

# MetricX-24: The Google Submission to the WMT 2024 Metrics Shared Task

Juraj Juraska, Daniel Deutsch, Mara Finkelstein and Markus Freitag

Google

{jjuraska,danddeutsch,marafin,freitag}@google.com

## Abstract

In this paper, we present the MetricX-24 submissions to the WMT24 Metrics Shared Task and provide details on the improvements we made over the previous version of MetricX. Our primary submission is a hybrid reference-based/-free metric, which can score a translation irrespective of whether it is given the source segment, the reference, or both. The metric is trained on previous WMT data in a two-stage fashion, first on the DA ratings only, then on a mixture of MQM and DA ratings. The training set in both stages is augmented with synthetic examples that we created to make the metric more robust to several common failure modes, such as fluent but unrelated translation, or undertranslation. We demonstrate the benefits of the individual modifications via an ablation study, and show a significant performance increase over MetricX-23 on the WMT23 MQM ratings, as well as our new synthetic challenge set.<sup>1</sup>

## 1 Introduction

Automatic evaluation metrics are critical to the development of machine translation (MT) systems. In recent years, the landscape of MT evaluation has changed dramatically since the use of lexical metrics, like BLEU (Papineni et al., 2002) and ChrF (Popović, 2015), that compared the tokens or characters of the candidate translation to a reference translation to predict a scalar score that represents the quality of the translation. Evaluation metrics based on neural networks opened up the door for more experimentation, and metrics now vary based on what type of output they produce, what they require as input for prediction, and whether they use a dedicated evaluation model or a general-purpose large language model.

This paper provides details on MetricX-24, the successor to MetricX-23. MetricX is a learned

<sup>1</sup>Our code and models can be found at <https://github.com/google-research/metricx>.

regression-based metric trained to predict a floating point score representing the quality of a translation. This year, we made four submissions to the WMT24 Metrics Shared Task, all based on the mT5 language model (Xue et al., 2021), which is further fine-tuned on direct assessment (DA) ratings, MQM ratings (Lommel et al., 2014; Freitag et al., 2021), and newly constructed synthetic data. The primary submission, denoted MetricX-24-Hybrid, is a hybrid reference-based/-free metric, which can score a translation irrespective of whether it is given the source segment, the reference, or both. The same model is thus the primary submission for both the reference-based evaluation and the quality estimation (QE) task, having predicted the scores once with and once without the references provided in the input. Our contrasting submissions, MetricX-24(-QE), are standalone reference-based/QE models, trained only for their specific task.

The key takeaways from our experiments, detailed in this report, include:

1. Learned metrics cannot reliably detect undertranslation, duplication, missing punctuation, and fluent but unrelated translation;
2. Adding a relatively small amount of synthetic data to the training set can boost the metric’s performance, especially on lower-quality translations with the above issues;
3. It is possible to effectively train a metric on a mixture of MQM and DA ratings, thus maintaining high performance on a larger set of language pairs;
4. Training a metric in the hybrid input mode, i.e., with and without the reference included in the input, allows it to learn to rely less on the reference when it is of poor quality.

## 2 Data

Developing MetricX-24, we relied solely on publicly available data from the WMT Metrics shared

tasks between 2015 and 2023. The translation ratings from these years come in two different flavors: (1) direct assessment (DA) scores on a scale from 0 to 100, collected in general from non-expert raters, and (2) MQM scores (Lommel et al., 2014; Freitag et al., 2021) on a scale from 0 to 25 (with 0 being the best), which are grounded on error spans and their corresponding severity levels, annotated by professional raters. MQM ratings have been collected as part of the WMT campaign only since 2020 and, because the annotations are considerably more time-consuming and expensive to obtain, they are only available for a few language pairs. The DA scores, on the other hand, offer a broader language coverage of nearly 50 language pairs, but the raw ratings are noisy (due to different rating strategies) and generally of lower quality. Therefore, it is often beneficial to  $z$ -normalize DA ratings per rater before training models on them, so as to make the ratings more comparable across different annotators. In contrast, models do not benefit from MQM scores being  $z$ -normalized because the scores come from a rather small group of annotators and they adhere to a rubric.

In the rest of this section, we provide details on which data we use for training and evaluation, as well as how the different datasets are preprocessed. Furthermore, we describe new synthetic data we created from the WMT datasets, with the goal of addressing some of MetricX’s known failure modes.

