



HAL
open science

xCOMET: Transparent Machine Translation Evaluation through Fine-grained Error Detection

Nuno Guerreiro, Ricardo Rei, Daan Van, Pierre Colombo, Luisa Coheur,
André Martins

► To cite this version:

Nuno Guerreiro, Ricardo Rei, Daan Van, Pierre Colombo, Luisa Coheur, et al.. xCOMET: Transparent Machine Translation Evaluation through Fine-grained Error Detection. Transactions of the Association for Computational Linguistics, 2023. hal-04574986

HAL Id: hal-04574986

<https://hal.science/hal-04574986v1>

Submitted on 14 May 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

xCOMET: Transparent Machine Translation Evaluation through Fine-grained Error Detection

Nuno M. Guerreiro^{*1,3,4,5}, Ricardo Rei^{*1,2,5}, Daan van Stigt¹,
Luisa Coheur^{2,5}, Pierre Colombo⁴, André F. T. Martins^{1,3,5}

¹Unbabel, Lisbon, Portugal, ²INESC-ID, Lisbon, Portugal

³Instituto de Telecomunicações, Lisbon, Portugal

⁴MICS, CentraleSupélec, Université Paris-Saclay, France

⁵Instituto Superior Técnico, University of Lisbon, Portugal

Abstract

Widely used learned metrics for machine translation evaluation, such as COMET and BLEURT, estimate the quality of a translation hypothesis by providing a single sentence-level score. As such, they offer little insight into translation errors (e.g., what are the errors and what is their severity). On the other hand, generative large language models (LLMs) are amplifying the adoption of more granular strategies to evaluation, attempting to detail and categorize translation errors. In this work, we introduce xCOMET, an open-source learned metric designed to bridge the gap between these approaches. xCOMET integrates both sentence-level evaluation and error span detection capabilities, exhibiting state-of-the-art performance across all types of evaluation (sentence-level, system-level, and error span detection). Moreover, it does so while highlighting and categorizing error spans, thus enriching the quality assessment. We also provide a robustness analysis with stress tests, and show that xCOMET is largely capable of identifying localized critical errors and hallucinations.

1 Introduction

Automatic metrics for machine translation evaluation are widely used by researchers and practitioners to evaluate the quality of translations and the systems generating them. Notably, *learned* neural metrics, such as COMET (Rei et al., 2020) and BLEURT (Sellam et al., 2020), have demonstrated significant improvements in terms of correlation with human judgements when compared to traditional metrics like BLEU (Papineni et al., 2002; Freitag et al., 2021b, 2022).

These metrics are trained to regress on scores obtained through human annotations, by predicting a single sentence-level score representing the quality of the translation hypothesis. However,

these single scores do not offer a detailed view into translation errors (e.g., it is not immediate which words or spans of words are wrongly translated). Moreover, as they are obtained by making use of highly complex pre-trained models, they can be difficult to interpret (Rei et al., 2023b; Leiter et al., 2023). One appealing strategy to bring a more detailed view into translation errors is to obtain finer-grained information on error spans through highlighting them and indicating their severity (Fonseca et al., 2019; Perrella et al., 2022; Bao et al., 2023). In fact, this is the strategy adopted in recent works that have employed generative large language models (LLMs) for machine translation evaluation: (i) identify errors within a given translation, subsequently (ii) categorize these errors according to their severity, and finally (iii) infer a sentence-level score from the predicted errors (Fernandes et al., 2023; Xu et al., 2023). However, these methods still lag behind dedicated learned metrics when using open LLMs, such as the LLaMA models (Touvron et al., 2023; Xu et al., 2023). As it stands, competitive performance with generative strategies remains contingent on utilizing large *proprietary, closed* LLMs such as PaLM-2 and GPT-4 (Fernandes et al., 2023).

In this work, we bridge the gap between these two approaches to machine translation evaluation by introducing xCOMET: a *learned* metric that simultaneously performs sentence-level evaluation and error span detection. Through extensive experiments, we show that our metrics leverage the strengths of both paradigms: they achieve state-of-the-art performance in all relevant vectors of evaluation (sentence-level, system-level, and error span prediction), while offering, via the predicted error spans, a lens through which we can analyze translation errors and better interpret the sentence-level scores. We achieve this by employing a curriculum during training that is focused on leveraging high-quality *publicly available* data at both the sentence-

^{*}Equal contribution. Corresponding authors:

✉ {nuno.guerreiro,ricardo.rei}@unbabel.com

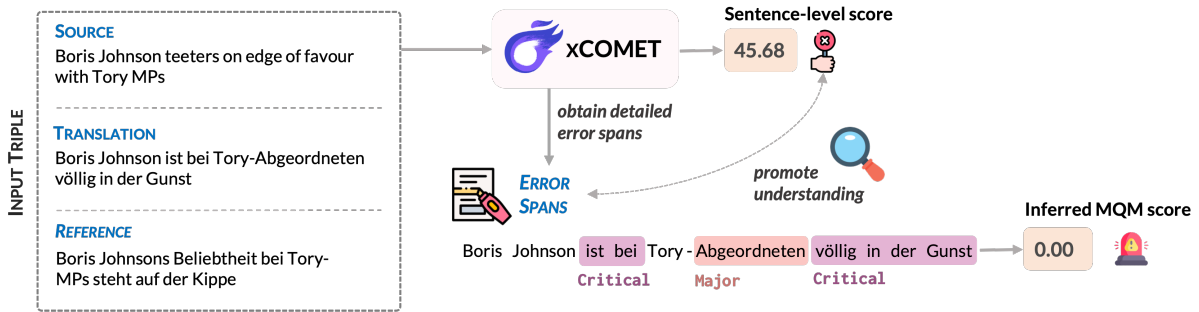


Figure 1: The **xCOMET** framework illustrated through a real example: the metric not only provides a sentence-level score, but also predicts translation error spans along with their respective severity. From these spans, we can infer MQM score (following the MQM typology) that informs and highly correlates with the sentence-level score (see Section 6). These spans complement the sentence-level score by providing a detailed view into the translation errors.

and error span level, complemented by the construction of synthetic data to enhance the metric’s robustness. Moreover, **xCOMET** is a unified metric (Wan et al., 2022b), accommodating all modes of evaluation within a single model. This enables the metric to be used even for quality estimation (when no reference is available), or for reference-only evaluation, similarly to BLEURT (when a source is not provided). Crucially, **xCOMET** also provides high-quality sentence-level scores that are directly inferred from the predicted error spans, in the style of AUTOMQM (Fernandes et al., 2023) and INSTRUCTSCORE (Xu et al., 2023).

Our contributions can be summarized as follows:

1. We introduce **xCOMET**, a novel evaluation metric that leverages the advantages of regression-based metrics and error span detection to offer a more detailed view of translation errors.
2. We show that **xCOMET** is a state-of-the-art metric at all relevant vectors of evaluation — sentence-level, system-level, and error span prediction — generally outperforming widely-used neural metrics and generative LLM-based machine translation evaluation.
3. We provide a comprehensive robustness analysis of **xCOMET**, showing that this new suite of metrics identifies the vast majority of localized critical errors and hallucinations.
4. We release two evaluation models: **xCOMET-XL**, with 3.5B parameters, and **xCOMET-XXL**, featuring 10.7B parameters.¹

¹The full suite of metrics (**xCOMET-XL** and **xCOMET-XXL**) will be released through HuggingFace Hub: <https://huggingface.co/Unbabel>.

2 Background

Methodologies for human assessment of translation quality. Human evaluation of machine translation is primarily conducted through three distinct approaches: post-edits (PE), direct assessments (DA), and the Multidimensional Quality Metrics (MQM) framework.

In post-edits (PE), *professional translators* are tasked with “fixing” a given translation, making minimal edits to improve its quality. Using this edited translation — often termed *post-edit* — we can evaluate the machine translation output by quantifying the number of edits, thus gauging the initial translation’s quality (Snover et al., 2006).

Direct assessments (DA) (Graham et al., 2013) are a simple and widely-used evaluation method. Annotators — *non-expert bilingual speakers* or *professional translators* — are asked to annotate each translation with a score ranging from 0 to 100 to reflect its adequacy and fluency, where a score of 100 corresponds to a perfect translation, and 0 corresponds to a completely inadequate one.

The Multidimensional Quality Metrics (MQM) framework (Lommel et al., 2014), on the other hand, offers a more comprehensive and systematic approach to MT evaluation. *Professional translators* highlight errors—typically in the form of error spans— within translations, attributing them severity ratings (e.g., *minor*, *major*, or *critical*) and categorical labels (e.g., *fluency*, *accuracy*). Figure 1 illustrates one such annotation. MQM annotations have gained prominence in recent years due to their capacity to offer detailed insights into translation errors, facilitating more fine-grained and accurate comparisons between translation systems (Freitag et al., 2021a). As such, the field of Automatic

Evaluation of MT has increasingly favoured comparisons using MQM annotations over traditional DA and PE methodologies (Freitag et al., 2021b, 2022; Zerva et al., 2022).

