statistical significance tests on corpus-level scores (Koehn 2004; Jon and Bojar 2023; Fucci et al. 2023) and can also be used to compute statistical metrics, but it is a more expensive approach in terms of computation requirements.

At a first sight, it is reasonable to believe that CLA and SLA provide the same results, so the main advantage of the latter would only be the less expensive way to compute statistical metrics. But there is a conceptual mathematical difference between these two aggregation methods, which we explain next.

#### 3.2 CLA and SLA as Differently Weighted Ratio Averages

We dive now into a mathematical explanation of the difference between corpus-level and segment-level aggregations, to demonstrate why these two methods may present differences in the results. We argue that the main difference between the two aggregation methods can be seen as a classical case of *ratio of averages vs. average of ratios*.

To understand this, let us adopt a simplified definition of BLEU as a ratio of the number of matches  $m$  by n-grams  $w$  in a corpus, as in the previous section. And let us also refer to the corpus-level and segment-level BLEU as simply BLEU and m-BLEU, respectively. Considering all the  $n$  sentences  $i$  of the corpus, it is evident that BLEU can be computed by the ratio between the sum of all partial matches  $m_i$  in each sentence  $i$  by the sum of n-grams in all sentences  $i$ ,  $w_i$ :

$$\text{BLEU} = \frac{m}{w} = \frac{\sum_{i=1}^n m_i}{\sum_{i=1}^n w_i} \quad (1)$$

Accordingly, m-BLEU is the average of ratios between the number of matches  $m_i$  and the number of words  $w_i$  of

each sentence  $i$ :

$$\text{m-BLEU} = \frac{1}{n} \sum_{i=1}^n \frac{m_i}{w_i} = \sum_{i=1}^n \left( \frac{1}{n} \right) \frac{m_i}{w_i} \quad (2)$$

It is easy to derive that BLEU is the weighted average of the sentence ratios by the proportional length of each sentence  $i$ :

$$\begin{aligned} \text{BLEU} &= \frac{\sum_{i=1}^n m_i}{\sum_{i=1}^n w_i} = \sum_{i=1}^n \frac{m_i}{w} = \sum_{i=1}^n \frac{w_i}{w} \frac{m_i}{w_i} \\ \text{BLEU} &= \sum_{i=1}^n \left( \frac{w_i}{w} \right) \frac{m_i}{w_i} \end{aligned} \quad (3)$$

As we see, m-BLEU weights the ratios equally with  $1/n$  weights while BLEU weights the ratios with the value  $w_i/w$  which is proportional to the length of each sentence  $i$ . Therefore, BLEU results are likely to be biased by the proportion of matches and candidate sentences lengths, while m-BLEU considers the performance independently of that.

## 4 Empirical Evaluation

In this section we present experiments aiming at investigating whether the choice of aggregation method actually impacts the score provided by a lexical metric. For that, we consider three different implementations of BLEU and chrF and, considering the outputs from 492 different systems, we present a detailed analysis on the distribution of scores provided by these metrics using different aggregation methods.

### 4.1 The Dataset

For this investigation we rely on the WMT 2023 Metrics Shared Task dataset (Freitag et al. 2023), or simply *WMT23 dataset*, comprising the results of 492 different MT systems, involving different languages and domains. This dataset contains 468,850 system outputs with the corresponding inputs and references, in a quite diverse setting, containing 147 different language pairs of 48 source languages and 44 target languages.

We converted the 468,850 raw entries to 492 system evaluations by grouping the data by dataset type (`challengesets2023` or `generaltest2023`), dataset (`challenge_ACES`, `challenge_DFKI`, `challenge_NRC-MSLC23`, or `generaltest2023`), language pair (147 options), and system (two systems for the `challengesets2023` dataset type and 14 systems for the `generaltest2023`).

### 4.2 Three Implementations of BLEU

We considered three different implementations for BLEU. Two of them are based on the corpus- (CLA) and segment-level aggregation (SLA), and the third relies on bootstrap-resampled scores (BRS) to provide more robust statistical estimates, so that we can use this approach as a reference point for statistically-reliable scores.

