

# Pitfalls and Outlooks in Using COMET

Vilém Zouhar<sup>★ 1</sup> Pinzhen Chen<sup>★ 2</sup> Tsz Kin Lam<sup>2</sup> Nikita Moghe<sup>2</sup> Barry Haddow<sup>2</sup>

<sup>1</sup>ETH Zurich <sup>2</sup>University of Edinburgh

vzouhar@ethz.ch {pinzhen.chen, tlam, nikita.moghe, bhaddow}@ed.ac.uk

## Abstract

The COMET metric has blazed a trail in the machine translation community given its strong correlation with human judgements of translation quality. Its success stems from being a pre-trained multilingual model finetuned for quality assessment. However, it being a neural metric also gives rise to a set of pitfalls that may not be widely known. We investigate these unexpected behaviours from three aspects: 1) technical: obsolete software versions and compute precision; 2) data: empty content, language mismatch, and translationese at test time as well as distribution and domain biases in training; 3) usage and reporting: multi-reference support and model referencing in the literature. All of these problems imply that COMET scores may be not incomparable between papers or technical setups and we put forward our perspective on fixing each issue. Furthermore, we release the `sacreCOMET` package that can generate a signature for the software and model configuration as well as an appropriate citation. The goal of this work is to help the community make more sound use of the COMET metric.

## 1 Introduction

Automated metrics provide a cheap and scalable way of evaluating and benchmarking NLP models. In machine translation (MT), the evaluation protocol has moved from string matching metrics (BLEU, TER, chrF, inter alia; Papineni et al., 2002; Snover et al., 2006; Popović, 2015) to trained neural metrics (Shimanaka et al., 2018; Takahashi et al., 2020; Rei et al., 2020a; Sellam et al., 2020) with COMET being widely adopted. The trained metrics have been shown to correlate much better with human judgement (Freitag et al., 2021, 2022b, 2023), making them more reliable in estimating translation quality and ranking translation systems.

Nonetheless, the solution to translation evaluation is yet to be perfected. One problem is the haphazard use of the metric. Previously, Post (2018) showed that different usages and implementations of BLEU, e.g. tokenization and smoothing, lead to inconsistencies in scores. We suspect that the use of COMET might be sensitive to misconfigurations too, resulting in unexpected behaviours. Furthermore, trained MT metrics are optimized on a limited amount of data (usually valid machine translations), leading to overfitting and reduced robustness against corner cases. Contributions of this work are listed as follows:

- we reveal nine problems spanning technical issues, data biases, and model reporting;
- we show that inconsistent use of COMET leads to non-comparable scores across papers or setups;
- we release the `sacreCOMET` package for better reporting and reproducibility;
- we provide directions for future work on building learned metrics.

## 2 Background and Setup

**Metric background.** Publicly available human judgements of translation quality come from shared task annotation campaigns, where translations are evaluated with some annotation protocol. From 2017, in WMT, the protocol was a variant of direct assessment (DA; Graham et al., 2013) which has annotators providing a number from 0 (lowest) to 100 (highest) as the segment quality. This has been subsequently replaced by MQM and ESA protocols (Lommel et al., 2014; Kocmi et al., 2024d), though DA remains the most abundant data source for neural metric training.

Automated metrics aim to yield scores that correlate with human judgements of translations. Most metric scores are computed at the segment level and then aggregated at the system level to e.g. obtain system comparison. The evaluation of metrics is done with respect to the human judgements.

<sup>★</sup>Equal contributions.

<sup>0</sup>Code: [github.com/PinzhenChen/sacreCOMET](https://github.com/PinzhenChen/sacreCOMET)

**COMET models.** Metrics such as chrF or BLEU are heuristic algorithms that match n-grams between the translation and the reference to compute a score. In contrast, COMET is a machine learning model fine-tuned from a pre-trained multilingual language model, e.g. XLM-R (Conneau et al., 2020), with an additional regression head. A reference-based COMET model learns to regress from a tuple of [source, hypothesis, reference] to the human judgement score (from previous evaluation campaigns) at the segment level. The quality estimation (reference-free) version of COMET is prepared by omitting the reference from the input.

Most issues in this work are demonstrated using two COMET checkpoints unless noted otherwise: the reference-based COMET<sub>22</sub><sup>DA</sup> and the reference-free COMET<sub>22</sub><sup>kiwiDA</sup> (Rei et al., 2022a). Both metrics output to a normalized range between 0 and 1. The COMET framework `unbabel-comet` is of version 2.2.2 except when we test how different software versions affect COMET scores.

**Data setup.** We base our experiments on the general domain translation and metrics shared tasks of WMT from 2023 (Kocmi et al., 2023; Freitag et al., 2023). The translation directions in the paper are centred around En↔De and En↔Zh, though we occasionally include other translation directions for demonstrative purposes.

Whenever possible, we divide all scores, including DA and model outputs, such that their output is between 0 and 1.

### 3 Problems

In this section, we identify and test nine possible pitfalls or curious behaviours with COMET which are not all well-studied. In three groups, these are:

- **Technicality:** obsolete Python and COMET software versions as well as compute precisions could lead to inaccurate score computation.
- **Training and test data:** COMET as a neural metric, might be derailed by empty hypotheses, language mismatch, and translationese at test time. It may also follow the training data biases.
- **Tool usage and score interpretation:** COMET has no defined way of equipping multiple references when available which leaves room for research. From a bibliometric perspective, we reveal that some literature omits a clear reference to the checkpoint version or citation.

In addition, we discuss some final issues that need more attention from the community.

Python	3.7.16	3.8.11	3.12.4
<code>unbabel-comet</code>	1.1.2	2.2.2	2.2.2
En→De	0.796	0.837	0.837
En→Zh	0.911	0.862	0.862
De→En	0.851	0.855	0.855
Zh→En	0.795	0.803	0.803

Table 1: COMET<sub>22</sub><sup>DA</sup> scores for WMT 23 Online-A under different package versions.

#### 3.1 Software versions [technical]

The official installation of the COMET package requires Python 3.8 or above.<sup>1</sup> We demonstrate that neglecting this would lead to unexpected scores because the same COMET checkpoint can produce vastly different scores with previous COMET framework versions that are no longer supported.

Under several Python versions, executing the following code leads to different COMET package (`unbabel-comet`) versions being installed. Running the framework for translation evaluation will subsequently result in false conclusions as shown in Table 1’s evaluation on WMT23 tests. The direct cause is that Python 3.7, which has been discontinued, only supports `unbabel-comet` versions up to 1.1.2. Nonetheless, we caution that the underlying factor is the version of `unbabel-comet` rather than Python.

```
$ pip install pip --upgrade
$ pip install unbabel-comet --upgrade

# will install
# unbabel-comet==1.1.2 under Python 3.7.16
# unbabel-comet==2.2.2 under Python 3.8.11
# unbabel-comet==2.2.2 under Python 3.12.4
```

**Recommendation.** Updating both Python and `unbabel-comet` to their latest versions is helpful and reporting the toolkit version aids reproducibility.

#### 3.2 Numerical precision [technical]

Model quantization represents a model using lower-numerical precision data types so that the model consumes less memory and model passes can be computed faster. Such improvement in inference is directly beneficial to deployment efficiency; it is also useful in other complex procedures involving COMET scoring, such as data filtering, re-ranking, and Minimum Bayes Risk (MBR) decoding (Kumar and Byrne, 2004).

Despite the aforementioned advantages, model quantization is not a feature supported by the cur-

<sup>1</sup>[github.com/Unbabel/COMET](https://github.com/Unbabel/COMET) 332dfb0 as of Aug 2024.