## 2.1 Training Data

**DA.** We utilize most of the DA data from the 2015–2022 period for training, with the following exceptions. As we observed during the development of the previous version of MetricX (Juraska et al., 2023), the into-English portion of the WMT21 DA ratings drags the model performance down. We confirmed this observation again this year and excluded these language pairs from the training data. With the gradually declining quality of DA ratings collected for WMT using the MTurk platform, we also exclude all into-English language pairs from WMT22.<sup>2</sup> Additionally, we exclude the en-zh language pair from WMT22, as we use the equivalent slice of data, but with MQM ratings, for evaluation. We use  $z$ -normalized ratings when training models on DA data only, but raw ratings

<sup>2</sup>One exception is zh-en, for which DA ratings were collected in two different ways, including using the same method and framework as the out-of-English language pairs (Kocmi et al., 2022).

when training on a mixture of MQM and DA data.

**MQM.** Besides the DA ratings, we also take advantage of the higher-quality MQM ratings from the years up to 2022 for training. These include four language pairs: en-de, en-ru, en-zh and zh-en.<sup>3</sup> We only use the conversation, e-commerce and social domains from WMT22 en-zh for training. In our experiments with different subsets of MQM ratings, we observed a consistent boost in performance with the 2020 data excluded, hence, our final models are only trained on MQM ratings from 2021 and 2022. We always train models on raw MQM ratings, i.e., using the 0–25 scale.

## 2.2 Evaluation Data

**MQM.** Our primary evaluation set consists of the WMT23 MQM ratings, which includes three language pairs: en-de, he-en and zh-en. Since the zh-en language pair is known to have low-quality references (Kocmi et al., 2023), we replace them with newly collected references. Note that this has no effect on the MQM ratings, as those were collected in a source-based fashion. Additionally, given the fact that one of the official WMT24 test language pairs is ja-zh, we reserve the news domain subset of the WMT22 en-zh ratings for evaluation, allowing us to assess our models’ performance on a language pair with Chinese as the target language.

**DA.** We use the WMT23 DA ratings as a secondary evaluation set, taking advantage of its better language coverage (8 language pairs). Nevertheless, with DA ratings generally following a significantly different distribution than MQM ratings, a higher correlation of the metric scores with these DA ratings does not necessarily imply better performance. For example, fine-tuning a model on zh-en MQM ratings results in lower performance than fine-tuning it on DA ratings, according to the zh-en DA evaluation set (but not the MQM one). Therefore, we only consider the WMT23 DA evaluation set in experiments where we mix MQM and DA training data together.

## 2.3 Synthetic Data

After seeing the initial benefits from the simple synthetic data used for training MetricX-23, we decided to construct a more comprehensive collection

<sup>3</sup>The en-zh MQM ratings, available at <https://github.com/google/wmt-mqm-human-evaluation>, were collected post-WMT22.

of synthetic training examples. They cover additional, less trivial failure modes of MetricX, i.e., translation issues commonly unrecognized by the metric. The DEMETR challenge set (Karpinska et al., 2022), which we relied on last year, does not cover several of the failure modes we created the new synthetic training examples for, hence we also constructed a set of test examples for each of them. Next, we describe how we designed both the training and the test synthetic datasets.

### 2.3.1 Training Sets

In order for the MetricX models to learn to identify certain types of bad translations that are not sufficiently (or at all) represented in the regular WMT training data, we generated synthetic examples that we augment the training data with. They were created by modifying examples from the DA datasets ranging from WMT15 to WMT22, comprising 49 language pairs. Table 1 provides an overview of the various failure modes that we considered, including brief descriptions of how we prepared the synthetic data to address them. Additional details regarding the creation process can be found in Appendix A.

### 2.3.2 Test Set

We constructed a new DEMETR-style test set based on the WMT23 DA dataset, with examples generated analogously to our synthetic training examples, as described in Table 1. Each synthetic example is paired with its original counterpart (although using the reference instead of the candidate translation whenever the synthetic translation was created from the reference), which allows for a metric to be evaluated on how frequently it ranks the pairs correctly.

