

Scaling up COMETKIWI: Unbabel-IST 2023 Submission for the Quality Estimation Shared Task

Ricardo Rei^{*1,2,4}, Nuno M. Guerreiro^{*1,3,4}, José Pombal¹, Daan van Stigt¹,
Marcos Treviso^{3,4}, Luisa Coheur^{2,4}, José G. C. de Souza¹, André F. T. Martins^{1,3,4}

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

³Instituto de Telecomunicações, Lisbon, Portugal

⁴Instituto Superior Técnico, University of Lisbon, Portugal

Abstract

We present the joint contribution of Unbabel and Instituto Superior Técnico to the WMT 2023 Shared Task on Quality Estimation (QE). Our team participated on all tasks: sentence- and word-level quality prediction (task 1) and fine-grained error span detection (task 2). For all tasks, we build on the COMETKIWI-22 model (Rei et al., 2022b). Our multilingual approaches are ranked first for all tasks, reaching state-of-the-art performance for quality estimation at word-, span- and sentence-level granularity. Compared to the previous state-of-the-art, COMETKIWI-22, we show large improvements in correlation with human judgements (up to 10 Spearman points). Moreover, we surpass the second-best multilingual submission to the shared-task with up to 3.8 absolute points.

1 Introduction

Quality Estimation (QE) is the task of automatically assigning a quality score to a machine translation output without depending on reference translations (Specia et al., 2018). This paper details the collaborative effort of Unbabel and Instituto Superior Técnico (IST) in the WMT23 Quality Estimation shared task, which encompassed two primary tasks: (i) sentence- and word-level quality prediction and (ii) fine-grained error span detection.

As of last year, some language pairs in the test set were absent from the training data. To address this, following a similar approach to the previous year, our systems were developed to achieve good multilingual generalization and to accommodate previously unseen languages. To achieve this, we start by leveraging the direct assessments (DA) labeled data obtained from the WMT Metrics shared task from 2017 to 2020, the MLQE-PE dataset (Fomicheva et al., 2022), and the training data (DA) specifically annotated for Indian languages in the 2023 shared task edition. In total, these datasets

encompass close to 1M annotations covering 38 language pairs. We start by constructing generic models using this corpus. These generic QE models were subsequently fine-tuned for this year’s subtasks.

For Task 1 – sentence-level, we fine-tuned our generic models exclusively with this year’s DA data. The architecture of these models remains consistent with our submission from the previous year, but we employ XLM-R XL and XXL as pretrained encoders (Conneau et al., 2020). For the word-level quality prediction task, we follow the successful approach of combining the sentence- and word-level signals into one loss during the finetuning step, which has yielded positive results in previous iterations (Rei et al., 2022b). For fine-grained error span detection, we conducted experiments exploring various approaches that build upon our word-level and sentence-level strategies. In terms of contrasting systems, we explored UnbabelQi¹ and GPT-4 (OpenAI, 2023). For GPT-4, we used a prompt designed to predict both the location and severity of errors in each translation, akin to the approach used in AutoMQM (Fernandes et al., 2023).

Overall, our main contributions are: (i) we introduce approaches for multilingual machine translation quality estimation that are consistently first-ranked at word-, span-, and sentence-level granularity; (ii) we explore different approaches to predict the span of problematic translations along with their error severities (OK, MINOR, MAJOR); (iii) we publicly release two of our best models for research purposes (COMETKIWI-XL² and -XXL³). To the best of our knowledge, these are the largest open-source QE models publicly released.

Our submitted systems attain the top multilingual results in all tasks: For Task 1 sentence-

¹<https://qi.unbabel.com/>

²<https://huggingface.co/Unbabel/wmt23-cometkiwi-da-xl>

³<https://huggingface.co/Unbabel/wmt23-cometkiwi-da-xxl>

*Equal contribution. ✉ ricardo.rei@unbabel.com

level prediction, our multilingual system achieves 59.4 Spearman correlation points, surpassing the second-best system by nearly 4 absolute points. For word-level, our system achieves a 31.7 MCC score, outperforming the second-best system by almost 2 absolute MCC points. For error span prediction, our multilingual system achieves a 22 F1.0 score, beating the second-best system by more than 5 F_1 points.