	Precision	COMET <sub>22</sub> <sup>DA</sup>	MAE	$\tau_c$	Acc	Time (s)
En→De	GPU FP32	0.822	10.4	0.274	0.885	113
	GPU FP16	0.822	10.4	0.274	0.885	55
	GPU FP32	0.822	10.4	0.274	0.885	2262
	CPU FP16	0.822	10.4	0.274	0.885	2403
	CPU QINT8	0.852	10.8	0.109	0.385	1856
De→En	GPU FP32	0.841	9.98	0.296	0.901	87
	GPU FP16	0.841	9.98	0.296	0.901	48
	GPU FP32	0.841	9.98	0.296	0.901	1674
	CPU FP16	0.841	9.99	0.296	0.901	1758
	CPU QINT8	0.860	10.7	0.164	0.516	1249
En→Zh	GPU FP32	0.842	11.7	0.290	0.933	111
	GPU FP16	0.842	11.7	0.290	0.933	102
	GPU FP32	0.842	11.7	0.290	0.933	1751
	CPU FP16	0.842	11.7	0.290	0.933	1710
	CPU QINT8	0.881	13.2	0.031	0.608	1258
Zh→En	GPU FP32	0.799	9.95	0.153	0.717	113
	GPU FP16	0.799	9.95	0.153	0.725	86
	GPU FP32	0.799	9.95	0.153	0.717	1936
	CPU FP16	0.799	9.95	0.153	0.717	1995
	CPU QINT8	0.872	10.5	0.081	0.475	1351

Table 2: System ranking with quantization on GPU and CPU. COMET<sub>22</sub><sup>DA</sup> is the absolute model score; MAE,  $\tau_c$ , and Acc are mean average error, correlation, and accuracy with respect to human judgements; Time refers to computation time in seconds.

rent COMET framework except in a concurrent work (Gowda et al., 2024). We make minimal modifications to the software and investigate the effect of numerical precision on COMET scores on both CPU (FP32, FP16, and QINT8) and GPU (FP32 and FP16). When using FP16, we first load the model weights to FP32, followed by `.half()` call. This is because loading the weights directly in FP16 still incorrectly results in FP32 precision. For CPU inference with dynamic QINT8, we apply the quantization module `torch.ao.quantization` from [PyTorch](#).

We use AMD Ryzen 9 5900X with NVIDIA GeForce RTX 3090 for GPU inference and a batch size of 8 in all settings (in practice a quantized model makes room for a larger batch size). Table 2 summarises the effect of numerical precision. In addition to reporting COMET scores, we also report (1) inference time in seconds (sec) as an efficiency measure; and (2) segment-level mean absolute error (MAE), segment-level Kendall’s tau-c ( $\tau_c$ ), and system-level pairwise accuracy (Acc). Everything is compared to the human DA scores either on segment- or system-level. Technically, Kendall’s  $\tau_c$  calculates rank correlation on an ordinal scale with adjustments for ties and pairwise accuracy computes the proportion of system pairs that have the same ordering by a metric as by humans.

Our results show that there is no meaningful difference between FP32 and FP16 in both CPU and GPU devices up to 3 significant figures. On GPU, FP16 is about 30% faster in time, but unsurprisingly it does not provide any speed-up on CPU. Interestingly, on the CPU, dynamic QINT8 gives systematically higher COMET scores and shorter running times than FP32 and FP16. However, the much lower  $\tau_c$  and pairwise accuracy indicate the lack of reliability at this precision. In addition to precision, we explored the effect of batch size and the choice of GPU or CPU during inference with results listed Appendix A. Whilst there are some fluctuations, they are mostly negligible. However, lower precision allows for higher batch size which usually directly corresponds to speed-up.

**Recommendation.** If GPU is available, it is feasible to run COMET with FP16 with a larger batch size for much faster inference without any quality loss. Otherwise, FP32 should be used.

### 3.3 Empty hypothesis [data]

An empty translation (a string of length 0) gets penalized heavily by string-based metrics because an empty string has zero surface overlap with the reference. However, neural metrics provide no such guarantee. We show that COMET assigns a positive instance-level score even if the hypothesis is an empty string as corroborated by Lo et al. (2023). Problematically, this score can even occasionally be higher than that of a genuine system translation.

In Table 3, we list COMET scores for system Online-A’s hypotheses at WMT23 and a file full of empty lines. Furthermore, we compare them with completely incorrect translations to explicate the score magnitude in two ways:

- *Random hypothesis*: we shuffle WMT22’s reference files at the sentence level in the respective translation directions. This provides us with high-quality human-written sentences. We sub-sample or over-sample if the number of lines in WMT22 is larger or smaller than the WMT23 size.
- *Random hypothesis (shuffled words)*: we further shuffle the words at each line in the sentence-shuffled files, generating nonsensical sentences. Sentence-shuffled hypotheses can be seen as fluent but extremely inadequate sentences whereas word-shuffled sentences are neither fluent nor adequate.

We observe that sentence-shuffled hypotheses attain comparable scores to empty ones, but word-shuffled hypotheses have the lowest scores across

Hypothesis ( $\downarrow$ )	COMET <sub>22</sub> <sup>DA</sup>				COMET <sub>22</sub> <sup>kiwiDA</sup>			
	En→De	En→Zh	De→En	Zh→En	En→De	En→Zh	De→En	Zh→En
Real system (Online-A)	0.837	0.862	0.855	0.803	0.800	0.791	0.794	0.787
Empty hypothesis	0.335	0.392	0.353	0.374	0.315	0.319	0.537	0.467
Random hypothesis (fluent)	0.373	0.434	0.334	0.350	0.333	0.341	0.447	0.391
Random hypothesis (shuffled words)	0.244	0.419	0.264	0.347	0.232	0.325	0.307	0.385

Table 3: Absolute average COMET scores for WMT23 Online-A, empty hypotheses, and random sentences. Random sentences are either fluent but irrelevant or perturbed with words shuffled and thus non-fluent.

the majority of the translation directions. Empty and shuffled hypotheses, despite having much lower COMET than the valid translations, would not be assigned zero scores by COMET<sub>22</sub><sup>DA</sup> or COMET<sub>22</sub><sup>kiwiDA</sup>, showing that COMET is more lenient than string overlap-based metrics in penalizing such irregularities.

We count the number of empty lines that score better than a translation from Online-A in Table 4. We observe roughly 0.25% such cases for most translation directions except for De→En’s COMET<sub>22</sub><sup>kiwiDA</sup> score situating at 1.45%. Further, in Figure 1 we plot the distributions of COMET scores for Zh→En’s empty and genuine translations with other translation directions in Appendix B. Noticeable overlaps are observed for COMET<sub>22</sub><sup>kiwiDA</sup> when translating into English.

Adhering to the DA protocol guidelines, this is not the proper behaviour because an empty hypothesis should receive a score of 0, to match *no meaning preserved*. This however is unsurprising with COMET, which has likely not seen empty hypotheses during training that would have received a score of 0 from a human annotator. Finally, even by relaxing the 0-score expectation, the metric should still assign the same score to all empty hypotheses regardless of the source. Since the distributions of empty hypotheses are nowhere close to a single vertical bar in Figure 1, it exposes the issue that segment-level COMET scores oddly hinge on the source sentence, as noted by Sun et al. (2020).

**Recommendation.** Force empty hypotheses to have 0 scores before aggregating. Also, a string-based metric like BLEU or chrF should be used to catch similarly malformed hypotheses.

### 3.4 Hypothesis language mismatch [data]

String overlap-based metrics can also score a hypothesis in a language different from the reference almost zero, especially with script mismatch. However, even for the reference-based COMET, there is no explicit way to enforce the intended target

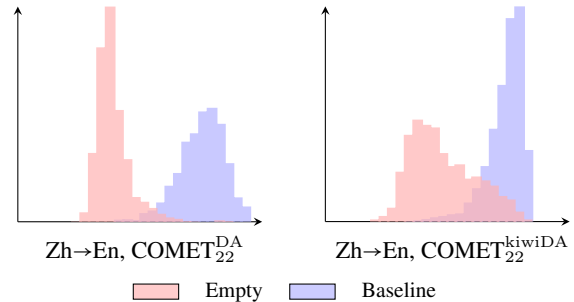


Figure 1: Distribution of instance-level scores for empty and baseline translations (x-axis: score; y-axis: count). See other translation directions in Appendix B.

	translation < empty	
	COMET <sub>22</sub> <sup>DA</sup>	COMET <sub>22</sub> <sup>kiwiDA</sup>
En→De	0 / 558	2 / 558
En→Ru	1 / 2075	0 / 2075
En→Uk	2 / 2075	1 / 2075
En→Zh	6 / 2075	2 / 2075
De→En	1 / 550	8 / 550
Ru→En	5 / 1724	6 / 1724
Uk→En	1 / 1827	5 / 1827
Zh→En	5 / 1977	1 / 1977

Table 4: Proportion of WMT23 Online-A’s translations that are worse than an empty line for the same source, displayed as “empty/total”.

language. This poses an increasingly pronounced problem, especially for multilingual translation models as well as the recent large language models, in which the generated language cannot be as easily controlled (Zhang et al., 2023).