## 3 Metric Descriptions

The MetricX-24 submissions to the WMT24 Metrics Shared Task build on top of the successful MetricX-23 (Juraska et al., 2023; Kocmi et al., 2023), with several major improvements. We start this section by summarizing the aspects this year’s submissions have in common with MetricX-23, then provide an overview of the modifications, and finally describe the differences between the individual submissions.

### 3.1 MetricX Model

MetricX is a learned metric, powered by a regression model trained to predict a floating point number that represents the quality of a given transla-

tion. The reference-based variant takes the candidate translation (hypothesis) and reference segments as input, and concatenates them, along with corresponding prefixes (“candidate:” and “reference:”, respectively). In contrast to the previous versions, MetricX-24 also prepends the source segment (along with the prefix “source:”) to the input, offering the model additional context to make a better prediction in the reference-based setting, which may be beneficial especially in cases where the reference is inadequate. The model then encodes this combined input and uses it to predict the translation quality score. The QE variant works in an analogous way, but taking only the source segment and the hypothesis as the input.

With MetricX-24, we continue to rely on mT5 (Xue et al., 2021) as the pretrained language model that we fine-tune on translation evaluation data. We refer the reader to Juraska et al. (2023) for details on how we adapted this encoder-decoder model to the regression task. Similar to MetricX-23, we fine-tune the model in two stages: first on DA ratings ( $z$ -normalized, aggregated per segment, negated, and finally clipped to the  $[-1.0, 1.0]$  range) and then further on raw MQM ratings. As a result, the metric produces scores in the  $[0, 25]$  range. The model is trained with a mean squared error (MSE) loss function. Further implementation details can be found in §4.

## 3.2 Design Improvements

We achieve some initial improvement in performance by simply including the WMT22 data in the training set – both the DA and the MQM ratings, which we previously used as the evaluation set when developing MetricX-23. The additional MQM ratings (including en-ru, a language pair not present in the older MQM data) are especially valuable, considering the scarcity of MQM data. Besides that, we introduce three major modifications to the training procedure and data in order to further improve MetricX’s performance, described throughout the rest of this section.

### 3.2.1 Training With Synthetic Data

Although we used synthetic training data alongside the DA and MQM ratings already for training MetricX-23, the synthetic examples covered only the two trivial cases of empty and reference-matching translations. As described in §2.3, we prepared a significantly more comprehensive synthetic training set for MetricX-24, which we combine

Failure mode	Synthetic candidate translation description	MQM score
Empty translation	Empty string.	25
Gibberish	Text of a similar length as the reference, generated by sampling words from the vocabulary built from all references in the data with a matching target language.	25
Fluent but unrelated translation	Arbitrary reference from the dataset of a similar length and in the same language.	25
Undertranslation	Candidate translation with an arbitrary sentence removed, if a multi-sentence segment, otherwise, candidate translation with 20–80% words removed from the end.	5–25
Duplication	Candidate translation duplicated, with a space in between.	25
Missing punctuation	Reference translation with the end punctuation removed (11 punctuation symbols considered, such as period, question mark, closing parenthesis or quotation mark).	1
Reference-matching translation	Reference translation itself (unlike the rest, these synthetic examples are meant to train the metric to predict a perfect score for translations matching the reference).	0

Table 1: Failure mode categories we prepared synthetic data for, along with brief descriptions of how we created the synthetic examples from the WMT data, and the MQM scores we label the training examples with.

with the DA and MQM data in both fine-tuning stages. We experimented with various ratios, and settled on 1:100 for each synthetic example category in the first stage and 1:5000 in the second stage. We evaluate the effects of adding the synthetic training data by measuring accuracy and average score differences on the synthetic test set, also described in §2.3.

### 3.2.2 Mixing DA and MQM Data

Next, we attempt to address the inevitable decline in MetricX performance on other languages after fine-tuning the model on MQM data, which only covers a few language pairs. The performance, as measured by the WMT23 DA evaluation set with 8 language pairs, quickly declines after starting to fine-tune on MQM ratings. While it is expected that the change in the general score distribution – caused by the switch from DA to MQM ratings – results in the Pearson correlations with the ground-truth scores dropping, we believe the model should be able to retain its system- and segment-level pairwise accuracy from the first stage of fine-tuning on DA data. Moreover, we observe a significant drop in system-level performance on the zh-en language pair of the MQM evaluation set, despite zh-en being present in the MQM training data.