2 Overview of the shared-task

QE systems are designed according to the granularity in which predictions are made (e.g., sentence- or word-level QE). In sentence-level QE, the goal is to predict a single quality score $\hat{y} \in \mathbb{R}$ given the whole source and its translation as input. Word-level QE works at a lower granularity level, with the goal of predicting binary quality labels $\hat{y}_i \in \{\text{OK}, \text{BAD}\}$ for all $1 \leq i \leq n$ *machine-translated words*, indicating whether that word is a translation error. In fine-grained error span detection, systems are tasked with flagging which parts of the segment, i.e., sequences of consecutive characters, contain errors. If an error span is found, the system has to point out its severity; in this shared task, an error span’s severity can be classified as MINOR or MAJOR. We sometimes refer to the parts of the segment that do not belong to an error span as being labelled as OK. We participated on all tasks of this year’s shared-task. We specify the language pairs and the released data below:

Task 1 – Sentence-level quality prediction: Submissions for this task were evaluated based on their correlation with Direct Assessment (DA) annotations for five language pairs: English→Marathi (*en-mr*), English→Hindi (*en-hi*), English→Tamil (*en-ta*), English→Telugu (*en-te*), and English→Gujarati (*en-gu*). Furthermore, they were evaluated using Multidimensional Quality Metrics (MQM) annotations for three language pairs: English→German (*en-de*), Chinese→English (*zh-en*), and English→Hebrew (*he-en*). Training data was made available for all language directions except for *he-en*.

Task 1 – Word-level quality prediction: Submissions for this task underwent evaluation based on tags inferred from post-editions for English→Farsi (*en-fa*) and English→Marathi (*en-mr*). Additionally, they were assessed using MQM annotations for *en-de*, *zh-en*, and *he-en*. No addi-

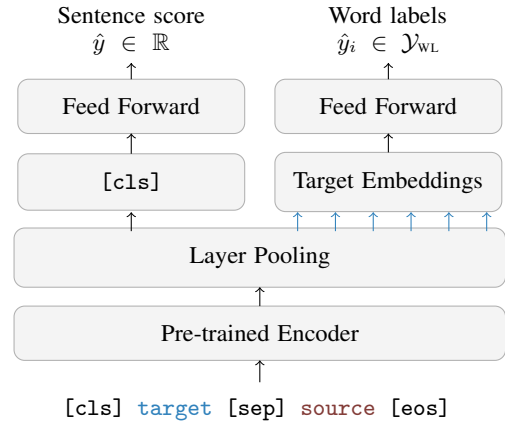


Figure 1: Our model follows COMETKIWI for sentence-level (left part) and word-level QE (right part). We represent the output space of the word-level head by \mathcal{Y}_{WL} .

tional training or development data with word-level tags were made available. To the best of our knowledge, no word-level data is available for *en-fa* and *he-en*.

Task 2 – Fine-grained error span detection: submissions were evaluated on error spans obtained via MQM annotations for 3 language pairs (*en-de*, *zh-en* and *he-en*). No training nor development data is available for *he-en*.

3 Implemented Systems

We largely follow the architecture of COMETKIWI (Rei et al., 2022b) – see Figure 1 for an illustration. We concatenate the machine translated sentence $\mathbf{t} = \langle t_1, \dots, t_n \rangle$ and its source sentence counterpart $\mathbf{s} = \langle s_1, \dots, s_m \rangle$ to serve as input to the encoder. This encoder then produces hidden state matrices $\mathbf{H}_0, \dots, \mathbf{H}_L$ for each layer $0 \leq \ell \leq L$, where $\mathbf{H}_\ell \in \mathbb{R}^{(n+m) \times d}$, where $\ell = 0$ corresponds to the embedding layer and d is the hidden size. Following this, all hidden states are fed to a scalar mix module (Peters et al., 2018) that learns a weighted sum of the hidden states of each layer of the encoder, producing a sequence of aggregated hidden states \mathbf{H}_{mix} as follows:

$$\mathbf{H}_{\text{mix}} = \lambda \sum_{\ell=0}^L \beta_\ell \mathbf{H}_\ell. \quad (1)$$