As a consequence, the first implementation is referred to as simply **BLEU**, consisting of the traditional BLEU score with CLA, here computed with the SacreBLEU tool with

default parameters (Post 2018b)<sup>2</sup>. For this metric, given a test set with samples and the corresponding reference MT outputs, we take all MT outputs and references at once, in a single list, and compute the score with the `corpus_bleu` function in Python code.

The second metric, referred to as **m-BLEU**, implements SLA. For this we simply compute the score of each segment (i.e. a test sample) with SacreBLEU’s `sentence_bleu` function, again with default parameters, and calculate the overall average to provide the system-level score.

The third metric, named **x-BLEU**, consists of implementing the BRS method, as in (Niu et al. 2020; Liu et al. 2021). This is an alternative to computing system-level scores with higher statistical robustness, where the resamplings represent varied rearrangements of the test set to compute corpus-level BLEU. We rely on 1,000 resamplings, with replacement, of 1,000 samples<sup>3</sup> for each system, and we applied the same SacreBLEU’s `corpus_bleu` function, with default parameters, on top of each resampled set, generating 1,000 scores for each system. We then provide the average of those 1,000 scores as the system-level score.

### 4.3 Three Implementations of chrF

In order to investigate whether our observation also generalize to other lexical metrics, we explored also three different implementations of chrF. Notice that chrF is quite similar to BLEU, where overlaps of character n-grams are computed instead of overlap of words, and the final score is based on an F1 score instead of the Precision-ish score used by BLEU. More importantly, the same variety of aggregation methods are possible to be used with chrF and those varieties can be put in practice with the `corpus_chrf` and `sentence_chrf` functions from SacreBLEU.

Consequently, the three different chrF implementations that we consider in this work are: **chrF**, the CLA version computed with `corpus_chrf` SacreBLEU’s function; **m-chrF**, the SLA implementation relying on the average of segment-level scores computed with `sentence_chrf`; and **x-chrF**, the BRS metric computed with averages of `corpus_chrf` on 1,000 resamplings with 1,000 samples. Notice that we always rely on the default SacreBLEU’s parameters for such functions.

### 4.4 CLA and SLA Correlate Weakly to Each Other, and SLA Correlates Strongly to BRS

Our first evaluation focused on investigating the distribution of scores provided by the three different aggregation methods, how these distributions correlate to each other, and how they correlate to more statistically-robust scores. We focus first on analysing BLEU and then show results with chrF to understand the impact of the base metric.

Figure 1, displaying a grid of 9 plots, presents the main results of this analysis. Notice that in the diagonal we show the histogram of the distribution of scores for each metric,

<sup>2</sup>nrefs:1;case:mixed;eff:no;tok:13a;smooth:exp; version:2.0.0

<sup>3</sup>Notice that 1,000 samples is roughly very close the mean of samples in the WMT23 dataset, which is 952.94.




Figure 1: Correlation plots of the different BLEU metrics to each other and, in the diagonal, the distribution of the scores of the 3 metrics.

considering the scores for the 492 systems, and in the upper and lower non-diagonal cells we present the pairwise scatter plots and Pearson correlations, in the -1 to 1 range, considering the metrics scores. The main intuition of computing Pearson correlation is to assess the linear correlation between the metrics. That is, a highly-positive correlation, i.e. a value close to 1.0, indicates that high scores in one metric correspond to high scores in another metric, and vice-versa. This correlation helps determine whether the metrics are aligned in capturing what the good and bad translation results are and what is in-between.

From both the distributions and correlation values, we can clearly see that x-BLEU and m-BLEU tend to produce much closer results to one-another than to BLEU. In the diagonal, we can see that the distribution of BLEU is right-skewed, while the other distributions are more centralized. The means of the distributions are 0.29, 0.43, and 0.39, for BLEU, x-BLEU, and m-BLEU; and the correlations of BLEU to x-BLEU and m-BLEU are of 0.53 and 0.48, respectively, and x-BLEU and m-BLEU present 0.95 of Pearson correlation.