In order to remedy these behaviors, we mix in a smaller proportion of DA ratings in the second-stage fine-tuning. That way the model is trained primarily on MQM ratings, but has a continued exposure to the additional 40+ language pairs from the first stage of fine-tuning. We experimented with different combinations of DA and MQM rating formats (e.g., raw vs.  $z$ -normalized, transformed to the MQM scale or not, etc.), and the one yielding

the best results was raw MQM ratings combined with raw DA ratings linearly transformed to the MQM scale of  $[0, 25]$ . Finally, we determined that a DA:MQM ratio of 1:4 works well for boosting the performance on the DA evaluation set back to the levels from the first stage of fine-tuning, without a significant negative impact on the model’s performance on the MQM evaluation set.<sup>4</sup>

### 3.2.3 Hybrid Input Mode

The third major modification we make to the training procedure when developing MetricX-24, is mixing training examples in three different formats: (1) source + hypothesis, (2) hypothesis + reference, and (3) source + hypothesis + reference. This allows the model to operate in both a QE and a reference-based mode (and the latter either with or without the source included). But perhaps more importantly, it gives the model an opportunity to learn how much weight to put on the source and the reference in different scenarios, or possibly to completely ignore the reference when it is of low quality. Such a hybrid model is then evaluated as a QE model by only passing it the source segment and the hypothesis as input, and as a reference-based model by additionally passing it the reference.

### 3.3 MetricX-24 Variants

There are four variants of MetricX-24 that we submitted to the WMT24 Metrics Shared Task:

- MetricX-24-Hybrid (primary)
- MetricX-24-Hybrid-QE (primary)
- MetricX-24

<sup>4</sup>After a more extensive post-submission experimentation, we determined the optimal ratio to be 1:10.

- MetricX-24-QE

Our primary reference-based and QE submissions are actually the same hybrid model, with the scores predicted with and without the references provided as part of the input. The secondary submissions are the standalone reference-based and QE counterparts of the hybrid model, i.e., only trained on examples with the references (as well as the source segments) included and on examples with the references omitted, respectively. Other than that, all of the submission models are identical in terms of training data mixtures, as described in §3.2.1 and §3.2.2, as well as training hyperparameters.

## 4 Experimental Setup

### 4.1 Meta-Evaluation

As mentioned in §2.2, our primary evaluation set consists of the MQM ratings from WMT23, as well as the news domain subset of the en-zh language pair from WMT22. Considering there is no into-English language pair among the official test sets this year, we focus primarily on en-de and en-zh when evaluating our models, but also keeping zh-en (the dataset with alternate references) in the mix, in order to ensure that we do not overfit the models to out-of-English language pairs. To evaluate our models, we calculate the agreements between their predicted scores and the human judgments of translation quality using the four different correlations from the WMT23 Metrics Shared Task (Freitag et al., 2023), detailed in Appendix B.

### 4.2 Checkpoint Selection

In both the first and the second stage of fine-tuning, we pick the best checkpoint  $c_{\text{best}}$  based on the following linear combination of segment- and system-level pairwise accuracy:

$$\arg \max_c \frac{3}{4} \sum_l \text{acc}_l^{\text{seg}}(c) + \frac{1}{4} \sum_l \text{acc}_l^{\text{sys}}(c),$$

where  $l \in \{\text{en-de}, \text{en-zh}, \text{zh-en}\}$ , and  $\text{acc}_l^{\text{seg}}(c)$  and  $\text{acc}_l^{\text{sys}}(c)$  are the segment- and the system-level pairwise accuracy calculated for checkpoint  $c$  on the language pair  $l$  of the evaluation set. We down-weight the system-level component to account for its greater variance and to thus avoid a checkpoint being picked due to a rare spike in system-level accuracy if segment-level accuracy is low.

### 4.3 Implementation Details

MetricX-24, similar to its predecessor, is implemented with TensorFlow (Abadi et al., 2015) and the T5X library (Roberts et al., 2023). All of the metric variants are based on mT5-XXL with 13B parameters. We defer further implementation details to Appendix C. We are publicly releasing our submissions, converted from TensorFlow to PyTorch (Paszke et al., 2019) checkpoints.<sup>5</sup>

## 5 Results and Discussion

Here we present the results of our experiments, focusing solely on fully trained models (i.e., those that went through both stages of fine-tuning) and modifications in the second stage. Since the ablation studies performed with reference-based and QE models show similar trends, we discuss the reference-based experiments in depth in this section, and provide the QE results in Appendix D.2. Due to limited resource availability, we were only able to run each experiment with one random seed.