Here λ is a scalar trainable parameter, $\beta \in \Delta^L$ is given by $\beta = \text{sparsemax}(\phi)$ using a sparse transformation (Martins and Astudillo, 2016), with $\phi \in \mathbb{R}^L$ as learnable parameters, and where we

denote by $\Delta^L := \{\beta \in \mathbb{R}^L : \mathbf{1}^\top \beta = 1, \beta \geq 0\}$ the probability simplex.⁴

For sentence-level models, we use the hidden state of the <cls> token as the sentence representation, which, in turn, is passed to a 2-layered feed-forward module in order to get a sentence score prediction $\hat{y} \in \mathbb{R}$. For word-level and error span detection models, we first retrieve the hidden state vectors associated with each token in t , and then pass them to a linear projection to get word-level predictions $\hat{y}_i \in \mathcal{Y}_{\text{WL}}, \forall 1 \leq i \leq n$. The output space of the word-level predictions is different depending on whether the models are constructed for word-level quality prediction ($\mathcal{Y}_{\text{WL}} = \{\text{OK}, \text{BAD}\}$), or error span detection ($\mathcal{Y}_{\text{WL}} = \{\text{OK}, \text{MINOR}, \text{MAJOR}\}$).

Pretrained multilingual encoders. Similarly to (Rei et al., 2022b), we employ InfoXML L (Chi et al., 2021).⁵ Additionally, we experiment with scaled-up multilingual encoders, including XLM-R XL,⁶ and XLM-R XXL.⁷ InfoXML L comprises 24 encoder blocks with 16 attention heads each, totaling 550M parameters. XLM-R XL and XLM-R XXL have 32 attention heads for each encoder block, 36 and 48 encoder blocks and a total of 3.5B and 10.7B parameters, respectively.

Generic models for all tasks. We create, for each model size, a generic model that will then be further adapted to each separate task. To train these models, we use the collective corpora from 2017 to 2019 DA annotations of the WMT Translation shared task, and the MLQE-PE corpus (Fomicheva et al., 2022). We include the human annotations respective to the language pairs of this year’s shared task for 7 different language pairs: DA annotations for *en-mr*, *en-hi*, *en-ta*, *en-te*, *en-gu*, and MQM annotations for *en-de* and *zh-en*. Overall, the generic models are trained on sentence-level quality prediction with over 940k samples with source, translation and quality score on 38 different language pairs. When presented with multiple DA scores for the same sentence pair, we used the z-score of the DAs for training but we first normalize the DAs between 0 and 1, where 1 represents a perfect

⁴As it has been shown in (Rei et al., 2022a) not all layers are relevant and thus, using sparsemax we learn to ignore layers that do not help in the task at hands.

⁵<https://huggingface.co/microsoft/infotxm-large>

⁶<https://huggingface.co/facebook/xlm-roberta-xl>

⁷<https://huggingface.co/facebook/xlm-roberta-xxl>

translation and 0 a random one.

Task adaptation. After having obtained the generic models, we will train models for each separate stream of the shared-task, i.e., sentence-level, word-level or error span prediction. To do so, we consider the multi-task optimization from Rei et al. (2022b) wherein sentence scores can be used alongside supervision from word-level tags. Formally,

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

$$\mathcal{L}_{\text{WL}}(\theta) = -\frac{1}{n} \sum_{i=1}^n w_{y_i} \log p_{\theta}(y_i) \quad (3)$$

$$\mathcal{L}(\theta) = \lambda_{\text{SL}} \mathcal{L}_{\text{SL}}(\theta) + \lambda_{\text{WL}} \mathcal{L}_{\text{WL}}(\theta), \quad (4)$$

where $w \in \mathbb{R}^{|\mathcal{Y}_{\text{WL}}|}$ represents the class weights given for the word-level tags,⁸ and $\lambda_{\text{SL}}, \lambda_{\text{WL}} \in \mathbb{R}_+$ are used to weigh the sentence and word-level losses, respectively. Note that $\lambda_{\text{SL}} = 1$ and $\lambda_{\text{WL}} = 0$ yields a fully sentence-level model, whereas $\lambda_{\text{SL}} = 0$ and $\lambda_{\text{WL}} = 1$ yields a word-level model.