Automatic metrics for translation evaluation.

Conventional automatic metrics for machine translation (MT) evaluation rely on *lexical*-based approaches, where the evaluation score is computed through statistics related to lexical overlap between a machine translation and a reference translation. Despite evidence indicating that these lexical metrics (e.g., BLEU (Papineni et al., 2002) and CHRF (Popović, 2015)) do not consistently align with human judgments, particularly when these are obtained through the MQM framework (Freitag et al., 2021b, 2022), they remain very popular. In fact, BLEU remains the most widely employed evaluation metric in machine translation to this day (Marie et al., 2021). On the other hand, *neural* metrics (e.g., COMET (Rei et al., 2020) and BLEURT (Sellam et al., 2020)) that rely on complex neural networks to estimate the quality of MT outputs are consistently among the best metrics for MT evaluation according to correlations with human judgments (Freitag et al., 2021b, 2022).

However, contrary to lexical metrics which offer a straightforward interpretation, it can often prove challenging to explain the score predicted by a *neural* metric to a given translation output. As such, there have been a series of efforts to bring interpretability to neural metrics by focusing on understanding the inner workings of neural metrics (Rei et al., 2023b; Leiter et al., 2023), or on constructing inherently interpretable neural metrics (e.g., MATESE (Perrella et al., 2022) and FG-TED (Bao et al., 2023)) by assigning a central role to the task of predicting *word-level* errors in a given translation, instead of *just* a sentence-level score.

More recently, with the rise of generative LLMs, some works have tried to frame the MT evaluation problem as a generative problem. This offers great flexibility, as the LLM can be prompted to either score the translation directly (Kocmi and Federmann, 2023), or to identify errors in the translation (e.g., in line with the MQM framework) (Fernandes et al., 2023; Xu et al., 2023).

3 Problem Statement

An automatic *metric* for translation evaluation aims at predicting the quality of a translated sentence, t , in light of a reference translation, r , for a given

source sentence, s . Here, we focus specifically on neural metrics that make use of a neural model, and typically operate under one of the following evaluation scenarios:

- **reference-only** (REF): the model evaluates the translation by processing it alongside a ground-truth reference sentence (BLEURT is an example of such a metric);
- **source-reference combined input** (SRC+REF): the model evaluates the translation by jointly processing it with both the source and the reference (COMET is an example of such a metric);
- **source-only** (SRC): the model evaluates the translation using only its corresponding source sequence (COMETKIWI (Rei et al., 2022b) is an example of such a model). This mode is commonly termed as *quality estimation* (QE) or *reference-free* evaluation (Specia et al., 2010).

In essence, the model’s input sequence consists of the translation t paired with some **additional input**—either r , $[r, s]$ or s —derived from the scenarios above. Given this input, the model may predict the quality of the translation at different granularities, e.g., sentence-level or word(span)-level.

Sentence-level prediction. The model is tasked to predict a single global score—typically between 0 and 1—for the translation, that represents how well it aligns with its context (i.e., source and/or reference sentence). These scores can be used for a broad range of tasks, such as gauging the quality of different translation systems (Freitag et al., 2022), identifying pathological translations (Guerreiro et al., 2023b), assisting the generation of translations by MT systems (Fernandes et al., 2022), or even acting as reward models for human alignment of language models (Gulcehre et al., 2023).

Word(span)-level prediction. In contrast, word-level (or span-level) predictions are more fine-grained, identifying individual words or phrases in the translation that may have errors or discrepancies—typically identifying them as OK/BAD or according to their severity, e.g., MINOR/MAJOR. These granular evaluations are more interpretable and assist in pinpointing specific issues, which can be particularly valuable for feedback and iterative translation improvements.

Our metric, xCOMET, emerges in a unique position in the landscape of MT evaluation metrics.

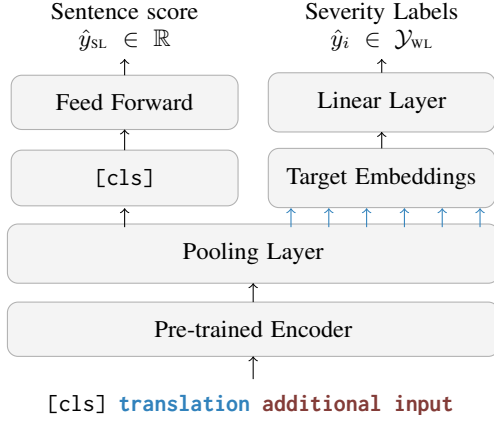


Figure 2: Architecture of xCOMET. The input to the model starts with a [cls] token followed by a **translation** and an **additional input** that will have the source, reference or both. After the pooling layer the [cls] token is passed to a feed-forward to produce a quality score while all subword pieces corresponding to the **translation** are passed to a linear layer that will classify them according to their severity levels, $\mathcal{Y}_{\text{WL}} = \{\text{OK}, \text{MIN}, \text{MAJ}, \text{CRIT}\}$.

It can simultaneously perform evaluation under all of the three scenarios (SRC, REF, SRC+REF) presented, and provide sentence-level scores and error span annotations that are in line with the MQM framework, thus bringing further transparency to the evaluation (see Figure 1 for an illustration). In the next section, we detail the design choices and methodology of xCOMET.

4 Design and Methodology of xCOMET

In this section, we describe the methodology behind xCOMET, outlining its model architecture, training settings and corpora, and learning curriculum. We detail how the model is designed to perform both regression and error span detection while adopting a unified input approach for enhanced flexibility and performance.

4.1 Model Architecture

xCOMET is built upon insights garnered from Unbabel-IST’s contributions to the WMT22 Metrics and QE shared tasks (Rei et al., 2022a,b). It is designed to concurrently handle two tasks: sentence-level regression and error span detection. Figure 2 illustrates its architecture. We follow the same architecture of the scaled-up version of COMETKIWI detailed in Rei et al. (2023a), which uses a large pre-trained encoder model as its backbone encoder model. Importantly, following

naturally from our multi-task setup, the model has two prediction heads: (i) a sentence-level *regression* head, which employs a feed-forward network to generate a sentence score, and (ii) a word-level *sequence tagger*, which applies a linear layer to assign labels to each token in a given translation.

We train two xCOMET versions — xCOMET-XL and xCOMET-XXL — using the XL (3.5B parameters) and XXL (10.7B parameters) versions of XLM-R (Goyal et al., 2021).²

4.2 Fully Unified Evaluation

xCOMET adopts a unified input approach (Wan et al., 2022b), allowing for all the evaluation scenarios described in Section 3—REF, SRC+REF, and SRC evaluation—under a single model. Thus, the input sequence consists of two parts: (i) the translated sentence $\mathbf{t} = [t_1, \dots, t_n]$ of length n , and (ii) an additional input containing information from the source, reference, or both.³ To do so, when a reference is available, we run three distinct forward passes (one for each evaluation scenario), each yielding sentence-level and word-level predictions.

4.2.1 Training time

For each forward-pass, we collect the sentence-level predictions $\{\hat{y}_{\text{SL}}^{\text{SRC}}, \hat{y}_{\text{SL}}^{\text{REF}}, \hat{y}_{\text{SL}}^{\text{SRC+REF}}\}$ and the word-level logits $\{\hat{\mathbf{y}}_{\text{WL}}^{\text{SRC}}, \hat{\mathbf{y}}_{\text{WL}}^{\text{REF}}, \hat{\mathbf{y}}_{\text{WL}}^{\text{SRC+REF}}\}$.⁴

As we have mentioned before, xCOMET models are trained with supervision from both sentence-level quality assessments, y_{SL} , and word-level severity tags, $\mathbf{y}_{\text{WL}} = [y_1, \dots, y_n]$, with $y_i \in \mathcal{Y}_{\text{WL}} = \{\text{OK}, \text{MIN}, \text{MAJ}, \text{CRIT}\}$. In the multi-task setting, we use the following loss \mathcal{L} for each input type ($\text{INPUT} \in \{\text{SRC}, \text{REF}, \text{SRC+REF}\}$):

$$\mathcal{L}_{\text{SL}}^{\text{INPUT}} = (y_{\text{SL}} - \hat{y}_{\text{SL}}^{\text{INPUT}})^2 \quad (1)$$

$$\mathcal{L}_{\text{WL}}^{\text{INPUT}} = -\frac{1}{n} \sum_{i=1}^n \alpha_{y_i} \log p(\hat{y}_i^{\text{INPUT}}) \quad (2)$$

$$\mathcal{L}^{\text{INPUT}} = (1 - \lambda) \mathcal{L}_{\text{SL}}^{\text{INPUT}} + \lambda \mathcal{L}_{\text{WL}}^{\text{INPUT}}, \quad (3)$$

$\alpha \in \mathbb{R}^{|\mathcal{Y}_{\text{WL}}|}$ represents the class weights given for each severity label and λ is used to weigh the combination of the sentence and word-level losses.