The results of a repetition of the experiments with chrF as the base metric is presented in Figure 2 and we observe quite similar outcomes. That is, chrF correlates weakly with the other metrics and m-chrF correlates strongly to x-chrF. However, we observe that, unlike the BLEU metrics, the distributions of chrF, x-chrF, and m-chrF are all centralized, presenting means of 0.54, 0.57, and 0.58, respectively. In terms of Pearson correlation, x-chrF and m-chrF present almost perfect correlations to each other with 0.99, while chrF correlates weakly to the other methods, with correlations of 0.56 and 0.55.

What it is very noticeable from these experiments, is that SLA presents a quite higher correlation to BRS than CLA does, which is strong evidence that SLA is not only as capable as BRS to compute statistically-robust scores but also a more-cost-effective option. Moreover, it also indicates that CLA is a quite less statistically-robust method.




Figure 2: Correlation plots of the different chrF metrics to each other and, in the diagonal, the distribution of the scores of the 3 metrics.




Figure 3: Correlation plots of the different implementations based on the BLEU and chrF metrics.

#### 4.5 Aggregation Methods Present Strong Cross-Metric Correlation

We now focus on investigating the correlations between the different aggregation methods of the different BLEU and chrF implementations. The experiments are analogous to those presented in the previous sections, so in Figure 3 we present the scatter plot and the Pearson correlation of each pair of similar implementations, i.e. BLEU vs. chrF, x-BLEU vs. x-chrF, and m-BLEU vs. m-chrF.

Interestingly, these results indicate that similar aggregation methods correlate strong with each other. Although the correlation of BLEU to chrF is the weakest, with 0.77, it is above the 0.50 to 0.55 correlation values that such metrics presented to the non corpus-level ones. And the non corpus-level metrics correlate quite strong to each other, since x-BLEU vs. x-chrF present a correlation of 0.88, and m-BLEU vs. m-chrF of 0.93. These results indicate that, as we mentioned in Section 3.2, corpus-level aggregation might be biased by the ratio between n-gram matches and sentence lengths, and that can explain why BLEU and chrF correlate strong to one another, and weakly to the other methods.

In order to gather additional evidence, in Figure 4 we present a cross-metric comparison between m-BLEU and m-chrF, against BLEU, chrF, x-BLEU, and x-chrF, respectively. Again, a weak correlation of BLEU and chrF to the non corpus-level metrics is seen, given that BLEU correlates




Figure 4: Correlation plots of the m-BLEU and m-chrF to the BLEU, chrF, x-BLEU, and x-chrF metrics.

weakly to x-chrF and m-chrF, and vice-versa. Moreover, m-BLEU and x-BLEU also correlate strongly to x-chrF and m-chrF, respectively, indicating that the corpus-level aggregation introduces a bias which can hinder the statistical robustness of lexical metrics.

#### 4.6 The Statistical Robustness of CLA and SLA

Given that one main outcome from the previous section is that corpus-level aggregation (CLA) metrics correlate weakly to their more statistically robust counterparts such as bootstrap-resampled scores (BRS) and that the segment-level aggregated (SLA) metrics correlate strongly to BRS, in this section we focus at investigating the statistical robustness of the aggregation methods.

Our methodology consists of using BRS as an upper bound for statistical robustness for system-level scores, since they rely on the well-known bootstrap resampling method, and on not only evaluating the correlation between the distributions of BRS-based scores against downsampled versions of both CLA- and SLA-based scores, but also the correlation of these downsampled metrics to each other. With this approach we believe we can understand the sensitiveness of the aggregation methods to the number of samples and how their statistical robustness is affected by the number of samples.

In greater detail, we created downsampled test sets for each of 492 datasets, considering three different sizes, i.e.  $N = \{1, 10, 100\}$ , and computed scores with both CLA- and SLA-based metrics on each of these downsampled test sets. Then, we computed the Pearson correlation between the scores of these downsampled versions of each aggregation method and between these scores and the more statistically-robust ones computed with BRS. This process were repeated 1,000 times for a better statistical estimate. The distribution of correlations, for both BLEU- and chrF-related evaluations, are presented in Figure 5.