Using unconstrained models. For error span detection, we evaluate UnbabelQi, an Unbabel demo QE system, alongside GPT4 (OpenAI, 2023). We prompt GPT4 to produce an MQM annotation for each source-target pair, based on five-shot examples which vary across language pairs but are consistent within segments of the same language pair. We also apply this system in Task 1, deriving a sentence-level score from error spans, in alignment with the MQM framework. This approach bears similarity to AutoMQM (Fernandes et al., 2023).

3.1 Task 1: Quality prediction

After the pretraining phase, we *further separately* adapt the generic models to the released DA and MQM data for this year’s shared task.

3.1.1 Sentence-level quality prediction

Adaptation for Sentence-level. In order to tailor our models to the language pairs featured in this year’s shared task, we conducted full fine-tuning until convergence on the released validation set. This fine-tuning exclusively leveraged the recently released Direct Assessment (DA) annotations for this year’s task. This approach yields additional improvements for those languages. In the case of the MQM language pairs, our preliminary experiments revealed that attempting significant performance

⁸These parameters help control how much we penalize the different granularities of word-level errors.

improvements on the MQM data led to noteworthy drops in correlations for the other language pairs using DAs. Consequently, for the MQM language pairs, we opted to employ the generic models as they are.

Ensembling models. Similarly to [Rei et al. \(2022b\)](#), we use Optuna ([Akiba et al., 2019](#)) to assemble four models – two XL and two XXL – into a single system. We do so by finding the optimal weights for each language pair among these four multilingual models, and combining their predictions according to those weights. Notably, the XXL models are generic models, whereas the two XL checkpoints were further optimized with this year’s shared task data. As expected, the XL models carry more weight for Indian languages, while the XXL generic models were deemed more crucial for MQM languages.

3.1.2 Word-level quality prediction

For the word-level QE tasks, we experimented with both the multi-task setting and word-labels only.

Training word-level models. This year, no training or development data with word-level tags were made available. As such, the training data for our models consists of the training data used in [Rei et al. \(2022b\)](#), combined with the development sets from the 2022 WMT Shared Task. As the word-level task was going to be tested in a zero-shot scenario for two out of five language pairs (*en-fa*, *he-en*), contrary to [Rei et al. \(2022b\)](#), we do not prepend a language prefix to the beginning of the source and target segments during training. Moreover, for the post-edit (PE) models, we removed samples from two language pairs (*ps-en* and *en-cs*) from the training data. We did so to assess, during validation, the models’ capability to generalise in a zero-shot scenario. For the MQM models, we used all available annotations, including those in *en-ru*.

Ensembling models. For word-level we followed a similar ensembling technique used for sentence-level. Specifically, we combined multiple systems trained with different hyperparameters, encoder size and pre-training setups. In the case of word-level predictions, we aggregate multiple predictions into OK/BAD tags by following the *ensemble-tags* procedure from [Rei et al. \(2022b\)](#). In this approach, we combine the predicted tags of each model: for every input segment, we get a combined tag, $\alpha \sum_{i \in \mathcal{M}} w_i c_i$, where c_i is the tag

predicted by the model and α is the weight for the BAD tag. We use Optuna to determine the optimal weights w_i for each model and the optimal BAD weight α for each LP. In the final submission, we combine six models (five PE models and one MQM model). Five of these models use InfoXLM as the encoder model, and one PE model uses XLM-R XL.⁹ Refer to Table 2 for the test set results.

3.2 Task 2: Fine-grained error span detection

In this task, we investigated three distinct approaches. The first approach extends word-level models by modifying their output predictions. More precisely, it involves transforming consecutively predicted BAD tags into character-level error spans, rather than categorizing individual words based on the first subword. To determine the error severities of these spans, we considered two options: labeling all the subwords within the span as either MINOR or MAJOR. Our best results were achieved with the latter approach.