²To the best of our knowledge, these represent the two largest open-source encoder-only models.

³Each input is always delimited by separators. For example, the unified input can be written as: `<s>translation</s></s>src</s></s>ref</s>`

⁴Here, for each $\text{INPUT} \in \{\text{SRC}, \text{REF}, \text{SRC+REF}\}$, we define $\hat{\mathbf{y}}_{\text{WL}}^{\text{INPUT}} = [\hat{y}_1^{\text{INPUT}}, \dots, \hat{y}_n^{\text{INPUT}}]$.

The final learning objective is the summation of the losses for each input type:

$$\mathcal{L} = \mathcal{L}^{\text{SRC}} + \mathcal{L}^{\text{REF}} + \mathcal{L}^{\text{SRC+REF}} \quad (4)$$

Furthermore, in line with preceding metrics constructed upon the COMET framework, our models use features such as gradual unfreezing, and discriminative learning rates. See Appendix B for full details and hyperparameters.

4.2.2 Inference time

Error span prediction. For each subword in the translation, we average the output distribution of the word-level linear layer obtained for each forward pass. Using this distribution, we predict a set of word-level tags $\hat{\mathbf{y}}_{\text{WL}} = [\hat{y}_1, \dots, \hat{y}_n]$. From these tags, we construct a list of *error spans*, S , by grouping adjacent subwords identified as errors. The severity of each span in S is defined according to the most severe error tag found within the span.

Sentence-level prediction. For each forward pass, we obtain the corresponding sentence-level scores: \hat{y}_{SRC} , \hat{y}_{REF} , and $\hat{y}_{\text{SRC+REF}}$. Additionally, we leverage the information coming from the predicted list of error spans, S , to infer an automated MQM score. To do so, we follow the MQM framework: we obtain the error counts for each severity level— c_{MIN} , c_{MAJ} , c_{CRIT} —and apply the predetermined severity penalty multipliers to define the error type penalty total, $e(S)$. Formally:

$$e(S) = c_{\text{MIN}} + 5 \times c_{\text{MAJ}} + 10 \times c_{\text{CRIT}}. \quad (5)$$

Finally, we obtain \hat{y}_{MQM} by capping and flipping the sign of $e(S)$:

$$\hat{y}_{\text{MQM}} = \begin{cases} \frac{25-e(S)}{25}, & \text{if } e(S) < 25. \\ 0, & \text{otherwise.} \end{cases} \quad (6)$$

Note that the predicted score \hat{y}_{MQM} is bounded between 0 and 1, with a score of 1 corresponding to a perfect translation.

We aggregate the scores to compute the final sentence-level score, \hat{y}_{SL} , through a weighted sum of the different sentence-level scores. Importantly, we also include the inferred MQM score \hat{y}_{MQM} to directly inform the final sentence-level prediction. Formally, given $\hat{\mathbf{y}} = [\hat{y}_{\text{SRC}}, \hat{y}_{\text{REF}}, \hat{y}_{\text{SRC+REF}}, \hat{y}_{\text{MQM}}]$:

$$\hat{y}_{\text{SL}} = \mathbf{w}^\top \hat{\mathbf{y}} \quad (7)$$

where \mathbf{w} is set to $[1/9, 1/3, 1/3, 2/9]$.⁵

4.3 Corpora

Our models are exclusively trained on publicly available DA and MQM annotations, most of which have been collected by WMT over the recent years.

DA data. We use DA annotations collected by WMT from 2017 to 2020, and the MLQE-PE dataset (Fomicheva et al., 2022). As the MLQE-PE dataset does not contain reference translations, we used the post-edit translations as reference translations. Overall, the corpus consists of around 1 million samples, spanning 36 language pairs.

MQM data. We collected the MQM annotations sourced from WMT from 2020 to 2022.⁶ We also used annotations sourced from other MQM-annotated datasets: (i) IndicMT (Sai B et al., 2023), which contains MQM annotations spanning 5 Indian languages, and (ii) DEMETR (Karpinska et al., 2022), a diagnostic dataset with perturbations spanning semantic, syntactic, and morphological error categories.

Corpora with MQM annotations are usually extremely unbalanced with critical errors being underrepresented. In term, this may lead to metrics dealing less well with pathological translations, such as critical errors and hallucinations (Amrhein and Sennrich, 2022; Raunak et al., 2022; Guerreiro et al., 2023b). As such, we augment the MQM corpus with *synthetic critical errors*. We create different types of detached and oscillatory hallucinations (Raunak et al., 2021; Guerreiro et al., 2023b): (i) detached hallucinations, where we replace the translation with a random sentence; (ii) other detached hallucinations, where we replace the true translation with an unrelated translation that is semantically similar to the source sentence⁷; and (iii) oscillatory hallucinations, where we randomly sample a n -gram from the translation (with n in $\{2, 3, 4\}$) and repeat it between 1 and 10 times. We provide examples of these synthetic hallucinations in Appendix A. Overall, our MQM corpus consists of 176K samples, spanning 14 language pairs.

⁵UNITE uniformly distributes the weight across the different sentence-level scores to obtain the final prediction. However, we found that, in practice, distributing the weight of each sentence-level prediction can lead to improved results.

⁶Here, we exclude the 2022 News domain annotations, which we reserved for testing.

⁷We measure cross-lingual similarity via the cosine similarity between the sentence embeddings obtained with the LaBSE encoder (Feng et al., 2022).

Scaling of sentence-level scores. While the sentence-level scores inferred from MQM annotations (through the procedure in Equation 6) are bounded between 0 and 1, DA annotations usually require z -normalization in order to mitigate variations in scoring strategies by different annotators (Bojar et al., 2017).⁸ Thus, as z -scores are inherently centered at 0 and unbounded, there is a scaling mismatch between the data samples.

Consequently, to circumvent this limitation, we employ min-max scaling on our DA corpus to set its range of scores to $[0, 1]$. To do so, we set a practical minimum and maximum z -score value. We obtain the minimum score by averaging the z -scores for translations with over 1 annotation, wherein all annotators unanimously scored them with an unnormalized 0 DA score, i.e., they deemed the translation as “random”. For determining a maximum value, we applied the same process for perfect translations, i.e., unnormalized 100 DA score.⁹

4.4 Training Curriculum

xCOMET models undergo a 3-phase curriculum training. Throughout these phases, the training emphasis alternates between sentence-level prediction and error span prediction by tweaking the parameter λ in Equation 3. The curriculum phases can be described as follows:

Phase I: The model is trained exclusively using the DA data. In this phase, the focus is exclusively set on sentence-level regression.

Phase II: In this stage, we introduce word-level supervision. To achieve this, the model is fine-tuned on our diverse MQM corpus, with most emphasis placed on the word-level task.

Phase III: The last training phase is aimed at unifying both tasks. The model is further fine-tuned using high-quality MQM data from (Freitag et al., 2021a), with a bigger emphasis set to sentence-level prediction.

We describe how we obtain the values of λ for Phases II and III in Appendix B.¹⁰

⁸This is particularly relevant for DA annotations, since these judgements typically come from non-expert annotators.

⁹This technique was initially introduced in BLEURT-20 (Pu et al., 2021).

¹⁰The achieved λ weights for Phases II and III were $\lambda = 0.983$ and $\lambda = 0.055$, respectively.

Interpretation of the curriculum. We first start by training a sentence-level metric — similar to UNITE (Wan et al., 2022a) — on the vastly available DA annotations. This first phase acts as a warm-up for subsequent stages. In fact, prior research has shown that models trained on DA annotations leverage token-level information that aligns with MQM error annotations (Rei et al., 2023b). When we move to the second phase, we assume that we have a metric that can perform sentence-level regression. Thus, the aim here shifts to integrating word-level supervision without compromising the previously acquired sentence-level prediction skills. To do so, we use the highly diverse corpora of MQM annotations and set most emphasis on the word-level task. Finally, we exclusively leverage a small corpus (around 25k samples) of very high-quality MQM annotations from (Freitag et al., 2021a) — each sample has three annotations from separate annotators — with additional synthetic hallucinations. Our focus here is to mitigate any potential decline in sentence-level regression capabilities during Phase II.