The results show that CLA-based metrics are statistically weak. That is, as we can see in both top and bottom left-most plots, the downsampled CLA- and SLA-based metrics correlate strongly only in the case of datasets with only one sample, i.e.  $N = 1$ , with correlation values approaching 1.0. Nevertheless, as the number of samples increases, their correlation reduces quite drastically, to about 0.7 with  $N = 10$

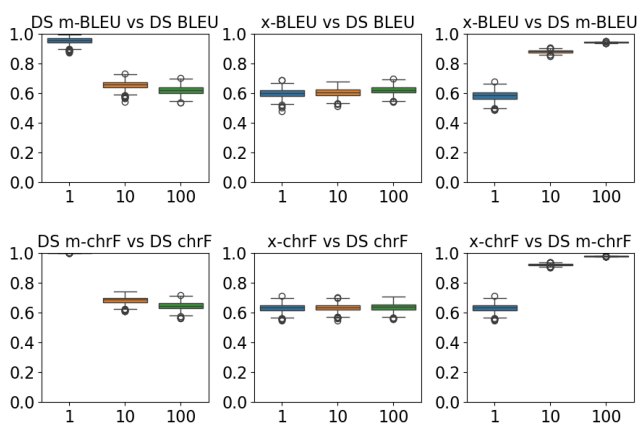


Figure 5: Boxplots of the correlations of scores comparing downsampled (DS) version of corpus- and segment-level aggregation. The x axis represents sample size, and the y axis Pearson correlation.

and about 0.6 with  $N = 100$ . That means that, statistically speaking, we can claim that those metrics differ considerably to each other.

We can go further and dare to claim that an evaluation of an entire test set with CLA is similar to conducting the evaluation of the same system with just a single sample. In the central plots of Figure 5, we can see that the correlation of their downsampled scores to the more statistically-robust BRS-based methods does not increase, or increases very insignificantly, with the number of samples. On the other hand, as we can observe in the right-most plots of the same figure, the correlation of donwsampled SLA-based metrics correlate weakly with the bootstrapped scores with  $N = 1$ , but the correlation increases significantly with  $N = 10$  and  $N = 100$ , showing that this aggregation method does not suffer from the same lack-of-statistical-robustness problem.

## 5 Impact of the Aggregation Method

In order to materialize the actual impact of the aggregation method in computing system-level scores, in this section we present an evaluation comparing the scores of the metrics described in Section 4 compared to ground-truth data from human judgements. For that we rely on three language pairs with Multidimensional Quality Measurements (MQM), used to benchmark metrics for the WMT23 Metrics Shared Task (Freitag et al. 2023). The human scores are computed with the weighted average of multiple dimensions. To present a more extensive evaluation, we also consider the eight language pairs annotated with Direct Assessment (DA) scores.

By evaluating MQM and DA data individually, we compute the mean Pearson correlation of the system scores from each language pair to the implementations of BLEU and chrF described in Section 4. In this case, we consider Pearson correlation values in the usual -1 to 1 range, but we also list the correlations converted to the 0 to 1 interval for a more straightforward comparison with the results presented in (Freitag et al. 2023). For comparison purposes, we also in-

	MQM		DA	
	metric	corr. ([0,1])	metric	corr ([0,1])
1	COMET	0.923 (0.962)	COMET	0.926 (0.963)
2	BLEURT	0.872 (0.936)	BLEURT	0.917 (0.959)
3	BERTScore	0.855 (0.928)	BERTScore	0.821 (0.911)
4	m-chrF	0.806 (0.903)	m-chrF	0.802 (0.901)
5	x-chrF	0.804 (0.902)	x-chrF	0.793 (0.897)
6	m-BLEU	0.776 (0.888)	m-BLEU	0.729 (0.865)
7	x-BLEU	0.762 (0.881)	x-BLEU	0.456 (0.728)
8	BLEU	0.425 (0.713)	chrF	0.285 (0.643)
9	chrF	0.392 (0.696)	BLEU	-0.006 (0.497)