We conduct experiments to understand if translation outputs in an incorrect language impact the score, and whether different mismatching languages can lead to distinct patterns. We use the translation directions En→Ru, En→Uk, and En→Zh in WMT23 which share the same English source input. Having Online-A’s output in all three directions, we substitute hypotheses in a particular translation direction with those from another direction. A similar hypothesis was presented by Amrhein et al. (2022) which suggested that COMET metrics are not robust to hypothesis language mis-

	En→Ru	En→Uk	En→Zh
Correct lang.	0.853	0.832	0.862
Incorrect target lang.	Uk: 0.797 Zh: 0.536	Ru: 0.807 Zh: 0.540	Ru: 0.655 Uk: 0.644
Empty hyp.	0.316	0.329	0.391
Random hyp.	0.463	0.472	0.435

Table 5: COMET<sub>22</sub><sup>DA</sup> scores for WMT23 Online-A’s output in (1) correct, (2, 3) incorrect language, (4) empty outputs, (5) random, but fluent, output.

match. Our experiment setup offers a more detailed evaluation setup than their contrastive setup.

Table 5 presents COMET<sub>22</sub><sup>DA</sup> scores for translations in correct and incorrect languages, as well as empty lines and random sentences in the correct language as “baselines”, deemed as completely wrong translations. The pattern shows that when the hypothesis is in a language distant from the reference, the COMET score declines much more than when the hypothesis is in a similar language. A more concerning issue is that even when the hypotheses’ language is completely wrong, the resulting COMET score can still be vastly higher than empty hypotheses or random sentences in the correct language. We omit COMET<sub>22</sub><sup>kiwiDA</sup> because it does not have a mechanism to read a reference (language) making it inherently incapable of distinguishing output languages.

**Recommendation.** Run language identification and set hypotheses in an unexpected language to have a 0 COMET score before aggregating them system-level. Also, check with a string overlap-based metric like BLEU or chrF.

Lang	Score	Lang	Score	Lang	Score
De→En	0.754	En→Is	0.666	Hi→Bn	0.910
Ps→En	0.670	En→Lt	0.600	De→Fr	0.792
Is→En	0.724	En→Ru	0.765	Fr→De	0.834
Pl→En	0.761	En→Iu	0.720	Zu→Xh	0.639
Ru→En	0.771	En→Ha	0.768	De→Cs	0.510
Ja→En	0.663	En→Ja	0.745	Xh→Zu	0.574
Ta→En	0.655	En→Pl	0.706	Bn→Hi	0.770
Zh→En	0.743	En→Gu	0.514		
Ha→En	0.641	En→Cs	0.767		
Km→En	0.659	En→Zh	0.775		
Lt→En	0.726	En→Fi	0.616		
Cs→En	0.740	En→Ta	0.709		
Gu→En	0.575	En→De	0.841		
Kk→En	0.649	En→Kk	0.574		
Iu→En	0.724				
Fi→En	0.719				

Table 6: Average human DA score for each translation direction in WMT data up to 2023 (inclusive).

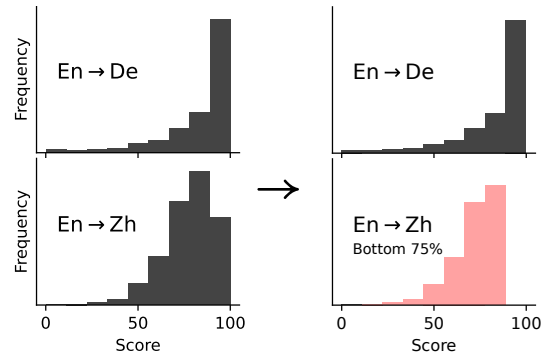


Figure 2: Setup of an experiment with bottom 75% of En→Zh scores which creates a bias in COMET<sub>22</sub><sup>DA</sup>. In the new data for En→Zh (bottom right) there are no translations with perfect scores. En→De data are unaffected.

### 3.5 Score distribution bias [data]

As the COMET metric is a machine learning model trained on human ratings of existing machine translations, it inherits many properties of statistical learning such as data (output score) distribution bias. The yearly WMT shared task receives submissions with varying quality and potentially varying quality ranges for different translation directions per year (from Koehn and Monz, 2006 to Freitag et al., 2023). This is attributed to diverse factors: the availability of data, source-target language similarity, the level of interest in languages, etc. The gap in translation quality will then propagate into skewed human judgement scores across translation directions. When a single COMET model learns to score all translation directions, it can overfit the score distribution w.r.t. a translation direction in addition to the quality of a translation hypothesis.

We first verify this in Table 6 which shows that WMT translation directions are associated with vastly different human DA scores (from 0.51 to 0.91). Empirically, we illustrate this issue using two high-resource directions En→De and En→Zh. As shown in Figure 2, we keep either the top- or bottom-75% of all scored translations to alter the score distribution for each direction, simulating the scenario where low- and high-performing system submissions are received for different directions. We then train different COMET models on the human train data before and after alteration as per Figure 2. Finally, we evaluate those checkpoints on the same test set and report results in Table 7.

As expected, for both En→De and En→Zh, training on the top or bottom-scoring data leads to increased or decreased COMET scores on the same set of hypotheses. Besides, we observe that altering scores in a particular translation direction incon-

Training data	En→De	En→Zh
All	0.770	0.770
Top-75% of En→De	0.790	0.770
Bot-75% of En→De	0.765	0.764
Top-75% of En→Zh	0.783	0.789
Bot-75% of En→Zh	0.772	0.751

Table 7: Average scores from COMET<sub>22</sub><sup>DA</sup> trained on data with top- or bottom-75% of scores kept in a particular direction.

sistently affects scores in another direction systematically. For example, removing the bottom 25% En→De scores “improves” the test-time score from 0.770 to 0.790 whilst En→Zh remains unaffected.

Finally, an empty translation should have the same score irrespective of the original source. Due to the statistical learning nature of the metric, this is not the case, as found in Section 3.3. Sun et al. (2020) and Zouhar et al. (2023) show that some meaningful correlation with human scores can be attained with just the source as the input. This shows, that there is a learned bias based on the prior difficulty of the source segment, which is undesirable for an objective evaluation metric.

**Implication.** A trivial conclusion is that COMET scores for different translation directions are not comparable. Nevertheless, we caution that the same phenomenon could happen for other features such as the domain, output style, etc. Although z-score re-scaling *could* mitigate this problem, it has not been a common practice since WMT22 and it would further contribute to non-objective scores (Knowles, 2021). Moreover, while z-scoring is straightforward for the translation direction, other latent, language-agnostic biases still exist.

### 3.6 Domain bias [data]

Neural metrics like COMET are biased towards particular domains, manifested by worse test performance on unseen domains (Zouhar et al., 2024a). Taking inspiration from previous work and our discussions on “latent biases”, we now raise a question—can we create adversarial hypotheses at test time to exploit the domain bias in training time? Specifically, different domains in the training data are associated with different score ranges. By pretending that a translation is in a particular domain, it might manipulate its COMET score.

To make it explicit to COMET during training, we prepend the target translation with a tag of its domain—in our case, the year the WMT data originated. Note that in each iteration of WMT, systems

<b>Train</b>	2020 Fire prevented from spreading
<b>Test</b>	2019 Now I have to tell you a nice story. 2020 Now I have to tell you a nice story. 2021 Now I have to tell you a nice story. 2022 Now I have to tell you a nice story. 2023 Now I have to tell you a nice story.

Table 8: An illustration of year-as-a-domain tagging during training and testing.

Tag	Train	Test
2018	unseen	0.736
2019	0.721	0.737
2020	0.735	0.744
2021	0.749	0.749
2022	unseen	0.747
2023	unseen	0.747
2024	unseen	0.739
2025	unseen	0.747

Table 9: Average COMET<sub>22</sub><sup>DA</sup> scores for subsets in training and predictions on test data. During testing, the whole test set had a single tag, e.g. 2024, irrespective of the data origin.

get higher overall DA scores (e.g. 0.721 for 2019 and 0.749 for 2021). Table 8 illustrates our setup: during training, we tag the scored translation data with its year; during testing, we trial various year prefixes to understand the effect.

One would expect the metric to produce the same score based solely on the translation quality. However, as shown in Table 9, by merely changing the year tag, we can influence the average score of the test set. During training, the model would be able to observe that 2019 is associated with the worst score and 2021 the best. During test time, the model follows this bias and also extrapolates it to upcoming years where it predicts an improvement in the average DA scores. While the differences appear small, they are on the same scale as the differences between years in the training data.

**Implication.** Our year-as-a-domain setting might be overly simple, but the vulnerability of COMET to latent biases cannot be neglected. Although Amrhein and Sennrich (2022) has shown that COMET is not sensitive to numbers, this work reveals that it can be systematically exploited in an artificial setting. We offer a more practical (adversarial) example that one may disguise biomedical domain translations as news translations to game COMET.