5 Experimental Setting

5.1 Evaluation

We test our metrics on the MQM annotations from the News domain from the WMT 2022 Metrics shared task. These annotations cover three language pairs: Chinese→English (zh-en), English→German (en-de), and English→Russian (en-ru).¹¹ We evaluate the metrics in terms of sentence-level, system-level, and error span prediction performance.

At the sentence-level, we report both the Pearson correlation coefficient (ρ) and Kendall’s Tau (τ) using the Perm-Both hypothesis test (Deutsch et al., 2021). We also evaluate the metrics on System-level Pairwise Accuracy (Kocmi et al., 2021). We base these evaluations on 200 re-sampling runs, with a significance level (p) set to 0.05. For error span prediction, we adopt the WMT23 Quality Estimation shared task evaluation methodology and compute F1 scores calculated at the character level, taking into account partial matches for both minor and major errors.¹²

¹¹The test set comprises 4,500 segments for en-de, 4,500 for en-ru, and 7,575 for zh-en, sourced from 15 different translation systems.

¹²We convert all critical errors into major errors, in order to match the guidelines (focused exclusively on minor and major errors) described in (Freitag et al., 2021a), that were used for

METRIC	zh-en		en-de		en-ru		Avg.	
	ρ	τ	ρ	τ	ρ	τ	ρ	τ
BLEURT-20	0.462	0.336	0.568	0.380	0.498	0.379	0.509	0.365
COMET-22	0.423	0.335	0.581	0.369	0.516	0.391	0.507	0.361
METRICX	0.573	0.415	0.640	0.405	0.581	0.444	0.598	0.421
GEMBA-GPT4-DA*	0.318	0.292	0.508	0.387	0.454	0.383	0.427	0.354
xCOMET-XL	0.556	0.399	0.653	0.414	0.611	0.448	0.607	0.420
xCOMET-XXL	0.554	0.390	0.644	0.435	0.628	0.470	0.609	0.432
<i>Predicted MQM scores from the error spans ($\hat{y} = \hat{y}_{MQM}$)</i>								
xCOMET-XL (MQM)	0.447	0.374	0.561	0.389	0.534	0.445	0.514	0.402
xCOMET-XXL (MQM)	0.446	0.332	0.597	0.415	0.533	0.439	0.525	0.395

Table 1: Segment-level Pearson (ρ) and Kendall-Tau (τ) (\uparrow) using the Perm-Both hypothesis test (Deutsch et al., 2021). Numbers in bold belong to the top-performing cluster according to statistical significance ($p < 0.05$).

5.2 Baselines

Sentence and system-level. We benchmark our metrics widely used *open* neural metrics: COMET-22 (Rei et al., 2022a)¹³ and BLEURT-20 (Pu et al., 2021). Additionally, we include METRICX, the best performing metric from WMT22 Metrics shared task (Freitag et al., 2022).¹⁴ Finally, we also include GEMBA (Kocmi and Federmann, 2023), which employs GPT4 (OpenAI, 2023) to evaluate translations following DA guidelines.

Error span prediction. We report results using GPT3.5 and GPT4 models, by prompting it in the style of AUTOMQM (Fernandes et al., 2023).¹⁵ We carefully select 5 shots that are held constant for all samples. This way, we can directly compare our results with state-of-the-art LLMs, which have been shown to be able to perform the task of error detection (Fernandes et al., 2023; Xu et al., 2023).

6 Correlations with Human Judgements

In this section, we present a standard performance analysis of our metrics in terms of correlations with human judgments. Overall, we find xCOMET to be a state-of-the-art in sentence-level and error span prediction, being competitive with generative LLMs in terms of system-level evaluation.

Sentence-level evaluation. Table 1 shows that both xCOMET metrics outperform other strong performing neural metrics, including the generative

annotating the zh-en and de-en test sets.

¹³Used checkpoint: Unbabel/wmt22-comet-da

¹⁴Specifically, we employ the metricx_xxl_mqm_2020 submission scores from the mt-metrics-eval package. Although the metric has not been released publicly, it is public that it is built upon the mT5-XXL (Xue et al., 2021) and has 13B parameters (Deutsch et al., 2023).

¹⁵We use the models from the OpenAI API (gpt-3.5-turbo and gpt-4) in October, 2023.

METRIC	zh-en	en-de	en-ru	Avg.
BLEURT-20	0.762	0.771	0.743	0.759
COMET-22	0.705	0.800	0.733	0.746
METRICX	0.762	0.781	0.724	0.756
GEMBA-GPT4-DA	0.752	0.848	0.876	0.825
xCOMET-XL	0.800	0.743	0.790	0.778
xCOMET-XXL	0.800	0.829	0.829	0.819
<i>MQM scores from the error spans ($\hat{y} = \hat{y}_{MQM}$)</i>				
xCOMET-XL (MQM)	0.781	0.762	0.762	0.768
xCOMET-XXL (MQM)	0.781	0.838	0.810	0.810

Table 2: System-level Pairwise Accuracy (\uparrow) (Kocmi et al., 2021) using the Perm-Both hypothesis test (Deutsch et al., 2021). Numbers in bold belong to the top-performing cluster according to statistical significance ($p < 0.05$).

approach leveraging GPT4 of GEMBA. In particular, xCOMET-XXL sets a new state-of-the-art for en-de and en-ru. Interestingly, we can see that, while scaling up the encoder model of the xCOMET metrics (from XL to XXL) holds better results, xCOMET-XL is very competitive.¹⁶ In fact, it outperforms METRICX, which runs at even a larger size than xCOMET-XXL. Finally, we can also observe that the MQM scores inferred exclusively from the predicted error spans also exhibit strong performance, outperforming widely used metrics BLEURT-20 and COMET-22. This is particularly relevant: the predicted error spans bring not only a more detailed view into translation errors but also provide high-quality sentence-level scores.

System-level evaluation. Table 2 shows results for system-level. Similarly to what we observed at the sentence-level, our metrics show consistently superior performance when compared to other dedicated neural metrics. Notably, although genera-

¹⁶Note that this corresponds to an increase in the parameter size of over 7 billion parameters.

METRIC	zh-en	en-de	en-ru	Avg.
• AutoMQM (GPT3.5)	0.143	0.160	0.166	0.156
• AutoMQM (GPT4)	0.248	0.257	0.281	0.262
• xCOMET-XL	0.237	0.290	0.281	0.269
• xCOMET-XXL	0.257	0.320	0.262	0.280
<i>Error spans detected with source-only</i>				
• xCOMET-XL (SRC)	0.208	0.264	0.252	0.242
• xCOMET-XXL (SRC)	0.229	0.298	0.238	0.255

Table 3: F1 scores (\uparrow) on error span detection for reference-free (•) and reference-based (•) evaluation.

SCORE	zh-en	en-de	en-ru	All
\hat{y}_{SRC}	0.73	0.75	0.79	0.78
\hat{y}_{REF}	0.75	0.74	0.75	0.77
$\hat{y}_{\text{SRC+REF}}$	0.78	0.79	0.82	0.82
$\hat{y}_{\text{SL}}^{\dagger}$	0.90	0.92	0.92	0.92

Table 4: Pearson correlations between the regression scores produced by xCOMET-XXL (\hat{y}_{SRC} , \hat{y}_{REF} , $\hat{y}_{\text{SRC+REF}}$, \hat{y}_{SL}) and the MQM inferred score, \hat{y}_{MQM} , computed from the identified error spans. \dagger The computation of \hat{y}_{SL} , contrary to the computation of the other regression scores, makes direct use of \hat{y}_{MQM} (see Equation 7).

tive approaches typically do much better at system-level evaluation when compared to dedicated models (Kocmi and Federmann, 2023; Fernandes et al., 2023), xCOMET-XXL remains competitive in all language pairs with GEMBA using GPT4. Finally, building on the findings at the sentence-level, the MQM scores inferred directly and exclusively from the predicted error spans also exhibit very competitive performance in terms of system-level accuracy.

Error span prediction. While we have highlighted the utility of the predicted error spans through the inferred sentence-level MQM scores, here we turn to evaluating them directly. Table 3 shows that the error spans predicted via xCOMET metrics outperform those obtained with both GPT3.5 and GPT4 despite being smaller in capacity relative to these models. In fact, our metrics achieve close performance to that of GPT4, even when a reference is not provided.