Table 1: Rankings of Pearson correlations from metric scores to human judgments, in the [-1,1] range. Pearson correlations in the [0,1] scale are provided inside parentheses, the scale used in (Freitag et al. 2023) which presents relevant previous results.

clude three neural metrics to present a better reference point related the impact of the aggregation method compared with these more robust metrics: COMET (Rei et al. 2020a), with *Unbabel/wmt22-cometkiwi-da* base model, BLEURT (Sellam, Das, and Parikh 2020a), using *BLEURT-20* base model; and BERTScore (Zhang\* et al. 2020), relying on *bert-base-multilingual-cased* base model.

The results are presented in Table 1, where we list the rankings of the metrics according to their correlation to the human-annotated data. The results, interestingly, show that the SLA-based methods, i.e. m-BLEU and m-chrF, not only considerably outperform BLEU and chrF, the CLA-based ones, but also they provide correlations which are much closer to those of the neural metrics. We can also see that, with chrF and MQM annotations, we can improve from a moderate-to-low correlation of 0.392 to a strong correlation of 0.806, which is just 0.049 correlation points weaker than the correlation of BERTScore and just 0.066 of BLEURT. With the DA annotations, we can observe that m-chrF performs even closer to the neural metrics. Notice that the SLA-based metrics also outperform their BRS counterparts, i.e. x-BLEU and x-chrF, demonstrating again the statistical robustness of the SLA method.

## 6 Final Discussion

In this work we presented a comparison of the traditional corpus-level aggregation against the less popular method based on averaging individual segment-level scores, showing that the latter can generate system evaluation scores which correlate stronger to human judgements and to neural metrics. We demonstrate that this difference happens because of a fundamental mathematical difference: CLA metrics are biased towards the performance on long sentences, considerably reducing the capability of lexical metrics to correlate with human judgements when corpus-level aggregation is considered. The main issue we observed is that corpus-level aggregation voids the statistical robustness of a test set-based evaluation, providing scores that are comparable to evaluations with a single sample.

One main outcome of this paper should be regarded as a clear recommendation to MT researchers: **stop using**

**corpus-level aggregation.** As we have shown, segment-level aggregation is not only better than corpus-level in terms of correlation to human judgments but also comparable, if not better, than robust statistical evaluations based on bootstrap resampling. Similarly, MT researchers should **use segment-level scores for statistical evaluation** instead of the expensive computations for bootstrap resampling, although one could also bootstrap resample segment-level scores to get even more robust estimates.

Finally, we would like to draw the attention for the vast application field that lexical metrics still have since neural metrics cover only about a hundred or so languages. Moreover, it is important to understand that some of the bad reputation which lexical metrics received in the past years might be under-deserved because of the wide, but wrong, adoption of corpus-level aggregation.

## 7 Limitations

One clear limitation of this work is relying on a single data source, which is the WMT23 Shared Task dataset. Nevertheless, since it is a very recent dataset, considerably large, and coming from a very well-known workshop on machine translation, we believe that this dataset is strong enough to experimentally provide evidence, as we have done in our empirical evaluations.

Another limitation is to rely on a single tool to compute BLEU and chrF, which is SacreBLEU, even though there are other implementations available. Again, since this tool can be viewed as a de facto standard BLEU implementation (Post 2018b), we believe that the tool is strong enough to experimentally prove our assumptions in the empirical evaluation.

Additionally, we have not thoroughly evaluated the metrics, in the sense of changing parameters of the metrics such as the maximum n-gram lengths and thus forth. We stayed with the default SacreBLEU’s parameters only. But again, given the level of impact the changing the aggregation method presents, as we show, we do not believe that simply changing the base metrics’ parameters would affect significantly the outcomes of this paper.

## Ethical Statement

We are not aware of any ethical issue that this paper might raise. All of the data used for this research are publicly-available and the outcomes of this paper are likely to contribute to improving the quality of the research of the whole machine translation community.