### 3.7 Lack of multi-reference support [usage]

In machine translation, there usually exist many valid translations for the same input. An effective metric should incorporate multiple ground truths

WMT23		COMET <sub>22</sub> <sup>DA</sup>				
		ref	ref (alt)	avg	max	agg
MQM	He→En	0.885	0.897	0.910	0.910	<b>0.949</b>
	En→De	0.974	0.936	0.974	<b>0.987</b>	0.974
	Zh→En	0.783	0.908	0.850	0.858	<b>0.950</b>
DA	En→De	0.885	<b>0.949</b>	0.910	0.897	0.910
	Zh→En	0.717	0.875	0.783	0.775	<b>0.883</b>
	De→En	0.901	0.912	<b>0.923</b>	<b>0.923</b>	0.912
	En→Zh	<b>0.933</b>	0.817	0.900	0.867	<b>0.933</b>
	Cs→Uk	0.846	0.802	<b>0.901</b>	0.890	0.868
	En→Cs	0.858	<b>0.875</b>	0.858	0.858	0.867
	En→Ja	<b>0.941</b>	0.824	0.934	0.934	0.926
	Ja→En	0.922	0.915	<b>0.928</b>	<b>0.928</b>	0.922

Table 10: Pairwise system-level accuracy for different strategies incorporating multiple references into COMET<sub>22</sub><sup>DA</sup>. Evaluation is carried out on WMT23 with DA or MQM scores as human ground truths.

if available, thereby enhancing the accuracy and robustness of its evaluation. Existing metrics like BLEU or chrF rely on surface-level overlap to capture the ground truth space from multiple references while metrics like ParBLEU (Bawden et al., 2020) can automatically generate paraphrases of a given reference to be included during evaluation.

By design, only one reference can be used in COMET. Whilst one may argue that representing a text reference in the neural space can ease the restriction on word choices, it might still be beneficial to use multiple references to overcome defects in the base embedding model. Therefore, we test whether COMET can explicitly and reliably leverage multiple references. We identify three distinct ways in which multiple references have been incorporated in COMET in previous literature (Rei et al., 2020b; Zouhar and Bojar, 2024):

- **max**: Taking the maximum over the scores from multiple passes with different references.
- **avg**: Averaging the scores from multiple passes with different references.
- **agg**: Obtaining an aggregate score per example as follows. A quadruplet of source  $s$ , hypothesis  $h$ , reference  $r$ , and alternative reference  $\hat{r}$  is fed to COMET six times in different [src, hyp, ref] arrangements:  $[s, h, r]$ ,  $[r, h, s]$ ,  $[s, h, \hat{r}]$ ,  $[\hat{r}, h, s]$ ,  $[r, h, \hat{r}]$ , as well as  $[\hat{r}, h, r]$ . Then, the average score from these six passes is multiplied by  $(1 - \sigma)$  where  $\sigma$  denotes the standard deviation.

We use additional references from the WMT23 test set if available (He→En) or the outputs from the best-scoring system in each direction in the metrics shared task (Freitag et al., 2023) as an alter-

native reference. We report pairwise system-level accuracy (Kocmi et al., 2021) for various translation directions in Table 10. Our results suggest that there is no single method that can consistently take advantage of the inclusion of multiple references with the existing COMET implementation. At a higher inference cost, the six-pass aggregation with COMET might have a tiny edge over other methods when MQM is treated as human ground truths, but it is also outperformed under DA by single-reference or other multi-reference methods.

As translation systems have greatly improved lately, the above pattern might be explained by Freitag et al. (2020)’s finding that high-quality translation outputs do not benefit from multi-reference evaluation. We also caution that these observations are highly dependent on the quality of the underlying references. As studied previously, obtaining high-quality references is not trivial (Freitag et al., 2020, 2023; Zouhar and Bojar, 2024). Our use of the top-performing system outputs as alternate references is fit for the purpose but not optimal.

**Recommendation.** Our recommendations for the inclusion of multiple references into COMET or even other neural metrics are aspirational as this topic warrants further investigation. Extending unified pre-training (Wan et al., 2022) with multiple references in the architecture as well as using training objectives more suitable for handling more than one references (Zheng et al., 2018; Fomicheva et al., 2020a) can be helpful.

### 3.8 Translationese [data]

COMET has been trained with human translations as references and machine translations as hypotheses, where both could be deemed “translationese” to a certain extent (Gellerstam, 1986).

**Translationese in references.** We first conduct an experiment to see if the translationese present in the reference would undermine system evaluation with COMET. We consider WMT’s official reference as a standard version and Freitag et al. (2020)’s paraphrased reference as a less translationese reference (we use their “paraphrased as-much-as-possible” version). Experiments are carried out under two settings: (1) WMT19 En→De submissions scored by COMET<sub>22</sub><sup>DA</sup>, and (2) WMT20 En→De submissions scored by COMET<sub>20</sub><sup>DA</sup> (Rei et al., 2020b). These two settings cover two scenarios—whether the test suite has been used in training the COMET model, or not. A breakdown

Please paraphrase the following text as much as possible. Provide the paraphrase without any explanation:

\$HYPOTHESIS

Figure 3: Prompt template used to request a paraphrase from GPT-4o, where \$HYPOTHESIS is replaced by individual hypotheses.

of COMET scores and rankings for individual systems are listed in Appendix C Tables 12 and 13.

The COMET scores decline dramatically when we switch the reference from the original one to the paraphrased one—aiming to reduce translationese. It means that COMET is indeed sensitive to such changes in the reference. Yet interestingly, the overall system ranking in either setting remains rather stable. We find a very high Kendall’s  $\tau_a$  of 0.9827 and 0.9833 on the system rankings in the two settings; pairwise accuracy computed against human judgements also maintained at 0.924. We conclude that translationese in the reference impacts absolute COMET scores but not system ranking. These patterns are consistent regardless of whether the model has been exposed to the test set.

**Translationese in hypotheses.** We then attempt to understand if a varying degree of translationese in the system outputs will influence system ranking by COMET. We run  $\text{COMET}_{22}^{\text{DA}}$  and  $\text{COMET}_{22}^{\text{kiwiDA}}$  on WMT19 En→De system outputs as well as their corresponding rephrased outputs against the same source and reference. To acquire paraphrases affordably, we shortlist the top-10 systems’ translations in the previous experiment and we prompt GPT-4o using the prompt outlined in Figure 3.<sup>2</sup> We do not feed the source sentence to prevent the model from revising the quality.

Appendix Table 14 shows that both models yield the same system ranking when the original hypotheses are scored. After substituting the hypotheses with their paraphrases, the ranking has changed more under  $\text{COMET}_{22}^{\text{kiwiDA}}$  which witnesses much lower pairwise accuracy and Kendall’s  $\tau_c$  compared to  $\text{COMET}_{22}^{\text{DA}}$ . This suggests that  $\text{COMET}_{22}^{\text{kiwiDA}}$  is more sensitive to potential changes in the degree of translationese than  $\text{COMET}_{22}^{\text{DA}}$ .

**Considerations.** We note the limitations of our experiments. First, we assume that the paraphrased references are as good as the original ones and less translationese, but we did not verify this when para-

phrasing the hypotheses. If the hypothesis quality has been affected, we also assume that the LLM paraphrasing process affects all system outputs in an equal magnitude. Second, the evaluations that anchor to human judgements assume that human evaluators provide assessment solely on the quality and do not overly insist on adequacy/translationese. Third, our comparison between  $\text{COMET}_{22}^{\text{DA}}$  and  $\text{COMET}_{22}^{\text{kiwiDA}}$  only shows that they do not behave the same in dealing with change in the degree of translationese in hypotheses.

### 3.9 Model reporting [usage]

Different COMET models can yield distinct results. Therefore it is important to always specify the specific model for sensible score interpretation and comparison. In this section, we examine to what extent this holds up in scientific literature.

We automate this bibliometric task with SemanticScholar API (Kinney et al., 2023). Starting with 1100 papers that cite one of the COMET papers, we obtain 417 papers from 2021 to 2024 that have an easily accessible PDF version.<sup>3</sup> We check if any of the *tables* contains the string `comet`. Within those papers, we check whether the COMET model information is contained in the PDF using a regular expression.<sup>4</sup> After further manual validation, we found that 50 of the examined papers do not report a specific COMET version. This establishes that *at least 12% of papers report COMET scores without specific model information*.

In addition, out of the almost 1000 papers running COMET in their evaluation, most only cite the first COMET paper (Rei et al., 2020a) instead of the paper that describes the specific models that are being used (Rei et al., 2020b, 2022b,a,c, 2023a,b; Glushkova et al., 2021; Wan et al., 2022; Alves et al., 2024; Guerreiro et al., 2023).