### 5.1 Training With Synthetic Data

We start by examining the benefits of including synthetic training examples, as described in 2.3. In Table 2, the bottom four rows – corresponding to the hybrid model – demonstrate the effects of progressively adding DA data only, synthetic data only, and finally both, to the training set in the second stage of fine-tuning.<sup>6</sup> We ended up not using the duplication synthetic training set, as we observed that the models learn to correctly identify such cases even without it.

The first thing to notice is that mixing in DA ratings actually improves the metric’s performance on the synthetic test set over fine-tuning on MQM ratings alone, especially in the unrelated, undertranslation and duplication failure modes. Adding synthetic data instead is, however, significantly more effective in general, boosting the accuracy to the 94–100% range in most categories. Finally, augmenting the training set with both the DA and the synthetic data results in an overall similar performance as with the synthetic data only.

Missing punctuation is one of two categories in which our metrics score not so close to perfect. In fact, the synthetic training examples appear not to be helpful in improving the performance at all. Our

<sup>5</sup><https://github.com/google-research/metricx>

<sup>6</sup>The models that did not include synthetic training data in the second stage, did not use it in the first stage either.

MetricX variant	+DA	+Synth	Empty transl.	Gibberish	Unrelated	Undertransl.	Duplication	Missing punct.	Ref-match
23	–	~	<b>100.00</b>	<b>100.00</b>	88.14	57.75	38.14	66.01	<b>94.00</b>
24	✓	✓	99.29	99.86	<b>99.29</b>	<b>98.75</b>	99.14	83.01	78.14
24-Hybrid	–	–	51.43	99.86	81.00	68.75	87.57	83.66	76.00
	✓	–	53.57	99.71	92.14	82.25	<b>99.57</b>	<b>85.62</b>	72.86
	–	✓	94.14	99.71	99.14	96.25	94.43	84.97	79.86
	✓	✓	97.29	99.71	98.71	96.25	99.43	82.35	75.14

Table 2: Accuracy of reference-based MetricX variants in all 7 categories of our synthetic test set. “23” is the baseline, the last row of “24-Hybrid” corresponds to our primary submission, and “24” is our secondary submission.

MetricX variant	+DA	+Synth	Segment-level pairwise accuracy				System-level pairwise accuracy			
			en-de	zh-en	zh-en <sup>†</sup>	en-zh	en-de	zh-en	zh-en <sup>†</sup>	en-zh
23	–	~	60.20	53.12	54.06	55.73	90.91	89.52	86.67	74.36
24	✓	✓	60.71	54.50	55.78	56.16	96.97	<b>92.38</b>	<b>95.00</b>	<b>88.46</b>
24-Hybrid	–	–	61.17	54.63	55.52	57.43	<b>100.00</b>	89.52	91.67	85.90
	✓	–	60.75	54.89	55.58	57.65	98.48	<b>92.38</b>	92.50	84.62
	–	✓	<b>61.75</b>	54.38	55.43	<b>57.73</b>	98.48	90.48	91.67	<b>88.46</b>
	✓	✓	61.11	<b>55.00</b>	<b>55.82</b>	57.02	98.48	<b>92.38</b>	94.17	85.90

Table 3: Meta-evaluation scores of reference-based MetricX variants on the WMT23 MQM evaluation set. “23” is the baseline, the last row of “24-Hybrid” corresponds to our primary submission, and “24” is our secondary submission. <sup>†</sup>Alternate references.

hypothesis is that using references to create this category of synthetic examples results in a significant proportion of misleading examples because we assume references to be perfect, but that is not always the case. That, combined with the fact that the removal of the punctuation symbol from the end of the segment warrants just a minor score change, means that some of the synthetic examples might have an unreasonably high ground-truth score associated with them, thus giving the model the opposite signal to what is desired.

The reference-matching translation synthetic training set appears not to be effective either, however, its benefits are somewhat concealed by the fact that mixing in DA data drags the performance in this category down. With the non-hybrid model, we observed a significantly bigger drop with DA data included (77% → 64%) and a greater increase with synthetic data included instead (77% → 83%). Granted, that is still far from perfect, however, expecting a 100% accuracy in this category equates to expecting that the candidate translation is never better than the reference, which, as we pointed out earlier, is not always true when judging the translation quality based on the source segment.