## References

- Bugliarello, E.; Mielke, S. J.; Anastasopoulos, A.; Cotterell, R.; and Okazaki, N. 2020. It’s Easier to Translate out of English than into it: Measuring Neural Translation Difficulty by Cross-Mutual Information. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 1640–1649. Online: Association for Computational Linguistics.
- Chen, M.; Duquenne, P.-A.; Andrews, P.; Kao, J.; Mourachko, A.; Schwenk, H.; and Costa-jussà, M. R.

2023. BLASER: A Text-Free Speech-to-Speech Translation Evaluation Metric. In Rogers, A.; Boyd-Graber, J.; and Okazaki, N., eds., *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 9064–9079. Toronto, Canada: Association for Computational Linguistics.
- Freitag, M.; Mathur, N.; Lo, C.-k.; Avramidis, E.; Rei, R.; Thompson, B.; Kocmi, T.; Blain, F.; Deutsch, D.; Stewart, C.; Zerva, C.; Castilho, S.; Lavie, A.; and Foster, G. 2023. Results of WMT23 Metrics Shared Task: Metrics Might Be Guilty but References Are Not Innocent. In Koehn, P.; Haddow, B.; Kocmi, T.; and Monz, C., eds., *Proceedings of the Eighth Conference on Machine Translation*, 578–628. Singapore: Association for Computational Linguistics.
- Freitag, M.; Rei, R.; Mathur, N.; Lo, C.-k.; Stewart, C.; Avramidis, E.; Kocmi, T.; Foster, G.; Lavie, A.; and Martins, A. F. T. 2022. Results of WMT22 Metrics Shared Task: Stop Using BLEU – Neural Metrics Are Better and More Robust. In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, 46–68. Abu Dhabi, United Arab Emirates (Hybrid): Association for Computational Linguistics.
- Freitag, M.; Rei, R.; Mathur, N.; Lo, C.-k.; Stewart, C.; Foster, G.; Lavie, A.; and Bojar, O. 2021. Results of the WMT21 Metrics Shared Task: Evaluating Metrics with Expert-based Human Evaluations on TED and News Domain. In *Proceedings of the Sixth Conference on Machine Translation*, 733–774. Online: Association for Computational Linguistics.
- Fucci, D.; Gaido, M.; Papi, S.; Cettolo, M.; Negri, M.; and Bentivogli, L. 2023. Integrating Language Models into Direct Speech Translation: An Inference-Time Solution to Control Gender Inflection. In Bouamor, H.; Pino, J.; and Bali, K., eds., *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 11505–11517. Singapore: Association for Computational Linguistics.
- Graham, Y.; and Baldwin, T. 2014. Testing for Significance of Increased Correlation with Human Judgment. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 172–176. Doha, Qatar: Association for Computational Linguistics.
- Jon, J.; and Bojar, O. 2023. Breeding Machine Translations: Evolutionary approach to survive and thrive in the world of automated evaluation. In Rogers, A.; Boyd-Graber, J.; and Okazaki, N., eds., *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2191–2212. Toronto, Canada: Association for Computational Linguistics.
- Juraska, J.; Finkelstein, M.; Deutsch, D.; Siddhant, A.; Mirzazadeh, M.; and Freitag, M. 2023. MetricX-23: The Google Submission to the WMT 2023 Metrics Shared Task. In Koehn, P.; Haddow, B.; Kocmi, T.; and Monz, C., eds., *Proceedings of the Eighth Conference on Machine Translation*, 756–767. Singapore: Association for Computational Linguistics.
- Koehn, P. 2004. Statistical Significance Tests for Machine Translation Evaluation. In *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing*, 388–395. Barcelona, Spain: Association for Computational Linguistics.
- Liu, Y.; Meng, F.; Chen, Y.; Xu, J.; and Zhou, J. 2021. Scheduled Sampling Based on Decoding Steps for Neural Machine Translation. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 3285–3296. Online and Punta Cana, Dominican Republic: Association for Computational Linguistics.
- Lorandi, M.; and Belz, A. 2024. High-quality Data-to-Text Generation for Severely Under-Resourced Languages with Out-of-the-box Large Language Models. In Graham, Y.; and Purver, M., eds., *Findings of the Association for Computational Linguistics: EACL 2024*, 1451–1461. St. Julian’s, Malta: Association for Computational Linguistics.
- Mathur, N.; Baldwin, T.; and Cohn, T. 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*, 4984–4997. Online: Association for Computational Linguistics.
- Niu, X.; Mathur, P.; Dinu, G.; and Al-Onaizan, Y. 2020. Evaluating Robustness to Input Perturbations for Neural Machine Translation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 8538–8544. Online: Association for Computational Linguistics.
- Papineni, K.; Roukos, S.; Ward, T.; and Zhu, W.-J. 2002a. Bleu: a Method for Automatic Evaluation of Machine Translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, 311–318. Philadelphia, Pennsylvania, USA: Association for Computational Linguistics.
- Papineni, K.; Roukos, S.; Ward, T.; and Zhu, W.-J. 2002b. Bleu: a Method for Automatic Evaluation of Machine Translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, 311–318. Philadelphia, Pennsylvania, USA: Association for Computational Linguistics.
- Popović, M. 2015. chrF: character n-gram F-score for automatic MT evaluation. In *Proceedings of the Tenth Workshop on Statistical Machine Translation*, 392–395. Lisbon, Portugal: Association for Computational Linguistics.
- Popović, M. 2017. chrF++: words helping character n-grams. In *Proceedings of the Second Conference on Machine Translation*, 612–618. Copenhagen, Denmark: Association for Computational Linguistics.
- Post, M. 2018a. A Call for Clarity in Reporting BLEU Scores. In *Proceedings of the Third Conference on Machine Translation: Research Papers*, 186–191. Brussels, Belgium: Association for Computational Linguistics.
- Post, M. 2018b. A Call for Clarity in Reporting BLEU Scores. In *Proceedings of the Third Conference on Machine Translation: Research Papers*, 186–191. Brussels, Belgium: Association for Computational Linguistics.