Interplay of error spans and sentence-level scores. Table 4 shows a strong correlation between the different score types predicted by xCOMET and the MQM inferred score derived exclusively from error spans. This interplay is highly important: the predicted error spans may be valuable, not just for the sake of accuracy but also for interpretability. Interestingly, these high correla-

tions with the predicted scores from each forward pass (\hat{y}_{SRC} , \hat{y}_{REF} , $\hat{y}_{\text{SRC+REF}}$) are obtained despite no explicit alignment mechanism governing the relationship between the predictions of the sentence-level and word-level heads. We hypothesize that it is thus the shared encoder that, during the multi-task training, aligns the representations between the two tasks. As such, xCOMET provides, through its predicted error spans, a potential lens through which we can better understand, contextualize, and even debug its own sentence-level predictions.

7 Robustness of xCOMET to pathological translations

In the previous section, we have shown that xCOMET metrics exhibit state-of-the-art correlations with human judgements when evaluating on high-quality MQM annotations. However, more often than not, these MQM annotations are highly unbalanced and contain little to no major or critical errors. As such, they may not offer a full picture of the metrics’ performance. In this section, we shift our focus to studying how xCOMET metrics behave when evaluating translations with localized major or critical errors, such as named-entity errors or mismatches in numbers, and highly pathological translations, such as hallucinations.

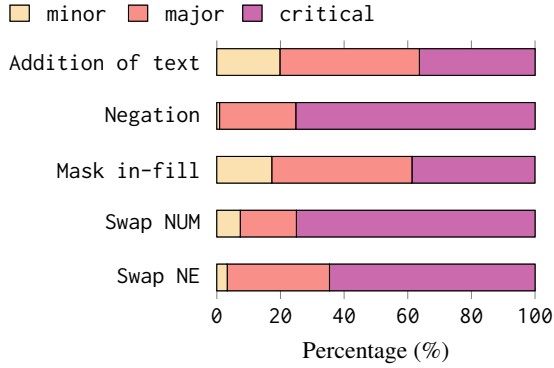
7.1 Localized errors

We employ SMAUG (Alves et al., 2022)¹⁷, a tool designed to generate synthetic data for stress-testing metrics, to create corrupted translations that contain major or critical errors. Concretely, we generate translations with the following pathologies: addition of text, negation errors, mask in-filling, named entity errors, and errors in numbers. For this evaluation, we use data from the WMT 2023 Metrics shared task.¹⁸ Specifically, we corrupt the released *synthetic* references for which the metrics found no errors.¹⁹ Moreover, as the full suite of SMAUG transformations can only be applied to English data, we focus solely on Chinese→English (zh-en) and Hebrew→English (he-en) translations. Full details about the corrupted data and examples are shown in Appendix C.

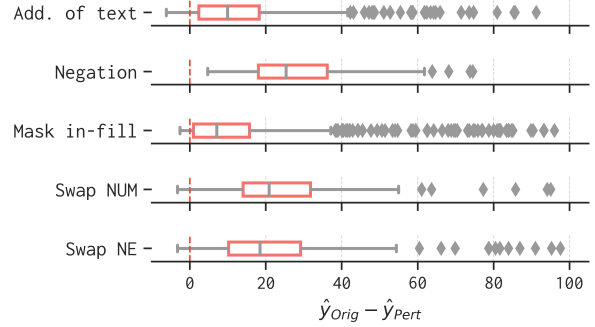
¹⁷<https://github.com/Unbabel/smaug>

¹⁸At the time of writing, the human MQM annotations for this data have not been released. Nevertheless, this is not prohibitive for conducting the analysis in this section.

¹⁹This allows us to isolate the effect of the perturbations. In case there are predicted error spans for the transformed translations, these are a result of the perturbation induced.



(a) Percent of error types on data with critical errors (for both zh-en and he-en data), as predicted by xCOMET-XXL.



(b) Impact of the perturbations, as measured by the difference in xCOMET-XXL ($\hat{y} = \hat{y}_{SL}$) between the original and the perturbed translation, on the zh-en data.

Figure 3: Analysis of xCOMET-XXL for data with localized critical errors in terms of (a) distribution of error severities for the predicted error spans, and (b) sensitivity of the sentence-level scores.

xCOMET predicts most localized errors as major or critical errors.

Table 5 shows that xCOMET metrics identify the vast majority of localized errors, with trends varying across scale and language pair. Generally, negation errors and mismatches in numbers are the most easily identified by the metrics. This is interesting: localized errors, such as mismatches in numbers and named-entity errors, had been pinpointed as weaknesses of previous COMET metrics (Amrhein and Sennrich, 2022; Raunak et al., 2022). This earlier limitation seems to now have been addressed successfully. In fact, the results in Figure 3a show that most of these errors are indeed predicted as critical errors. One plausible hypothesis for these improvements is the incorporation of datasets that contain several negative translations, such as DEMETR, MLQE-PE, and synthetic hallucinations into our training set.

xCOMET sentence-level scores are sensitive to localized perturbations.

Figure 3b²⁰ shows that localized errors can lead to significant decreases in the predicted sentence-level scores, with perturbation-wise trends mirroring those of the error span predictions: the most pronounced decreases are found for negation errors and mismatches in numbers and named-entities (median decreases of around 20 points). The distribution of the decreases in quality also reveals two relevant trends: (i) localized perturbations can cause xCOMET-XXL to shift from a score of a perfect translation to that of an unrelated translation, and (ii) the behavior of xCOMET-XXL is not perfect and

²⁰Results for he-en and xCOMET-XL can be found in Appendix C.

ERROR	zh-en		he-en	
	XL	XXL	XL	XXL
Add. of text	3.66	10.7	6.15	7.35
Negation	0.20	0.20	3.89	4.90
Mask in-fill	5.01	17.0	4.78	3.92
Swap NUM	3.19	2.88	0.16	0.00
Swap NE	3.66	6.94	9.81	7.01
All	2.24	10.7	9.81	7.00

Table 5: Percentage (%) of translations, segmented by perturbation type, that are predicted to have no errors (\downarrow). We show results for both zh-en and he-en language pairs across xCOMET (XL and XXL) sizes.

can be further improved: in some rare cases, perturbations may actually lead increase in the score. Nevertheless, upon closer inspection, we found that, for over 90% of these cases, the decrease is smaller than 1 point.

7.2 Hallucinations

Hallucinations lie at the extreme-end of machine translation pathologies (Raunak et al., 2021), and can have devastating impact when models are deployed *in the wild*. Yet, these translations are often overlooked when assessing the performance of different translation systems. Their rarity means that performance, usually judged according to an aggregated corpus-level score, may remain largely unperturbed by a very small number of hallucinations.²¹

²¹For example, hallucination rates with state-of-the-art translation systems on high-resource language pairs are extremely rare: Raunak et al. (2022) found only five hallucinations in over 100k translations when using Microsoft’s translation system by way of their paid public APIs; and, Guerreiro et al. (2023a) found zero hallucinations with a 12B M2M

METRIC	All	Full Det.	Osc.
• BLEURT-20	0.824	0.892	0.799
• COMET-22	0.829	0.878	0.883
• COMETKIWI-XXL	0.839	0.834	0.902
• xCOMET-XL	0.865	0.907	0.922
• xCOMET-XXL	0.890	0.964	0.844
<i>QE scores from the error spans ($\hat{y} = \hat{y}_{\text{SRC}}$)</i>			
• xCOMET-XL (SRC)	0.885	0.924	0.944
• xCOMET-XXL (SRC)	0.902	0.959	0.866

Table 6: Hallucination detection performance on the de-en hallucination benchmark from Guerreiro et al. (2023b) as measured by AUROC (\uparrow) for reference-free (\bullet) and reference-based (\bullet) quality metrics. We report results for all the dataset, for fully detached, and oscillatory hallucinations separately.

In this section, we want to assess how the xCOMET metrics rank hallucinations among other translations. To do so, we will use the German \rightarrow English hallucination benchmark introduced in Guerreiro et al. (2023b). This benchmark involves over 3.4k translations of different error types, including omissions, named-entity errors, and hallucinations (oscillatory, fully, and strongly detached). For a metric that has not been trained explicitly to rank translations, the benchmark is quite challenging: hallucinations should be ranked below other severe errors and incorrect translations. We provide examples of the hallucinations in the dataset in Appendix D.

xCOMET metrics can distinguish hallucinations from other translations. The results in Table 6 show that both xCOMET metrics largely rank hallucinations lower than other errors. This is especially true for the most severe type of hallucination (fully detached), for which the AUROC exceeds 95 for the XXL metric. In fact, Figure 4 reveals that xCOMET-XXL assigns over 90% of these fully detached hallucinations a score under 10. Relative to previous metrics, xCOMET achieves overall improvements. Interestingly, we also find that SRC-based evaluation (i.e., without the use of a reference translation) can reap benefits in this scenario. We hypothesize that this is due to the metric over-relying on the reference when it is made available (Rei et al., 2023b). While hallucinations contain content that is detached from the source, some of their text may still overlap (even if just lexically) with the reference text (e.g., in strongly

model (Fan et al., 2021) when tested on more than 35k translations. This trend, however, was not found for medium and low-resource translation directions, where hallucination rates can be well above 10% for some language pairs.