The second approach leverages xCOMET ([Guerreiro et al., 2023](#)) in conjunction with a pseudo-reference obtained from DeepL or Google Translate.¹⁰ Similar to our models from Task 1 word-level, xCOMET is trained with a multitask objective. Additionally, xCOMET is simultaneously optimized for both reference-free and reference-based evaluation, following UNITE ([Wan et al., 2022](#)). During inference, xCOMET can leverage a reference translation to enhance error identification. Since we employ a pseudo-reference that may contain translation errors, we initially assess the quality of the pseudo-reference using a generic QE system from Task 1 (*reference_score*). For all pseudo-references with a score below 0.5, we run xCOMET with QE-only input. For pseudo-references scoring above 0.5, the input weights for xCOMET are determined as follows:

$$\begin{aligned} \text{diff} &= 1 - \text{reference_score} \\ \text{src_weight} &= 2 \cdot \text{diff} \\ \text{ref_weight} &= (1 - \text{src_weight}) \cdot 0.4 \\ \text{uni_weight} &= (1 - \text{src_weight}) \cdot 0.6 \end{aligned}$$

Here, *src_weight* represents the weight assigned to the source-only input, *ref_weight* de-

⁹We found it hard to obtain performance boosts by scaling up to XLM-R XL on the word-level task. As such, we did not experiment with XLM-R XXL.

¹⁰We choose the best translation using the generic XXL model from task 1.

Encoder	DA					MQM			avg.
	en-mr	en-hi	en-ta	en-te	en-gu	en-de	zh-en	he-en [†]	
2nd place (Yan et al., 2023)									0.556
InfoXLM L	<i>CometKiwi-22 (Rei et al., 2022b)</i>					0.413	0.476	0.619	0.485
	<i>Generic models</i>								
InfoXLM L	0.661	0.505	0.641	0.282	0.661	0.422	0.448	0.610	0.529
XLM-R XL	0.664	0.536	0.607	0.335	0.637	0.422	0.469	0.624	0.537
XLM-R XXL	0.685	0.520	0.670	0.326	0.655	0.443	0.476	0.662	0.555
	<i>Further adapted models for sentence-level</i>								
XLM-R XL	0.684	0.583	0.682	0.386	0.683	0.434	0.441	0.696	0.574
XLM-R XXL	0.693	0.555	0.738	0.359	0.701	0.434	0.457	0.661	0.575
	<i>Final Ensemble</i>								
Ensemble 4x	0.702	0.598	0.739	0.389	0.714	0.448	0.493	0.668	0.594
	<i>GPT4-based model</i>								
GPT4-QE	0.379	0.212	0.146	0.174	0.297	0.442	0.412	0.488	0.319

Table 1: Results for sentence-level QE in terms of Spearman correlation. We represent zero-shot LPs with †.

notes the typical metric input (reference-only input), and `uni_weight` represents a unified input where the model receives all three sentences (translation, source, and reference). Notably, for pseudo-references with a QE score of 1, we rely solely on a reference-only input and the unified input. We refer to this approach as xCOMET-PS-REF.

We also contrast the aforementioned approaches with two unconstrained QE systems: UnbabelQI and GPT-4, as mentioned in Section 3. We refer to these approaches as UNBABELQI and GPT4-QE, respectively.

4 Experimental Results

We present the results on the official test set for each of the tasks for multiple model/data configurations. Sentence-level submissions were evaluated using the Spearman rank correlation. Pearson and Kendall correlation were also used as secondary metrics, but here we report only Spearman since it was the primary metric used to rank systems. word-level submission were evaluated using MCC, F_1 -OK, and F_1 -BAD, but we report only MCC as it was considered the main metric. Error span detection was evaluated using F_1 score in which the positive labels are all the characters belonging to erroneous spans. Furthermore, each true positive is downweighted to half if the system failed to classify the error span’s severity (e.g., MINOR instead of MAJOR). The submitted systems were independently evaluated on in-domain and zero-shot LPs for direct assessments and MQM.