Rei, R.; Farinha, A. C.; de Souza, J. G.; Ramos, P. G.; Martins, A. F.; Coheur, L.; and Lavie, A. 2022. Searching for COMETINHO: The Little Metric That Could. In *Proceedings of the 23rd Annual Conference of the European Association for Machine Translation*, 61–70. Ghent, Belgium: European Association for Machine Translation.

Rei, R.; Stewart, C.; Farinha, A. C.; and Lavie, A. 2020a. COMET: A Neural Framework for MT Evaluation. In Webber, B.; Cohn, T.; He, Y.; and Liu, Y., eds., *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2685–2702. Online: Association for Computational Linguistics.

Rei, R.; Stewart, C.; Farinha, A. C.; and Lavie, A. 2020b. COMET: A Neural Framework for MT Evaluation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2685–2702. Online: Association for Computational Linguistics.

Sellam, T.; Das, D.; and Parikh, A. 2020a. BLEURT: Learning Robust Metrics for Text Generation. In Jurafsky, D.; Chai, J.; Schluter, N.; and Tetreault, J., eds., *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 7881–7892. Online: Association for Computational Linguistics.

Sellam, T.; Das, D.; and Parikh, A. P. 2020b. BLEURT: Learning Robust Metrics for Text Generation. In *Proceedings of ACL*.

Wan, Y.; Liu, D.; Yang, B.; Zhang, H.; Chen, B.; Wong, D.; and Chao, L. 2022. UniTE: Unified Translation Evaluation. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 8117–8127. Dublin, Ireland: Association for Computational Linguistics.

Zhang\*, T.; Kishore\*, V.; Wu\*, F.; Weinberger, K. Q.; and Artzi, Y. 2020. BERTScore: Evaluating Text Generation with BERT. In *International Conference on Learning Representations*.

Zouhar, V.; Ding, S.; Currey, A.; Badeka, T.; Wang, J.; and Thompson, B. 2024. Fine-Tuned Machine Translation Metrics Struggle in Unseen Domains. In Ku, L.-W.; Martins, A.; and Srikumar, V., eds., *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, 488–500. Bangkok, Thailand: Association for Computational Linguistics.