**Recommendation.** Always report the COMET version, ideally with a link. Also, cite the affiliated COMET paper as opposed to the first paper (Rei et al., 2020a), because different checkpoints have variations in training regimes that might be crucial in analysing the evaluation outcome.

<sup>3</sup>Papers in 2020 did not have to report the specific model as there was only one available at the time. Further, we acknowledge potential bias to only papers with available PDFs.

<sup>4</sup>Case-insensitive: `"comet[ \-](da|20|21|22|23)|wmt(20|21|22|23)\-comet|xcomet\-\wmt\-\da\-\estimator"`

<sup>2</sup>We accessed gpt-4o-2024-08-06 via API in Aug 2024.



### 3.10 Discussions on other issues

**Interpretation of significance.** Statistical hypothesis test merely shows how likely the difference between the average of two model’s scores on the same test set is caused by random fluctuation. [Kocmi et al. \(2024c\)](#) shows the significance of a difference between two metric scores can be made arbitrarily high and one can force  $p \rightarrow 0$  by using a sizeable test set. This tells little about whether this difference is meaningful to a human reader. For this reason, we stress the use of `mt-thresholds` that converts differences in metric scores to how perceivable they are by human annotators.<sup>5</sup>

**Averaging and subtracting COMET scores.** Research nowadays favours experiments on multiple translation directions, as multilingual translation models and large language models become trendy. Recent papers are more often seen to report a (macro-)average COMET score as an aggregate measure across many directions, usually in  $Xx \rightarrow En$ ,  $En \rightarrow Xx$ , and  $All \rightarrow All$ . Whilst indicative, this is not entirely scientific because (1) the score range is inherently distinct for each translation direction and (2) there is no assurance that the scores are on a linear scale. Consequently, an outlying score in a single direction can distort the average, leading to a false claim. Likewise, absolute COMET score differences are not comparable if from different base numbers or directions. We suggest that, when multiple translation directions are of interest, in addition to averaged scores, practitioners can report the number of wins (against another system) as another aggregation of individual scores.

**Optimizing to COMET.** Owing to COMET’s strong correlation with human judgement, recent works investigate the feasibility of using it in translation modelling directly. These strategies either include COMET in a distillation workflow ([Finkelstein and Freitag, 2024](#); [Guttmann et al., 2024](#)), as a data filtering method ([Peter et al., 2023](#)), as a decoding method ([Freitag et al., 2022a](#); [Fernandes et al., 2022](#); [Vernikos and Popescu-Belis, 2024](#)) or as a training objective ([Yan et al., 2023](#)).

Nevertheless, COMET may cease to be a good measure if practitioners over-optimize a system towards it. As [Yan et al. \(2023\)](#) demonstrated, a model trained towards COMET can generate “universal translations” (hallucinations) preferred by COMET regardless of the source sentence. Re-

cently, using COMET-based MBR decoding has become prevalent in shared tasks. For example, Unbabel-Tower70B at WMT24 also used MBR and dominated all automatic metrics but not so much under human evaluation ([Kocmi et al., 2024a,b](#)). MBR decoding could be seen as an automatic way to exploit bias in the scoring method, (currently COMET in most cases), so practitioners need to be aware of its shortcomings, disclose the use of such, and base system building on multiple (less correlated) metrics ([Jon et al., 2023](#)).

A novel issue is using automated metrics in human evaluation ([Zouhar et al., 2024b](#)) that collects data for metric training. This might create a similar effect as translationese in machine translation. The data could be biased by the particular quality estimator that is assisting annotators in the data collection process.

**Sensitivity to sentence segmentation.** Like most MT metrics, COMET works at the sentence level, but sometimes sentence-segmented input is not available. This is often the case in speech translation (ST) where sentence segmentation is treated as part of the task ([Ahmad et al., 2024](#)). To address the problem of mismatching segmentation between the system output and the reference, a common solution in ST is to re-segment the output using a minimum error rate method ([Matusov et al., 2005](#)) in order to force-align it with the reference. Forced alignment can introduce segmentation errors resulting in truncated (thus grammatically incorrect) sentences. There is evidence that COMET, as a metric reliant on sentence embeddings, is more sensitive to segmentation errors than string-based metrics, like BLEU, which rely purely on  $n$ -gram overlaps with no linguistic notion of a sentence ([Amrhein and Haddow, 2022](#)). In a recent comparison of COMET with human ranking, [Sperber et al. \(2024\)](#) suggested that COMET-based ranking is robust to segmentation errors but that a “more thorough study of this issue is needed”.

**Other metrics.** Our work focused on COMET, the current most popular family of MT metrics. Nonetheless, our recommendations could apply to other neural metrics, like MetricX-23 ([Juraska et al., 2023](#)), because many issues we outlined are due to their statistical learning nature. Even beyond this, metric reporting and score interpretation in practice, e.g. software usage or averaging across multiple directions, can be problematic for string-matching metrics like BLEU or chrF too.

<sup>5</sup>[kocmitom.github.io/MT-Thresholds](https://kocmitom.github.io/MT-Thresholds)

## 4 The SacreCOMET Package

To help alleviate problems in Sections 3.1 (software version), 3.2 (compute precision), and 3.9 (model reporting), we release a simple package `sacreCOMET` with two functionalities. Given a model name, the first functionality attempts to find the appropriate citation including a link to the paper and a BibTeX:

```
$ pip install sacrecomet
$ sacrecomet cite Unbabel/xcomet-xl

https://arxiv.org/abs/2310.10482
@misc{guerreiro2023xcomet,
  title={xCOMET: Transparent Machine
    Translation Evaluation through Fine-
    grained Error Detection},
  ...
```

The second functionality semi-automatically detects the local software versions to generate a signature for better reproducibility. Both functionalities can also be run in interactive mode.

```
$ sacrecomet --model unite-mup --prec fp32
Python3.11.8|Comet2.2.2|fp32|unite-mup
```

## 5 Future Work on Learned Metrics

- **Fixing data bias:** Learned metrics are sensitive to the training data distribution. Future metrics should aim to reduce the bias caused by the data selection process such that they are applicable to a range of MT systems.
- **Interpretability across languages:** Currently, practitioners cannot compare, or pedantically, aggregate scores in different translation directions. It would be useful to unify the scores to a single scale that can be interpreted independent of the language (similar to Kocmi et al., 2024c), e.g. to indicate  $X\%$  of segments are production-ready.
- **Confidence-aware metrics:** As seen in works of Glushkova et al. (2021); Fomicheva et al. (2020b), it is possible to build metrics that output a confidence interval, though its usage in evaluation and modelling remains scarce.
- **Inference speed:** Learned metrics are getting better but at the cost of bulky models and increased inference time. These overheads should be taken into account when developing new models, such as the work of Rei et al. (2022b).

- **Representations:** Current COMET models are built upon off-the-shelf multilingual encoder models which are likely trained on human-written texts. However, this could bring in a domain mismatch—when translation hypotheses act as the input to metric models, they are not human-written but machine-translated.
- **Robustness:** Metrics should have the correct behaviour even in corner cases, be it empty output or incorrect language. Mapping all inputs (Amrhein et al., 2022, *inter alia*), including partial or adversarial ones, evaluating the metrics, and coming up with methods to make them more robust would increase the metrics’ adoption and trust.
- **Built-in QE:** In production, machine translation and quality estimation are commonly two different processes. In many applications, however, a single QE model is used for a single MT model. Quantifying how much is QE adaptation to a particular MT model useful is beneficial for a holistic understanding of QE metrics. Further, proposing methods for supervised quality estimation built into the MT could ease industry adoption. Beyond the work of Tomani et al. (2024), this remains largely unexplored.

- **Noise-aware training:** Human annotations are notoriously noisy. At the scale of WMT data, poor-quality annotations are unavoidable. The inter-annotator agreement for even robust annotations, such as ESA, remains low at  $\tau_c \approx 0.3$ . The effect of data quality on learned metrics is so far unknown and methods for noise/uncertainty-aware training are under-studied.

## 6 Conclusion

COMET is currently one of the most powerful automatic metrics/quality estimators for machine translation, consistently achieving the top correlation with human judgement. In comparison to previous metrics, it is a statistical learning model and thus inherits all the related problems in addition to possible technical misconfigurations. We urge practitioners to consider more deeply the use of COMET in non-standard scenarios especially where such training bias might come into play. Beyond these issues, there has been confusion in the literature in reporting the correct COMET model and its correct setting. For improved consistency, we release an easy-to-use tool to assist practitioners.

## Acknowledgments

The work has received funding from UK Research and Innovation (UKRI) under the UK government’s Horizon Europe funding guarantee [grant numbers 10052546 and 10039436]. We thank the anonymous reviewers for their insightful suggestions.