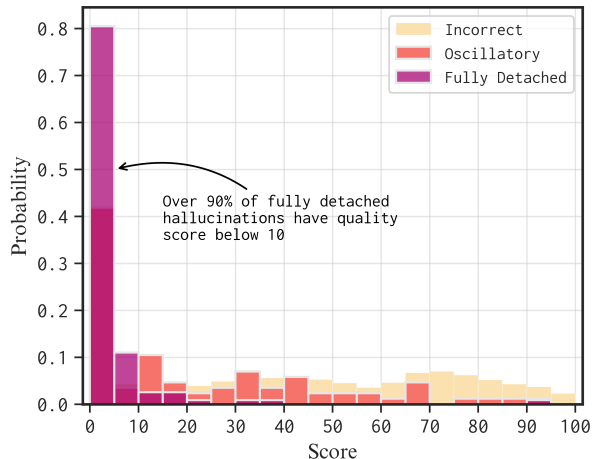


Figure 4: Category-wise distribution of xCOMET-XXL scores on the hallucination benchmark.

detached or oscillatory hallucinations), leading to higher scores. In future work, it would be interesting to explore whether these trends hold for other language pairs, including low-resource ones, through the use of multilingual hallucination benchmarks like HalOmi (Dale et al., 2023).²²

8 Conclusions

We introduced xCOMET: a novel suite of metrics for machine translation evaluation that effectively combines sentence-level prediction with fine-grained error span prediction. Through extensive experiments, we have shown that xCOMET is a state-of-the-art metric at all relevant vectors of evaluation: sentence-level, system-level, and error span prediction. Notably, through xCOMET’s capabilities to predict error spans, we can not only obtain useful signals for downstream prediction (either directly through error span prediction or by informing sentence-level scores) but also gain access to a lens through which we can better understand and interpret its predictions. Finally, we also stress-tested the suite of metrics by analyzing their behavior on scoring localized critical errors and hallucinations: xCOMET metrics identify the vast majority of localized errors and can appropriately penalize the severity of hallucinations.

We hope xCOMET can serve as a further step towards more detailed and informed machine translation evaluation. The full suite of metrics (xCOMET-XL and xCOMET-XXL) will be made available through the HuggingFace Hub: <https://huggingface.co/Unbabel>.

²²The HalOmi benchmark is yet to be publicly released.

Acknowledgements

We are grateful to José Pombal, José G. C. de Souza and Sweta Agrawal for their valuable feedback and discussions.

This work was supported by the Portuguese Recovery and Resilience Plan (PRR) through project C645008882-00000055, Center for Responsible AI, by the European Research Council (DECOLLAGE, ERC-2022-CoG 101088763), by EU’s Horizon Europe Research and Innovation Actions (UTTER, contract 101070631), and by the Fundação para a Ciência e Tecnologia (contracts UIDB/50021/2020 and UIDB/50008/2020). We also thank the HPC resources from GENCI-IDRIS (Grants 2023–AD011014714, 2023–AD0110146A68R1 and AD011012377R2).

References

- Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. 2019. [Optuna: A next-generation hyperparameter optimization framework](#). In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*.
- Duarte Alves, Ricardo Rei, Ana C Farinha, José G. C. de Souza, and André F. T. Martins. 2022. [Robust MT evaluation with sentence-level multilingual augmentation](#). In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 469–478, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- 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)*, pages 1125–1141, Online only. Association for Computational Linguistics.
- Keqin Bao, Yu Wan, Dayiheng Liu, Baosong Yang, Wenqiang Lei, Xiangnan He, Derek F. Wong, and Jun Xie. 2023. [Towards fine-grained information: Identifying the type and location of translation errors](#).
- Ondřej Bojar, Rajen Chatterjee, Christian Federmann, Yvette Graham, Barry Haddow, Shujian Huang, Matthias Huck, Philipp Koehn, Qun Liu, Varvara Logacheva, Christof Monz, Matteo Negri, Matt Post, Raphael Rubino, Lucia Specia, and Marco Turchi. 2017. [Findings of the 2017 conference on machine translation \(WMT17\)](#). In *Proceedings of the Second Conference on Machine Translation*, pages 169–214, Copenhagen, Denmark. Association for Computational Linguistics.
- David Dale, Elena Voita, Janice Lam, Prangthip Hansanti, Christophe Ropers, Elahe Kalbassi, Cynthia Gao, Loïc Barrault, and Marta R. Costa-jussà. 2023. [Halomi: A manually annotated benchmark for multilingual hallucination and omission detection in machine translation](#).
- Daniel Deutsch, Rotem Dror, and Dan Roth. 2021. [A statistical analysis of summarization evaluation metrics using resampling methods](#). *Transactions of the Association for Computational Linguistics*, 9:1132–1146.
- Daniel Deutsch, Juraj Juraska, Mara Finkelstein, and Markus Freitag. 2023. [Training and meta-evaluating machine translation evaluation metrics at the paragraph level](#).
- Angela Fan, Shruti Bhosale, Holger Schwenk, Zhiyi Ma, Ahmed El-Kishky, Siddharth Goyal, Mandeep Baines, Onur Celebi, Guillaume Wenzek, Vishrav Chaudhary, Naman Goyal, Tom Birch, Vitaliy Liptchinsky, Sergey Edunov, Edouard Grave, Michael Auli, and Armand Joulin. 2021. [Beyond english-centric multilingual machine translation](#). *J. Mach. Learn. Res.*, 22(1).
- Fangxiaoyu Feng, Yinfei Yang, Daniel Cer, Naveen Ariavzhagan, and Wei Wang. 2022. [Language-agnostic BERT sentence embedding](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 878–891, Dublin, Ireland. Association for Computational Linguistics.
- Patrick Fernandes, Daniel Deutsch, Mara Finkelstein, Parker Riley, André F. T. Martins, Graham Neubig, Ankush Garg, Jonathan H. Clark, Markus Freitag, and Orhan Firat. 2023. [The devil is in the errors: Leveraging large language models for fine-grained machine translation evaluation](#).
- 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*, pages 1396–1412, Seattle, United States. Association for Computational Linguistics.
- Marina Fomicheva, Shuo Sun, Erick Fonseca, Chrysoula Zerva, Frédéric Blain, Vishrav Chaudhary, Francisco Guzmán, Nina Lopatina, Lucia Specia, and André F. T. Martins. 2022. [MLQE-PE: A multilingual quality estimation and post-editing dataset](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 4963–4974, Marseille, France. European Language Resources Association.
- Erick Fonseca, Lisa Yankovskaya, André F. T. Martins, Mark Fishel, and Christian Federmann. 2019. [Findings of the WMT 2019 shared tasks on quality](#)