4.1 Quality Estimation

Sentence-level. Results for sentence-level are presented in Table 1. Results indicate that retraining the system from the previous year, specifically COMETKIWI with InfoXLM, using data that encompasses this year’s DA, leads to significant improvements. Remarkably, this improvement in correlations is achieved while maintaining the same level of correlations for *en-de* (a high-resource language pair for which both models share the same data) and *he-en*, a language pair that both models had not seen during training. Surprisingly, there was a drop in correlations for *zh-en* even though both models saw the same *zh-en* data. Nevertheless, the overall performance of the newly retrained version improved by 4.4 Spearman points.

As anticipated, among the three backbone transformers, the XXL model is the top performer, with significant improvements across all language pairs when compared to InfoXLM. Moreover, additional finetuning on this year’s training data results in further improvements for the Indian languages. Notably, concerning the MQM data, this supplementary finetuning step not only preserves performance but sometimes even increases it. Similar to last year, the ensemble of high-performing models once again makes up our best submission.

Finally, despite performing well in Task 2, GPT4-QE shows poor correlations at sentence-level prediction with the exception of the *en-de* for which GPT4-QE, although lagging behind the ensemble approach, surpasses our individual models.

Method	Post-edit		MQM			avg.
	en-fa [†]	en-mr	en-de	zh-en	he-en [†]	
Baseline (Rei et al., 2022b)	0.293	0.287	0.179	0.225	0.275	0.226
2nd place (Yan et al., 2023)						0.298
<i>Adapted models for word-level</i>						
PE model (InfoXLM L)	0.343	0.343	0.227	0.253	0.382	0.310
PE model (XLM-R XL)	0.325	0.344	0.255	0.197	0.306	0.285
MQM model (InfoXLM L)	0.296	0.252	0.215	0.269	0.334	0.273
<i>Final Ensemble</i>						
Ensemble PE + MQM	0.345	0.347	0.246	0.302	0.402	0.317

Table 2: Results for word-level QE in terms of MCC for the post-edit and MQM LPs. The ensemble is composed by multiple post-edit and MQM models. We represent zero-shot LPs with †.

Word-level. We report the best individual systems Table 2. Our best individual systems were trained on top of the InfoXLM L generic model. For PE models, we used multi-task objective in Eq. 4, as we found that combining the sentence-level and word-level loss was beneficial. However, for MQM models, we trained word-level only models, by setting $\lambda_{SL} = 0.0$ and $\lambda_{WL} = 1.0$.

Interestingly, we found that PE models are very competitive on MQM language pairs. For example, the best overall performance for *he-en* was actually obtained with a PE word-level model. This is also reflected on the Optuna weights obtained for our final ensemble, wherein the weights of the PE models are significantly higher than those of the MQM models for all language pairs but *en-de*. In fact, our final ensemble for *en-zh* and *en-he* consists solely of PE models trained with different learning rates, λ_{SL} , λ_{WL} and w . Further investigation on two different vectors may lead to improved word-level models: (i) balancing DA and MQM word-level annotations, and (ii) appropriately leveraging the larger capacity of scaled up encoder models.

Fine-grained error span detection. Results for fine-grained error span detection are shown in Table 3. Using a word-level model to obtain error span predictions leads to reasonable performance, comparable to our unconstrained submission, UNBABELQI, a model directly tasked with error span detection. That said, xCOMET-PS-REF, an error span detection model, surpassed both of the previous approaches. We attribute the improved performance to this system being an ensemble of two significantly larger models, and to the usage of a pseudo-reference. We found the latter to be particularly beneficial on *he-en*, a language pair for which we had no training data.

Method	en-de	zh-en	he-en [†]	avg.
2nd place (Li et al., 2023)				0.165
Baseline	0.167	0.219	0.083	0.156
WORD-LEVEL	0.235	0.272	0.105	0.204
xCOMET-PS-REF	0.259	0.270	0.125	0.218
UNBABELQI	0.249	0.227	0.111	0.196
GPT4-QE	0.273	0.265	0.121	0.220

Table 3: Results for fine-grained error span detection (Task 2). Evaluation metric is F_1 score. We represent zero-shot LPs with †. The first two systems are constrained while the other two are unconstrained submissions.