## References

- Ibrahim Said Ahmad, Antonios Anastasopoulos, Ondřej Bojar, Claudia Borg, Marine Carpuat, Roldano Cattoni, Mauro Cettolo, William Chen, Qianqian Dong, Marcello Federico, Barry Haddow, Dávid Javorský, Mateusz Krubiński, Tsz Kin Lam, Xutai Ma, Prashant Mathur, Evgeny Matusov, Chandresh Maurya, John McCrae, Kenton Murray, Satoshi Nakamura, Matteo Negri, Jan Niehues, Xing Niu, Atul Kr. Ojha, John Ortega, Sara Papi, Peter Polák, Adam Pospíšil, Pavel Pecina, Elizabeth Salesky, Nivedita Sethiya, Balaram Sarkar, Jiatong Shi, Claytone Siksote, Matthias Sperber, Sebastian Stüker, Katsuhito Sudoh, Brian Thompson, Alex Waibel, Shinji Watanabe, Patrick Wilken, Petr Zemánek, and Rodolfo Zevallos. 2024. [Findings of the IWSLT 2024 evaluation campaign](#). In *Proceedings of the 21st International Conference on Spoken Language Translation (IWSLT 2024)*.
- Duarte M. Alves, José Pombal, Nuno M. Guerreiro, Pedro H. Martins, João Alves, Amin Farajian, Ben Peters, Ricardo Rei, Patrick Fernandes, Sweta Agrawal, Pierre Colombo, José G. C. de Souza, and André F. T. Martins. 2024. [Tower: An open multilingual large language model for translation-related tasks](#). Preprint, arXiv:2402.17733.
- Chantal Amrhein and Barry Haddow. 2022. [Don’t discard fixed-window audio segmentation in speech-to-text translation](#). In *Proceedings of the Seventh Conference on Machine Translation (WMT)*.
- Chantal Amrhein, Nikita Moghe, and Liane Guillou. 2022. [ACES: Translation accuracy challenge sets for evaluating machine translation metrics](#). In *Proceedings of the Seventh Conference on Machine Translation (WMT)*.
- Chantal Amrhein and Rico Sennrich. 2022. [Identifying weaknesses in machine translation metrics through minimum Bayes risk decoding: A case study for COMET](#). In *Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*.
- Rachel Bawden, Biao Zhang, Andre Tättar, and Matt Post. 2020. [ParBLEU: Augmenting metrics with automatic paraphrases for the WMT’20 metrics shared task](#). In *Proceedings of the Fifth Conference on Machine Translation*.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*.
- Patrick Fernandes, António Farinhas, Ricardo Rei, José G. C. de Souza, Perez Ogayo, Graham Neubig, and Andre Martins. 2022. [Quality-aware decoding for neural machine translation](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*.
- Mara Finkelstein and Markus Freitag. 2024. [MBR and QE finetuning: Training-time distillation of the best and most expensive decoding methods](#). In *The Twelfth International Conference on Learning Representations*.
- Marina Fomicheva, Lucia Specia, and Francisco Guzmán. 2020a. [Multi-hypothesis machine translation evaluation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*.
- Marina Fomicheva, Shuo Sun, Lisa Yankovskaya, Frédéric Blain, Francisco Guzmán, Mark Fishel, Nikolaos Aletras, Vishrav Chaudhary, and Lucia Specia. 2020b. [Unsupervised Quality Estimation for Neural Machine Translation](#). *Transactions of the Association for Computational Linguistics*, 8:539–555.
- Markus Freitag, David Grangier, and Isaac Caswell. 2020. [BLEU might be guilty but references are not innocent](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Markus Freitag, David Grangier, Qijun Tan, and Bowen Liang. 2022a. [High quality rather than high model probability: Minimum Bayes risk decoding with neural metrics](#). *Transactions of the Association for Computational Linguistics*, 10.
- Markus Freitag, Nitika Mathur, Chi-kiu Lo, Eleftherios Avramidis, Ricardo Rei, Brian Thompson, Tom Kocmi, Frederic Blain, Daniel Deutsch, Craig Stewart, Chrysoula Zerva, Sheila Castilho, Alon Lavie, and George Foster. 2023. [Results of WMT23 metrics shared task: Metrics might be guilty but references are not innocent](#). In *Proceedings of the Eighth Conference on Machine Translation*.
- Markus Freitag, Ricardo Rei, Nitika Mathur, Chi-kiu Lo, Craig Stewart, Eleftherios Avramidis, Tom Kocmi, George Foster, Alon Lavie, and André F. T. Martins. 2022b. [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)*.
- Markus Freitag, Ricardo Rei, Nitika Mathur, Chi-kiu Lo, Craig Stewart, George Foster, Alon Lavie, and Ondřej

- Bojar. 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*.
- Martin Gellerstam. 1986. Translationese in Swedish novels translated from English. In L. Wollin and H. Lindquist, editors, *Translation studies in Scandinavia: Proceedings from the Scandinavian Symposium on Translation Theory (SSOTT) II*, number 75 in Lund Studies in English, page 88–95. CWK Gleerup, Lund.
- Taisiya Glushkova, Chrysoula Zerva, Ricardo Rei, and André F. T. Martins. 2021. [Uncertainty-aware machine translation evaluation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*.
- Thamme Gowda, Roman Grundkiewicz, Elijah Rippeth, Matt Post, and Marcin Junczys-Dowmunt. 2024. [Py-marian: Fast neural machine translation and evaluation in python](#). *Preprint*, arXiv:2408.11853.
- Yvette Graham, Timothy Baldwin, Alistair Moffat, and Justin Zobel. 2013. [Continuous measurement scales in human evaluation of machine translation](#). In *Proceedings of the 7th Linguistic Annotation Workshop and Interoperability with Discourse*.
- Nuno M. Guerreiro, Ricardo Rei, Daan van Stigt, Luisa Coheur, Pierre Colombo, and André F. T. Martins. 2023. [xcomet: Transparent machine translation evaluation through fine-grained error detection](#). *Preprint*, arXiv:2310.10482.
- Kamil Guttman, Mikołaj Pokrywka, Adrian Charkiewicz, and Artur Nowakowski. 2024. [Chasing COMET: Leveraging minimum bayes risk decoding for self-improving machine translation](#). *arXiv preprint arXiv:2405.11937*.
- Josef Jon, Martin Popel, and Ondřej Bojar. 2023. [CUNI at WMT23 general translation task: MT and a genetic algorithm](#). In *Proceedings of the Eighth Conference on Machine Translation*.
- Juraj Juraska, Mara Finkelstein, Daniel Deutsch, Aditya Siddhant, Mehdi Mirzazadeh, and Markus Freitag. 2023. [MetricX-23: The Google submission to the WMT 2023 metrics shared task](#). In *Proceedings of the Eighth Conference on Machine Translation*.
- Rodney Michael Kinney, Chloe Anastasiades, Russell Authur, Iz Beltagy, Jonathan Bragg, Alexandra Buraczynski, Isabel Cachola, Stefan Candra, Yoganand Chandrasekhar, Arman Cohan, Miles Crawford, Doug Downey, Jason Dunkelberger, Oren Etzioni, Rob Evans, Sergey Feldman, Joseph Gorney, David W. Graham, F.Q. Hu, Regan Huff, Daniel King, Sebastian Kohlmeier, Bailey Kuehl, Michael Langan, Daniel Lin, Haokun Liu, Kyle Lo, Jaron Lochner, Kelsey MacMillan, Tyler C. Murray, Christopher Newell, Smita R Rao, Shaurya Rohatgi, Paul Sayre, Zejiang Shen, Amanpreet Singh, Luca Soldaini, Shivashankar Subramanian, A. Tanaka, Alex D Wade, Linda M. Wagner, Lucy Lu Wang, Christopher Wilhelm, Caroline Wu, Jiangjiang Yang, Angele Zamaron, Madeleine van Zuylen, and Daniel S. Weld. 2023. [The semantic scholar open data platform](#). *ArXiv*, abs/2301.10140.
- Rebecca Knowles. 2021. [On the stability of system rankings at WMT](#). In *Proceedings of the Sixth Conference on Machine Translation*.
- Tom Kocmi, Eleftherios Avramidis, Rachel Bawden, Ondřej Bojar, Anton Dvorkovich, Christian Federmann, Mark Fishel, Markus Freitag, Thamme Gowda, Roman Grundkiewicz, Barry Haddow, Marzena Karpinska, Philipp Koehn, Benjamin Marie, Kenton Murray, Masaaki Nagata, Martin Popel, Maja Popovic, Mariya Shmatova, Steinþór Steingrímsson, and Vilém Zouhar. 2024a. [Preliminary wmt24 ranking of general mt systems and llms](#). *Preprint*, arXiv:2407.19884.
- Tom Kocmi, Eleftherios Avramidis, Rachel Bawden, Ondřej Bojar, Anton Dvorkovich, Christian Federmann, Mark Fishel, Markus Freitag, Thamme Gowda, Roman Grundkiewicz, Barry Haddow, Philipp Koehn, Benjamin Marie, Christof Monz, Makoto Morishita, Kenton Murray, Makoto Nagata, Toshiaki Nakazawa, Martin Popel, Maja Popović, and Mariya Shmatova. 2023. [Findings of the 2023 conference on machine translation \(WMT23\): LLMs are here but not quite there yet](#). In *Proceedings of the Eighth Conference on Machine Translation*.
- Tom Kocmi, Eleftherios Avramidis, Rachel Bawden, Ondřej Bojar, Anton Dvorkovich, Christian Federmann, Mark Fishel, Markus Freitag, Thamme Gowda, Roman Grundkiewicz, Barry Haddow, Marzena Karpinska, Philipp Koehn, Benjamin Marie, Christof Monz, Kenton Murray, Masaaki Nagata, Martin Popel, Maja Popović, Mariya Shmatova, Steinþór Steingrímsson, and Vilém Zouhar. 2024b. [Findings of the WMT24 general machine translation shared task: The llm era is here but MT is not solved yet](#). To be published at WMT 2024.
- Tom Kocmi, Christian Federmann, Roman Grundkiewicz, Marcin Junczys-Dowmunt, Hitokazu Matsushita, and Arul Menezes. 2021. [To ship or not to ship: An extensive evaluation of automatic metrics for machine translation](#). In *Proceedings of the Sixth Conference on Machine Translation*.
- Tom Kocmi, Vilém Zouhar, Christian Federmann, and Matt Post. 2024c. [Navigating the metrics maze: Reconciling score magnitudes and accuracies](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.
- Tom Kocmi, Vilém Zouhar, Eleftherios Avramidis, Roman Grundkiewicz, Marzena Karpinska, Maja Popović, Mrinmaya Sachan, and Mariya Shmatova. 2024d. [Error span annotation: A balanced approach for human evaluation of machine translation](#). *Preprint*, arXiv:2406.11580.