- estimation. In *Proceedings of the Fourth Conference on Machine Translation (Volume 3: Shared Task Papers, Day 2)*, pages 1–10, Florence, Italy. Association for Computational Linguistics.
- Markus Freitag, George Foster, David Grangier, Viresh Ratnakar, Qijun Tan, and Wolfgang Macherey. 2021a. [Experts, errors, and context: A large-scale study of human evaluation for machine translation](#). *Transactions of the Association for Computational Linguistics*, 9:1460–1474.
- Markus Freitag, Ricardo Rei, Nitika Mathur, Chi-kiu Lo, Craig Stewart, Eleftherios Avramidis, Tom Kocmi, George Foster, Alon Lavie, and André F. T. Martins. 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)*, pages 46–68, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Markus Freitag, Ricardo Rei, Nitika Mathur, Chi-kiu Lo, Craig Stewart, George Foster, Alon Lavie, and Ondřej Bojar. 2021b. [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*, pages 733–774, Online. Association for Computational Linguistics.
- Naman Goyal, Jingfei Du, Myle Ott, Giri Anantharaman, and Alexis Conneau. 2021. [Larger-scale transformers for multilingual masked language modeling](#). In *Proceedings of the 6th Workshop on Representation Learning for NLP (RepL4NLP-2021)*, pages 29–33, Online. Association for Computational Linguistics.
- 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*, pages 33–41, Sofia, Bulgaria. Association for Computational Linguistics.
- Nuno M. Guerreiro, Duarte Alves, Jonas Waldendorf, Barry Haddow, Alexandra Birch, Pierre Colombo, and André F. T. Martins. 2023a. [Hallucinations in large multilingual translation models](#).
- Nuno M. Guerreiro, Elena Voita, and André Martins. 2023b. [Looking for a needle in a haystack: A comprehensive study of hallucinations in neural machine translation](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 1059–1075, Dubrovnik, Croatia. Association for Computational Linguistics.
- Caglar Gulcehre, Tom Le Paine, Srivatsan Srinivasan, Ksenia Konyushkova, Lotte Weerts, Abhishek Sharma, Aditya Siddhant, Alex Ahern, Miaosen Wang, Chenjie Gu, Wolfgang Macherey, Arnaud Doucet, Orhan Firat, and Nando de Freitas. 2023. [Reinforced self-training \(rest\) for language modeling](#).
- Marzena Karpinska, Nishant Raj, Katherine Thai, Yixiao Song, Ankita Gupta, and Mohit Iyyer. 2022. [DEMETR: Diagnosing evaluation metrics for translation](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 9540–9561, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Tom Kocmi and Christian Federmann. 2023. [Large language models are state-of-the-art evaluators of translation quality](#). In *Proceedings of the 24th Annual Conference of the European Association for Machine Translation*, pages 193–203, Tampere, Finland. European Association for Machine Translation.
- 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*, pages 478–494, Online. Association for Computational Linguistics.
- Christoph Leiter, Piyawat Lertvittayakumjorn, M. Fomicheva, Wei Zhao, Yang Gao, and Steffen Eger. 2023. [Towards explainable evaluation metrics for machine translation](#). *ArXiv*, abs/2306.13041.
- Arle Lommel, Aljoscha Burchardt, and Hans Uszkorait. 2014. [Multidimensional quality metrics \(MQM\): A framework for declaring and describing translation quality metrics](#). *Tradumàtica: tecnologies de la traducció*, 0:455–463.
- Benjamin Marie, Atsushi Fujita, and Raphael Rubino. 2021. [Scientific credibility of machine translation research: A meta-evaluation of 769 papers](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 7297–7306, Online. Association for Computational Linguistics.
- OpenAI. 2023. [Gpt-4 technical report](#).
- 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*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Stefano Perrella, Lorenzo Proietti, Alessandro Scirè, Niccolò Campolungo, and Roberto Navigli. 2022. [MaTESe: Machine translation evaluation as a sequence tagging problem](#). In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 569–577, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.

- Maja Popović. 2015. [chrF: character n-gram F-score for automatic MT evaluation](#). In *Proceedings of the Tenth Workshop on Statistical Machine Translation*, pages 392–395, Lisbon, Portugal. Association for Computational Linguistics.
- Amy Pu, Hyung Won Chung, Ankur Parikh, Sebastian Gehrmann, and Thibault Sellam. 2021. [Learning compact metrics for MT](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 751–762, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Vikas Raunak, Arul Menezes, and Marcin Junczys-Dowmunt. 2021. [The curious case of hallucinations in neural machine translation](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1172–1183, Online. Association for Computational Linguistics.
- Vikas Raunak, Matt Post, and Arul Menezes. 2022. [SALTED: A framework for SALient long-tail translation error detection](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 5163–5179, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- 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)*, pages 578–585, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- 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](#).
- Ricardo Rei, Nuno M. Guerreiro, Marcos Treviso, Luisa Coheur, Alon Lavie, and André Martins. 2023b. [The inside story: Towards better understanding of machine translation neural evaluation metrics](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 1089–1105, Toronto, Canada. Association for Computational Linguistics.
- Ricardo Rei, Craig Stewart, Ana C Farinha, and Alon Lavie. 2020. [COMET: A neural framework for MT evaluation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2685–2702, Online. Association for Computational Linguistics.
- 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. 2022b. [CometKiwi: IST-unbabel 2022 submission for the quality estimation shared task](#). In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 634–645, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Ananya Sai B, Tanay Dixit, Vignesh Nagarajan, Anoop Kunchukuttan, Pratyush Kumar, Mitesh M. Khapra, and Raj Dabre. 2023. [IndicMT eval: A dataset to meta-evaluate machine translation metrics for Indian languages](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14210–14228, Toronto, Canada. Association for Computational Linguistics.
- Thibault Sellam, Dipanjan Das, and Ankur Parikh. 2020. [BLEURT: Learning robust metrics for text generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7881–7892, Online. Association for Computational Linguistics.
- Matthew Snover, Bonnie Dorr, Richard Schwartz, Linnea Micciulla, and John Makhoul. 2006. [A study of translation edit rate with targeted human annotation](#). In *Proceedings of Association for Machine Translation in the Americas*, pages 223–231.
- Lucia Specia, Dhvaj Raj, and Marco Turchi. 2010. Machine translation evaluation versus quality estimation. *Machine translation*, 24:39–50.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023. [Llama: Open and efficient foundation language models](#).
- Dazhen Wan, Zheng Zhang, Qi Zhu, Lizi Liao, and Minlie Huang. 2022a. [A unified dialogue user simulator for few-shot data augmentation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 3788–3799, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Yu Wan, Dayiheng Liu, Baosong Yang, Haibo Zhang, Boxing Chen, Derek Wong, and Lidia Chao. 2022b. [UniTE: Unified translation evaluation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8117–8127, Dublin, Ireland. Association for Computational Linguistics.
- Wenda Xu, Danqing Wang, Liangming Pan, Zhenqiao Song, Markus Freitag, William Yang Wang, and Lei Li. 2023. [Instructscore: Towards explainable text generation evaluation with automatic feedback](#).
- Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2021. [mT5: A massively multilingual pre-trained text-to-text transformer](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics:*

Human Language Technologies, pages 483–498, Online. Association for Computational Linguistics.

Chrysoula Zerva, Frédéric Blain, Ricardo Rei, Piyawat Lertvittayakumjorn, José G. C. de Souza, Steffen Eger, Diptesh Kanojia, Duarte Alves, Constantin Orăsan, Marina Fomicheva, André F. T. Martins, and Lucia Specia. 2022. [Findings of the WMT 2022 shared task on quality estimation](#). In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 69–99, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.

Supplemental Material

A Examples of synthetic hallucinations in the training data

We show in Table 7 one example for each type of synthetic hallucination that we use to augment xCOMET’s training data.

Source: Touristen in Portugal in Panik versetzt, nachdem ein tieffliegender Militärjet Strand überfliegt
Reference: "Tourists in Portugal are left terrified as a low-flying military jet flies skims beach"
Detached Hallucination (random): And best with a deal. I am cautiously optimistic that this will work.
Source: Komet entdeckt: Interstellarer Gast kreuzt durch unser Sonnensystem
Reference: Comet discovered: An interstellar guest crosses through our solar system
Detached Hallucination (semantically similar): Comet crossed by other star solar system
Source: Wie ist jetzt die Situation auf der Insel?
Reference: What is the situation on the island now?
Oscillatory Hallucination: What is the situation on the island the island the island the island the island the island the island now?

Table 7: Examples of synthetically-generated hallucinations for Phases II and III of xCOMET’s training.

B Hyperparameters

B.1 Setting the parameter λ for curriculum training

To obtain λ for each Phase, we do hyperparameter tuning with Optuna (Akiba et al., 2019), running over 20 trials changing λ in adequate intervals depending on the objective of the phase (e.g., values closer to 1 for Phase II, and values closer to 0 for Phase III).

B.2 Hyperparameters used for each phase

You can find the hyperparameters used to train models from Phase I to Phase III in Listings 1 and 2.²³ Note that the only difference between the two phases, apart from the training data (which we describe in 4.3), is the value of the λ parameter. Regarding the class weights α , we have also optimized this parameter using Optuna (Akiba et al., 2019), and keep them fixed throughout Phases II and III. As expected, the *optimal* weights reflect the class-inbalance, assigning the smallest weight to the OK tag and the largest to the MAJ and CRIT tags.

C SMAUG data

C.1 Data construction

In Section 7.1 we analysed how xCOMET behaves when the translation hypothesis contains localized major/critical errors using synthetic data created with SMAUG (Alves et al., 2022). The analysed pathologies are: hallucinations via addition of text, negation errors, hallucinations via mask in-fill, swap of numbers and swap of named entities.²⁴ Table 8 presents a summary of the examples created using

²³We omit the config arguments that are respective to the encoder model, such as `word_layer` and `hidden_sizes`.

²⁴On the hallucinations via addition of text and via mask in-fill, SMAUG encourages to add tokens or replace masked tokens with other tokens that preserve fluency. As such, there may be some perturbed translations that are close to paraphrases of the original translation. We hypothesize this is the reason why there is a bigger percentage of minor errors when compared to other error types (e.g., negation errors).