Overall, thanks to the new synthetic training data, MetricX-24 (hybrid or not) is clearly more robust to the failure modes than MetricX-23 (see first row

in the table), with the reference-matching translations being an exception. That might have to do with the absence of WMT22 data in the training set of MetricX-23, or the only synthetic examples present therein being those of empty and reference-matching translations.

## 5.2 Mixing DA and MQM Data

We already discussed the effects of adding DA data to the training set in the second stage of fine-tuning in terms of the synthetic test set performance; let us now have a look at the correlations with human MQM scores. Comparing the first two rows of the “24-Hybrid” section in Table 3, we see that there are just relatively minor changes in either direction across all language pairs, with score differences within the expected variance between runs.

What the table does not show, however, is the huge jump in all correlations across all language pairs on the WMT23 DA evaluation set, typically back to the levels from the first stage of fine-tuning on DA data only, or above. Segment- and system-level pairwise accuracy increases by up to 2 and 5 points, respectively, and Pearson’s  $r$  sees improvements of up to 10 points. These are valuable gains, considering we achieved them without sacrificing the performance on the MQM evaluation set. An overview of the results and a more detailed

analysis on the DA evaluation set can be found in Appendix D.1.

### 5.3 Hybrid Input Mode

To wrap up the evaluation, we discuss the performance difference between MetricX-24 and MetricX-24-Hybrid (rows 2 and 6 in Table 3). At the system level, the hybrid variant lags slightly behind in zh-en and en-zh, but it makes up for it by outperforming the non-hybrid across the board at the segment level. Notably, the hybrid metric achieves an almost 1% higher segment-level accuracy on en-zh, and the 0.5% boost on zh-en (with original references) may be evidence for the hybrid model handling examples with poor-quality references better, especially considering the accuracy difference on the zh-en set with alternate references is only 0.04%. The other performance differences between the two models are largely insignificant.

Finally, comparing our primary submission with MetricX-23 (row 1 in the table), we can see consistent gains of 1–2 points in segment-level accuracy, and substantially bigger gains at the system level, with the accuracy on en-zh improving by a whopping 11.5 points. We conclude that this a significant improvement over our last year’s submission, ranked overall second in the WMT23 Metrics Shared Task.

## 6 Related Work

Traditionally, evaluation metrics predict a scalar quality score for the translation. This type of metric includes BLEU, ChrF, MetricX (Juraska et al., 2023), BLEURT (Sellam et al., 2020; Pu et al., 2021), COMET (Rei et al., 2020, 2022a), COMETKiwi (Rei et al., 2022b), Prism (Thompson and Post, 2020), and more. While these metrics have historically been the dominant category of metric, newly proposed methods provide structured (Perrella et al., 2022; Fernandes et al., 2023; Kocmi and Federmann, 2023; Guerreiro et al., 2023) or natural language explanations (Xu et al., 2023) for the predicted scores.

Then, evaluation metrics are considered to be reference-based or reference-free (also known as “quality estimation”) depending on whether or not they require a reference to evaluate a translation. Metric developers usually train separate models for each type of metric (e.g., COMET and COMETKiwi, or MetricX-23 and MetricX-23-QE), but some opt for combining both tasks into a single

model (Wan et al., 2022; Guerreiro et al., 2023), which is the approach we took in this work with our hybrid model.

Finally, while most metrics like MetricX-24 use a dedicated model for scoring translations, some recent works have begun to leverage general-purpose large language models instead (Fernandes et al., 2023; Kocmi and Federmann, 2023; Xu et al., 2023; Leiter et al., 2023; Leiter and Eger, 2024). While LLM-based metrics have achieved strong system-level performance, using a learned dedicated model was the best approach at the segment-level in last year’s Metrics Shared Task (Freitag et al., 2023).