The best approach in terms of average F_1 was GPT4-QE, mostly due to the improved performance on *en-de*. While this is a promising finding for LLM-based quality estimation systems, there are limitations. First, obtaining a sentence-level score from the error spans (as per the MQM framework) leads to poor correlations with human judgements derived from DA (see Table 1) and with low-resource language-pairs like *he-en*. Second, despite being useful in practice and leading to gains in F_1 , it is hard to control GPT’s precision and recall. We found that the number of examples included in the prompt, their ordering, and the number of errors within each example led to noticeable changes in the system’s propensity to flag errors. Thirdly, running QE with a system such as GPT-4 is expensive and slow even for a shared task exercise.

5 Final Remarks

We describe Unbabel and IST joint submission to WMT23 QE shared task. Our approaches correlate well with human judgements for all the three granularities of translation quality prediction, ranking first in all multilingual tasks and surpassing the pre-

vious state-of-the-art model, COMETKIWI-22, by up to 10 Spearman correlation points. Overall, our models follow the same architecture of last year’s participation, COMETKIWI. However, this year we leverage more data and larger encoder models. Our best final systems are ensembles of different models trained on DA, post-edits or MQM scores that complement each other. Interestingly, our best systems surpass GPT-4 by a large margin for sentence-level translation quality prediction, and they are comparable to GPT-4 at error span detection.

Acknowledgements

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 (DECOL-LAGE, 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).

References

- Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. 2019. [Optuna: A next-generation hyperparameter optimization framework](#). In *Proceedings of the 25rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*.
- Zewen Chi, Li Dong, Furu Wei, Nan Yang, Saksham Singhal, Wenhui Wang, Xia Song, Xian-Ling Mao, Heyan Huang, and Ming Zhou. 2021. [InfoXLM: An information-theoretic framework for cross-lingual language model pre-training](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3576–3588, Online. Association for Computational Linguistics.
- 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*, pages 8440–8451, Online. Association for Computational Linguistics.
- Patrick Fernandes, Daniel Deutsch, Mara Finkelstein, Parker Riley, André F. T. Martins, Graham Neubig, Ankush Garg, J. Clark, Markus Freitag, and Orhan Firat. 2023. [The devil is in the errors: Leveraging large language models for fine-grained machine translation evaluation](#). *ArXiv*, abs/2308.07286.
- Marina Fomicheva, Shuo Sun, Erick Fonseca, 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 Language Resources and Evaluation Conference*, pages 4963–4974, Marseille, France. European Language Resources Association.
- 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](#).
- Yuang Li, Chang Su, Ming Zhu, Xinglin Piao, Mengyao Lyu, Min Zhang, and Hao Yang. 2023. [HW-TSC 2023 Submission for the Quality Estimation Shared Task](#). In *Proceedings of the Eighth Conference on Machine Translation*, Singapore. Association for Computational Linguistics.
- Andre Martins and Ramon Astudillo. 2016. [From softmax to sparsemax: A sparse model of attention and multi-label classification](#). In *International Conference on Machine Learning*, pages 1614–1623.
- OpenAI. 2023. [GPT-4 technical report](#).
- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep contextualized word representations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237, New Orleans, Louisiana. Association for Computational Linguistics.
- Ricardo Rei, Ana C Farinha, José G.C. de Souza, Pedro G. Ramos, André F.T. Martins, Luisa Coheur, and Alon Lavie. 2022a. [Searching for COMETINHO: The little metric that could](#). In *Proceedings of the 23rd Annual Conference of the European Association for Machine Translation*, pages 61–70, Ghent, Belgium. European Association for 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. 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.
- Lucia Specia, Carolina Scarton, and Gustavo Henrique Paetzold. 2018. [Quality estimation for machine translation](#). *Synthesis Lectures on Human Language Technologies*, 11(1):1–162.
- 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)*, pages 8117–8127, Dublin, Ireland. Association for Computational Linguistics.

Zeyu Yan, Wenbo Zhang, Qiaobo Deng, Hongbao Mao, Jie Cai, and Zhengyu He. 2023. IOL Research’s Submission for WMT 2023 Quality Estimation Shared Task. In *Proceedings of the Eighth Conference on Machine Translation*, Singapore. Association for Computational Linguistics.