Listing 1: Relevant hyperparameters used for training Phase - I models (Section 4.1) using COMET framework.

```
class_path: UnifiedMetric
init_args:
nr_frozen_epochs: 0.3
keep_embeddings_frozen: True
optimizer: AdamW
encoder_learning_rate: 1.83e-06
learning_rate: 3.66e-06
layerwise_decay: 0.983
sent_layer: mix
layer_transformation: sparsemax
layer_norm: False
loss: mse
dropout: 0.1
batch_size: 32
activations: Tanh
input_segments:
- mt
- src
- ref
word_level_training: False
```

Listing 2: Relevant hyperparameters used for training Phase II and III models (Section 4.1) using COMET framework. Note that the only difference between the two phases is the loss λ parameter (Eq. 3).

```
class_path: UnifiedMetric
init_args:
nr_frozen_epochs: 0.3
keep_embeddings_frozen: True
optimizer: AdamW
encoder_learning_rate: 1.0e-06
learning_rate: 3.66e-06
layerwise_decay: 0.983
sent_layer: mix
layer_transformation: sparsemax
layer_norm: False
loss: mse
dropout: 0.1
batch_size: 32
activations: Tanh
input_segments:
- mt
- src
- ref
word_level_training: true
loss_lambda: 0.983 (II) / 0.055 (III)
error_labels:
- minor
- major
- critical
cross_entropy_weights:
- 0.08
- 0.486
- 0.505
- 0.533
```

ERROR	zh-en	he-en
Add. of text	1516	1490
Negation	498	637
Mask in-fill	1716	1659
Swap NUM	313	214
Swap NE	519	586
Total	4762	4586

Table 8: Number of examples for each category, synthetically-created using SMAUG (Alves et al., 2022) for zh-en and he-en

SMAUG. We also provide examples of each error category in 9.

C.2 Additional results

We provide additional results on the impact of SMAUG perturbations for zh-en and he-en data for both the XL and XXL xCOMET models in Figure 5.

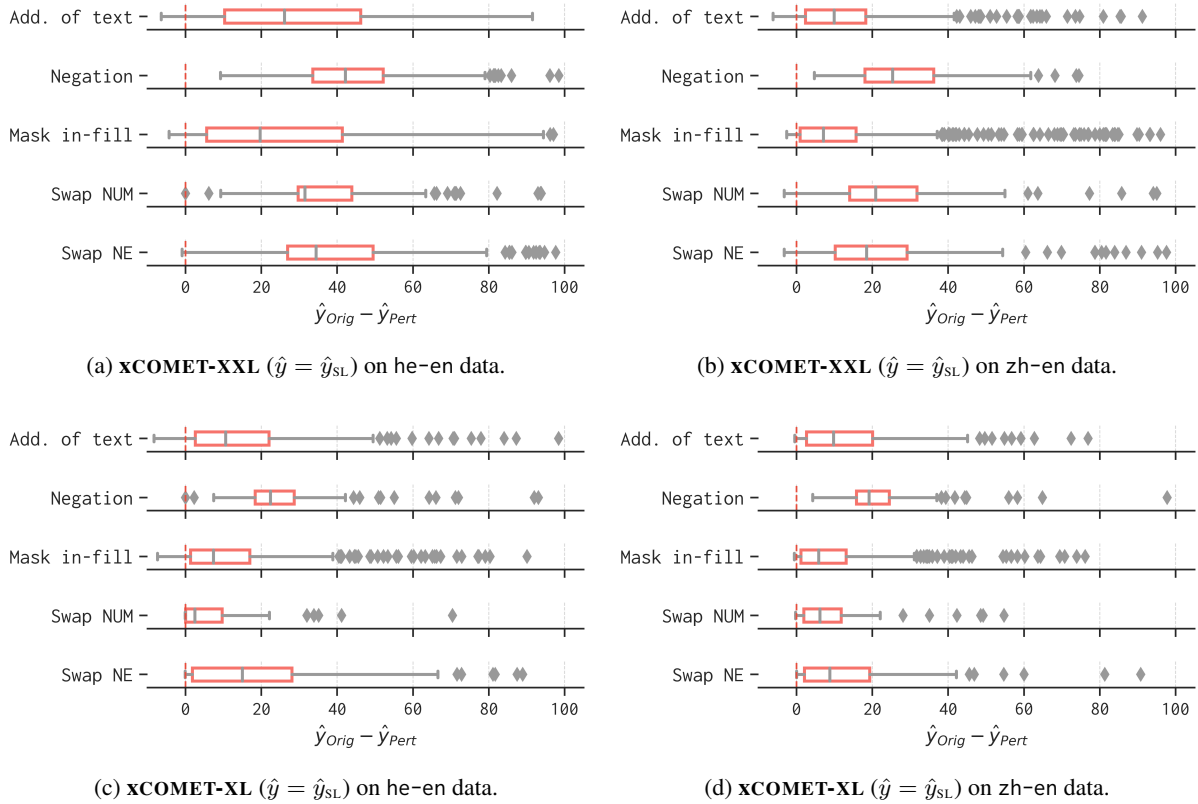


Figure 5: Impact of the perturbations, as measured by the difference in xCOMET-XXL ($\hat{y} = \hat{y}_{SL}$) between the original and the perturbed translation with both models.

D Examples of hallucinations from the benchmark of Guerreiro et al. (2023b)

We provide in Table 10 examples of hallucinations from the benchmark that we used, alongside the predicted error spans by xCOMET.

Source:

“这无疑是在火上浇油，支持了对华强硬派的观点”。

Translation:

"This undoubtedly added fuel to the fire, supporting the views of the hardliners towards China."

Reference:

"This undoubtedly adds fuel to the fire and supports the views of the hardliners towards China."

Add. of text:

he said the media was not "responsible" for such behavior. "This undoubtedly added fuels to the fire, thereby supporting the views of the hardliners towards China."

Source:

请注意本文件的某些内容可能涉及专利。

Translation:

Please be aware that this text may include references to patents in some places.

Reference:

Please be aware that this text may include references to patents in some places.

Negation:

Please be aware that this text **does not** include references to patents in some places.

Source:

任何本文件适用的产品均应满足实现级别 1 的要求。

Translation:

Any product applicable to this document should meet the requirements of implementation level 1.

Reference:

Each item to which this record relates must follow to the Execution Level 1 details.

Mask in-fill:

Any product applicable to **the project, which does not** meet the requirements of implementation level 1.

Source:

算我倒霉，懒得为了个200块的东西被你们折腾。

Translation:

I'm unfortunate because I'm too lazy to be hurled by you for 200 yuan stuff.

Reference:

I'm unfortunate because I'm too lazy to be hurled by you for 200 yuan stuff.

Swap NUM:

I'm unfortunate because I'm too lazy to be hurled by you for the **5** yuan stuff.

Source:

本文件由中华人民共和国工业和信息化部提出并归口。

Translation:

This document was proposed and managed by the Ministry of Industry and Information Technology of the People's Republic of China.

Reference:

The People's Republic of China's Ministry of Industry and Information Technology is responsible for the proposal and administration of this document.

Swap NE:

This document was proposed and managed by the **Government** of the People's Republic of China.

Table 9: Synthetically-generated errors (highlighted in pink) created with SMAUG (Alves et al., 2022) for zh-en to assess whether xCOMET can detect localized major and critical errors.

Source:	Touristen in Portugal in Panik versetzt, nachdem ein tieffliegender Militärjet Strand überfliegt
Translation:	Is the whole pancreas removed from each pancreas of pancreas?
Reference:	Is the entire pancreas always removed in pancreatic cancer operations?
Oscillatory hallucination with xCOMET error span predictions:	Is the CRIT: whole pancreas removed from each pancreas of pancreas?
<hr/>	
Source:	Das Teilabonnement für international tätige Juristen.
Translation:	The sub-sub-sub-sub-scription for international lawyers.
Reference:	Partial subscriptions for internationally active lawyers.
Oscillatory hallucination with xCOMET error span predictions:	The CRIT: sub-sub-sub-sub-scription for international lawyers.
<hr/>	
Source:	Nur der Stuck an der Decke uns das Wandabendmahl konnten nicht restauriert werden.
Translation:	The room was very small and the bathroom was very small.
Reference:	Only the stuccos of the ceiling and the wall could not be restored.
Fully detached hallucination with xCOMET error span predictions:	CRIT: The room was very small and the bathroom was very small.
<hr/>	
Source:	Empfehlenswert gleich mit der Zimmerreservierung zu buchen!
Translation:	The staff were very friendly and helpful. The room was clean and comfortable.
Reference:	We recommend booking your treatments together with the hotel booking!
Fully detached hallucination with xCOMET error span predictions:	CRIT: The staff were very friendly and helpful. The room was clean and comfortable.

Table 10: Examples of predictions of xCOMET-XXL for the hallucination benchmark of [Guerreiro et al. \(2023b\)](#).