- Philipp Koehn and Christof Monz. 2006. [Manual and automatic evaluation of machine translation between European languages](#). In *Proceedings on the Workshop on Statistical Machine Translation*.
- Shankar Kumar and William Byrne. 2004. [Minimum Bayes-risk decoding for statistical machine translation](#). In *Proceedings of the Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics: HLT-NAACL 2004*.
- Chi-kiu Lo, Samuel Larkin, and Rebecca Knowles. 2023. [Metric score landscape challenge \(MSLC23\): Understanding metrics' performance on a wider landscape of translation quality](#). In *Proceedings of the Eighth Conference on Machine Translation*.
- Arle Lommel, Aljoscha Burchardt, Maja Popović, Kim Harris, Eleftherios Avramidis, and Hans Uszkoreit. 2014. [Using a new analytic measure for the annotation and analysis of MT errors on real data](#). In *Proceedings of the 17th Annual Conference of the European Association for Machine Translation*.
- E. Matusov, G. Leusch, O. Bender, and H. Ney. 2005. [Evaluating Machine Translation Output with Automatic Sentence Segmentation](#). In *Proceedings of the 2nd International Workshop on Spoken Language Translation (IWSLT)*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*.
- Jan-Thorsten Peter, David Vilar, Daniel Deutsch, Mara Finkelstein, Juraj Juraska, and Markus Freitag. 2023. [There's no data like better data: Using QE metrics for MT data filtering](#). In *Proceedings of the Eighth Conference on Machine Translation*.
- Maja Popović. 2015. [chrF: character n-gram F-score for automatic MT evaluation](#). In *Proceedings of the Tenth Workshop on Statistical Machine Translation*.
- Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*.
- Ricardo Rei, José G. C. de Souza, Duarte Alves, Chrysoula Zerva, Ana C Farinha, Taisiya Glushkova, Alon Lavie, Luisa Coheur, and André F. T. Martins. 2022a. [COMET-22: Unbabel-IST 2022 submission for the metrics shared task](#). In *Proceedings of the Seventh Conference on Machine Translation (WMT)*.
- Ricardo Rei, Ana C Farinha, José G.C. de Souza, Pedro G. Ramos, André F.T. Martins, Luisa Coheur, and Alon Lavie. 2022b. [Searching for COMETINHO: The little metric that could](#). In *Proceedings of the 23rd Annual Conference of the European Association for Machine Translation*.
- Ricardo Rei, Nuno M. Guerreiro, José Pombal, Daan van Stigt, Marcos Treviso, Luisa Coheur, José G. C. de Souza, and André F. T. Martins. 2023a. [Scaling up cometkiwi: Unbabel-ist 2023 submission for the quality estimation shared task](#). *Preprint*, arXiv:2309.11925.
- Ricardo Rei, Nuno M. Guerreiro, Marcos Treviso, Luisa Coheur, Alon Lavie, and André F. T. Martins. 2023b. [The inside story: Towards better understanding of machine translation neural evaluation metrics](#). *Preprint*, arXiv:2305.11806.
- Ricardo Rei, Craig Stewart, Ana C Farinha, and Alon Lavie. 2020a. [COMET: A neural framework for MT evaluation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Ricardo Rei, Craig Stewart, Ana C Farinha, and Alon Lavie. 2020b. [Unbabel's participation in the WMT20 metrics shared task](#). In *Proceedings of the Fifth Conference on Machine Translation*.
- Ricardo Rei, Marcos Treviso, Nuno M. Guerreiro, Chrysoula Zerva, Ana C Farinha, Christine Maroti, José G. C. de Souza, Taisiya Glushkova, Duarte Alves, Luisa Coheur, Alon Lavie, and André F. T. Martins. 2022c. [CometKiwi: IST-unbabel 2022 submission for the quality estimation shared task](#). In *Proceedings of the Seventh Conference on Machine Translation (WMT)*.
- 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*.
- Hiroki Shimanaka, Tomoyuki Kajiwara, and Mamoru Komachi. 2018. [RUSE: Regressor using sentence embeddings for automatic machine translation evaluation](#). In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*.
- Matthew Snover, Bonnie Dorr, Rich Schwartz, Linnea Micciulla, and John Makhoul. 2006. [A study of translation edit rate with targeted human annotation](#). In *Proceedings of the 7th Conference of the Association for Machine Translation in the Americas: Technical Papers*.
- Matthias Sperber, Ondřej Bojar, Barry Haddow, Dávid Javorský, Xutai Ma, Matteo Negri, Jan Niehues, Peter Polák, Elizabeth Salesky, Katsuhito Sudoh, and Marco Turchi. 2024. [Evaluating the IWSLT2023 speech translation tasks: Human annotations, automatic metrics, and segmentation](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*.
- Shuo Sun, Francisco Guzmán, and Lucia Specia. 2020. [Are we estimating or guesstimating translation quality?](#) In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*.

- Kosuke Takahashi, Katsuhito Sudoh, and Satoshi Nakamura. 2020. [Automatic machine translation evaluation using source language inputs and cross-lingual language model](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*.
- Christian Tomani, David Vilar, Markus Freitag, Colin Cherry, Subhajit Naskar, Mara Finkelstein, Xavier Garcia, and Daniel Cremers. 2024. [Quality-aware translation models: Efficient generation and quality estimation in a single model](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.
- Giorgos Vernikos and Andrei Popescu-Belis. 2024. [Don't rank, combine! combining machine translation hypotheses using quality estimation](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.
- Yu Wan, Dayiheng Liu, Baosong Yang, Haibo Zhang, Boxing Chen, Derek Wong, and Lidia Chao. 2022. [UniTE: Unified translation evaluation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.
- Yiming Yan, Tao Wang, Chengqi Zhao, Shujian Huang, Jiajun Chen, and Mingxuan Wang. 2023. [BLEURT has universal translations: An analysis of automatic metrics by minimum risk training](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.
- Biao Zhang, Barry Haddow, and Alexandra Birch. 2023. [Prompting large language model for machine translation: a case study](#). In *Proceedings of the 40th International Conference on Machine Learning*.
- Renjie Zheng, Mingbo Ma, and Liang Huang. 2018. [Multi-reference training with pseudo-references for neural translation and text generation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*.
- Vilém Zouhar and Ondřej Bojar. 2024. [Quality and quantity of machine translation references for automatic metrics](#). In *Proceedings of the Fourth Workshop on Human Evaluation of NLP Systems (HumEval) @ LREC-COLING 2024*.
- Vilém Zouhar, Shehzaad Dhuliawala, Wangchunshu Zhou, Nico Daheim, Tom Kocmi, Yuchen Eleanor Jiang, and Mrinmaya Sachan. 2023. [Poor man's quality estimation: Predicting reference-based MT metrics without the reference](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*.
- Vilém Zouhar, Shuoyang Ding, Anna Currey, Tatyana Badeka, Jenyuan Wang, and Brian Thompson. 2024a. [Fine-tuned machine translation metrics struggle in unseen domains](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*.
- Vilém Zouhar, Tom Kocmi, and Mrinmaya Sachan. 2024b. [AI-assisted human evaluation of machine translation](#). *Preprint*, arXiv:2406.12419.

## A Batch size and GPU/CPU

We run the test inference on a combination of GPU or CPU with varying batch sizes (BS, 1 or 100). Results in Table 11 demonstrate that the tiny effects of these choices are negligible for COMET reporting.

Difference	MAE
BS=1, GPU – BS=1, GPU	0
BS=1, GPU – BS=64 GPU	$2 \times 10^{-7}$
BS=1, GPU – BS=1, CPU	$4 \times 10^{-7}$
BS=64, GPU – BS=64, CPU	$4 \times 10^{-7}$

Table 11: MAE between segment-level  $\text{COMET}_{22}^{\text{DA}}$  scores under various inference settings. The “BS=1, GPU” setting in the first row was run twice.

## B Distribution of COMET scores for empty and valid hypothesis

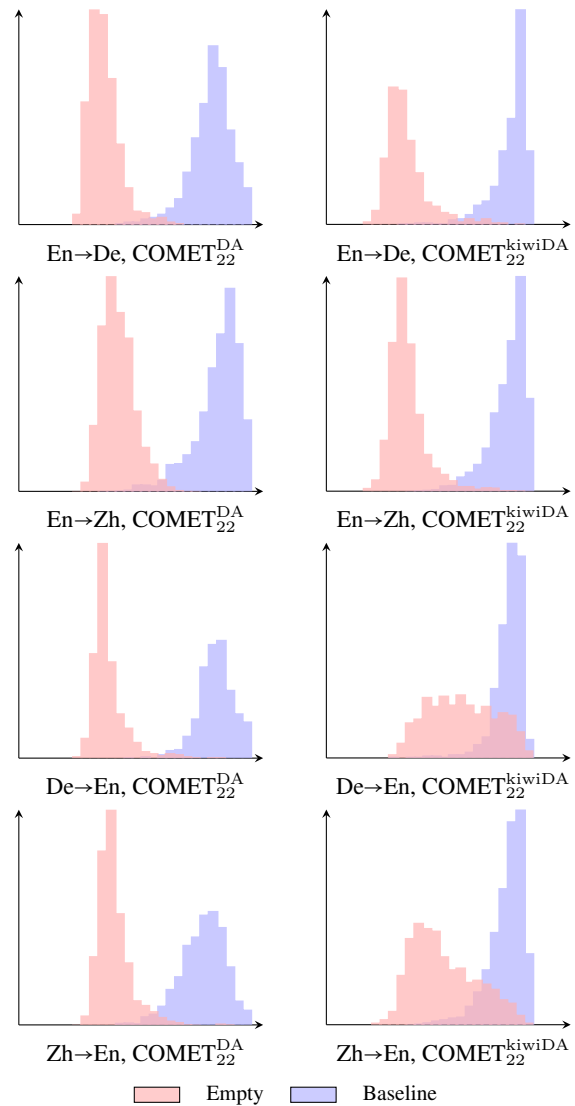


Figure 4: Distribution of instance-level scores for empty and baseline translations (x-axis: score; y-axis: count).

## C System rankings before and after paraphrasing references or hypotheses

WMT20 En→De	COMET <sub>20</sub> <sup>DA</sup>		Ranking	
	original ref	paraphrased ref	original ref	paraphrased ref
Sys-1069	0.508	0.305	9	9
Sys-832	0.540	0.333	8	8
Sys-1535	0.560	0.356	5	5
Online-A	0.499	0.288	10	10
Online-B	0.554	0.351	<b>7</b>	<b>6</b>
Online-G	0.268	0.052	14	14
Online-Z	0.329	0.121	13	13
Sys-73	0.405	0.192	12	12
Sys-1520	0.563	0.360	4	4
Sys-890	0.578	0.371	3	3
Sys-1136	0.472	0.264	11	11
Sys-388	0.102	-0.074	16	16
Sys-737	0.555	0.351	<b>6</b>	<b>7</b>
Ref-A	0.878	0.526	1	1
Ref-B	0.591	0.446	2	2
Sys-179	0.189	0.000	15	15

$\tau_c = 0.9833$   
 Acc = 0.924    Acc = 0.924

Table 12: Results for WMT20 En→De submissions evaluated against the original or human-paraphrased reference. Kendall’s  $\tau_c$  is measured between two evaluations based on original and paraphrased references; pairwise system-level accuracy (Acc) is measured against human DA scores.

WMT19 En→De	COMET <sub>22</sub> <sup>DA</sup>		Ranking	
	original ref	paraphrased ref	original ref	paraphrased ref
Sys-6862	0.867	0.817	2	2
Sys-6820	0.834	0.780	12	12
Sys-6819	0.843	0.789	9	9
Sys-6651	0.847	0.794	7	7
Sys-6926	0.852	0.797	<b>4</b>	<b>5</b>
Sys-6808	0.869	0.818	1	1
Sys-6785	0.837	0.785	11	11
Sys-6974	0.866	0.814	3	3
Sys-6763	0.851	0.797	<b>5</b>	<b>6</b>
Sys-6674	0.811	0.756	17	17
Sys-6508	0.804	0.752	19	19
Sys-6731	0.850	0.797	<b>6</b>	<b>4</b>
Sys-6871	0.809	0.756	18	18
Sys-6479	0.833	0.777	13	13
Sys-6823	0.845	0.792	8	8
Sys-6790	0.386	0.364	22	22
Sys-6981	0.826	0.774	14	14
Online-A	0.815	0.761	16	16
Online-B	0.838	0.784	10	10
Online-G	0.795	0.738	20	20
Online-X	0.728	0.673	21	21
Online-Y	0.821	0.762	15	15

$\tau_c = 0.9827$   
 Acc = 0.875    Acc = 0.845

Table 13: Results for WMT19 En→De submissions evaluated against the original or human-paraphrased reference. Kendall’s  $\tau_c$  is measured between two evaluations based on original and paraphrased references; pairwise system-level accuracy (Acc) is measured against human DA scores.



WMT19 En→De	COMET <sub>22</sub> <sup>DA</sup>		Ranking		COMET <sub>22</sub> <sup>kiwiDA</sup>		Ranking		
	orig. hyp	para. hyp	orig. hyp	para. hyp	orig. hyp	para. hyp	orig. hyp	para. hyp	
Sys-6823	0.845	0.840	<b>8</b>	<b>5</b>	0.821	0.824	<b>8</b>	<b>4</b>	
Sys-6862	0.867	0.842	2	2	0.840	0.828	<b>2</b>	<b>3</b>	
Sys-6819	0.843	0.835	9	9	0.815	0.817	<b>9</b>	<b>8</b>	
Sys-6808	0.869	0.843	1	1	0.840	0.828	<b>1</b>	<b>2</b>	
Sys-6974	0.866	0.842	3	3	0.838	0.829	<b>3</b>	<b>1</b>	
Sys-6651	0.847	0.835	<b>7</b>	<b>8</b>	0.821	0.817	<b>7</b>	<b>9</b>	
Sys-6926	0.852	0.839	<b>4</b>	<b>6</b>	0.824	0.822	<b>4</b>	<b>6</b>	
Sys-6763	0.851	0.841	<b>5</b>	<b>4</b>	0.823	0.824	5	5	
Online-B	0.838	0.830	10	10	0.805	0.810	10	10	
Sys-6731	0.850	0.837	<b>6</b>	<b>7</b>	0.822	0.820	<b>6</b>	<b>7</b>	
				$\tau_c = 0.822$ Acc = 0.911					
						$\tau_c = 0.644$ Acc = 0.822			

Table 14: Results for WMT19 En→De system outputs and LLM-paraphrased outputs evaluated against the original reference. Both Kendall’s  $\tau_c$  and pairwise system-level accuracy (Acc) are measured between two evaluations based on original and paraphrased references.