## 7 Conclusion

We presented in detail our approach to training MetricX-24, a regression-based MT evaluation metric. We submitted four versions of MetricX-24 to the WMT24 Metrics Shared Task, including a reference-based and a QE variant, as well as a new hybrid variant evaluated with and without the references. By evaluating on the WMT23 MQM dataset, we showed all of them to significantly outperform our last year’s submission, MetricX-23. In addition, we made MetricX-24 more robust to various types of bad translations, which do not frequently occur in the WMT data, such as undertranslation, or fluent but unrelated translation. Finally, by combining DA and MQM ratings together in the final stage of fine-tuning, we were able to dramatically increase the performance on the WMT23 DA dataset covering 8 language pairs, while maintaining the high correlations with the MQM ratings at the same time.

## References

- Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. 2015. [TensorFlow: Large-scale machine learning on heterogeneous systems](#). Software available from tensorflow.org.
- Daniel Deutsch, George Foster, and Markus Freitag. 2023. [Ties matter: Meta-evaluating modern metrics](#)

- with pairwise accuracy and tie calibration. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12914–12929, Singapore. Association for Computational Linguistics.
- Patrick Fernandes, Daniel Deutsch, Mara Finkelstein, Parker Riley, André Martins, Graham Neubig, Ankush Garg, Jonathan Clark, Markus Freitag, and Orhan Firat. 2023. [The devil is in the errors: Leveraging large language models for fine-grained machine translation evaluation](#). In *Proceedings of the Eighth Conference on Machine Translation*, pages 1066–1083, Singapore. Association for Computational Linguistics.
- Markus Freitag, George Foster, David Grangier, Viresh Ratnakar, Qijun Tan, and Wolfgang Macherey. 2021. [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, Nitika Mathur, Chi-kiu Lo, Eleftherios Avramidis, Ricardo Rei, Brian Thompson, Tom Kocmi, Frederic Blain, Daniel Deutsch, Craig Stewart, Chrysoula Zerva, Sheila Castilho, Alon Lavie, and George Foster. 2023. [Results of WMT23 metrics shared task: Metrics might be guilty but references are not innocent](#). In *Proceedings of the Eighth Conference on Machine Translation*, pages 578–628, Singapore. Association for Computational Linguistics.
- Nuno M Guerreiro, Ricardo Rei, Daan van Stigt, Luisa Coheur, Pierre Colombo, and André FT Martins. 2023. [XCOMET: Transparent machine translation evaluation through fine-grained error detection](#). *arXiv preprint arXiv:2310.10482*.
- Juraj Juraska, Mara Finkelstein, Daniel Deutsch, Aditya Siddhant, Mehdi Mirzazadeh, and Markus Freitag. 2023. [MetricX-23: The Google submission to the WMT 2023 metrics shared task](#). In *Proceedings of the Eighth Conference on Machine Translation*, pages 756–767, Singapore. Association for Computational Linguistics.
- 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, Eleftherios Avramidis, Rachel Bawden, Ondřej Bojar, Anton Dvorkovich, Christian Federmann, Mark Fishel, Markus Freitag, Thamme Gowda, Roman Grundkiewicz, Barry Haddow, Philipp Koehn, Benjamin Marie, Christof Monz, Makoto Morishita, Kenton Murray, Makoto Nagata, Toshiaki Nakazawa, Martin Popel, Maja Popović, and Mariya Shmatova. 2023. [Findings of the 2023 conference on machine translation \(WMT23\): LLMs are here but not quite there yet](#). In *Proceedings of the Eighth Conference on Machine Translation*, pages 1–42, Singapore. Association for Computational Linguistics.
- Tom Kocmi, Rachel Bawden, Ondřej Bojar, Anton Dvorkovich, Christian Federmann, Mark Fishel, Thamme Gowda, Yvette Graham, Roman Grundkiewicz, Barry Haddow, Rebecca Knowles, Philipp Koehn, Christof Monz, Makoto Morishita, Masaaki Nagata, Toshiaki Nakazawa, Michal Novák, Martin Popel, and Maja Popović. 2022. [Findings of the 2022 conference on machine translation \(WMT22\)](#). In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 1–45, Abu Dhabi, United Arab Emirates (Hybrid). 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 and Steffen Eger. 2024. [PrExMe! Large Scale Prompt Exploration of Open Source LLMs for Machine Translation and Summarization Evaluation](#). *arXiv preprint arXiv:2406.18528*.
- Christoph Leiter, Juri Opitz, Daniel Deutsch, Yang Gao, Rotem Dror, and Steffen Eger. 2023. [The Eval4NLP 2023 Shared Task on Prompting Large Language Models as Explainable Metrics](#). In *Proceedings of the 4th Workshop on Evaluation and Comparison of NLP Systems*, pages 117–138, Bali, Indonesia. Association for Computational Linguistics.
- Arle Lommel, Hans Uszkoreit, and Aljoscha Burchardt. 2014. [Multidimensional quality metrics \(mqm\): A framework for declaring and describing translation quality metrics](#). *Tradumática*, (12):0455–463.
- 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.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor



- Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. [Pytorch: An imperative style, high-performance deep learning library](#).
- 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.
- 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, 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.
- Adam Roberts, Hyung Won Chung, Gaurav Mishra, Anselm Levskaya, James Bradbury, Daniel Andor, Sharan Narang, Brian Lester, Colin Gaffney, Afroz Mohiuddin, et al. 2023. Scaling up models and data with t5x and seqio. *Journal of Machine Learning Research*, 24(377):1–8.
- 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.
- Noam Shazeer and Mitchell Stern. 2018. Adafactor: Adaptive learning rates with sublinear memory cost. In *International Conference on Machine Learning*, pages 4596–4604. PMLR.
- Brian Thompson and Matt Post. 2020. [Automatic Machine Translation Evaluation in Many Languages via Zero-Shot Paraphrasing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 90–121, Online. Association for Computational Linguistics.
- Yu Wan, Keqin Bao, Dayiheng Liu, Baosong Yang, Derek F. Wong, Lidia S. Chao, Wenqiang Lei, and Jun Xie. 2022. [Alibaba-translate China’s submission for WMT2022 metrics shared task](#). In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 586–592, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Wenda Xu, Danqing Wang, Liangming Pan, Zhenqiao Song, Markus Freitag, William Wang, and Lei Li. 2023. [INSTRUCTSCORE: Towards Explainable Text Generation Evaluation with Automatic Feedback](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 5967–5994, Singapore. Association for Computational Linguistics.
- 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.

## A Synthetic Data Creation

We sample 500 examples from each language pair, whose candidate translations (hypotheses) we then manipulate in different ways to create the synthetic examples for each failure mode category. The missing punctuation category is an exception, with a stratified sample across the 11 end-punctuation symbols, rather than language pairs, and 250 examples each.

In general, the synthetic examples have the candidate translation manipulated, turning it into a worse, or an outright bad, translation. One exception is the reference-matching category, whose purpose is to actually teach the metric to score translations that match the reference highly, which it does not learn to do reliably when only trained on the WMT data. Table 4 shows a few concrete examples from the synthetic training set.

## B Meta-Evaluation Details

**System-Level.** At the system level, we measure pairwise ranking accuracy (Kocmi et al., 2021), as well as Pearson’s  $r$ . Pairwise accuracy assesses how well a metric ranks MT systems by calculating the proportion of all possible pairs of systems that are ranked the same by the metric and human scores. Pearson’s  $r$ , on the other hand, captures how strong the linear relationship is between the metric and human scores for MT systems. We obtain the system-level scores (both metric and human) as the mean segment-level score for each system.

**Segment-Level.** At the segment level, we use the group-by-item pairwise accuracy with tie calibration, as described by Deutsch et al. (2023), and the no-grouping Pearson’s  $r$ . The pairwise accuracy calculates the proportion of all possible pairs of translations for the same source segment that are ranked the same by the metric and human, then averages the accuracies over all input segments. At the same time, it rewards correct tie predictions by introducing ties for any two translations with a metric score difference below an automatically determined threshold. The no-grouping Pearson’s  $r$  quantifies the linear relationship between the metric and human scores across all translations from every system and document.

## C Implementation Details

Having increased the maximum segment length from 256 to 512 SPM tokens, and including up to three segments (source, hypothesis and reference) in the model’s input, each training run requires 256 TPUs. Using a batch size of 256, we train our models for 16K steps in the first stage, using a learning rate of 0.001 with an inverse square root decay after the first 2K steps. We then fine-tune the best checkpoint for another 8K steps in the second stage, lowering the learning rate to 0.0002 and decaying it after 1K steps. The models are trained using the Adafactor optimizer (Shazeer and Stern, 2018).

## D Additional Results

### D.1 Mixing DA and MQM Data

Table 5 compiles the results of the meta-evaluation of a group of reference-based models on the WMT23 DA evaluation set. All of the models are standalone reference-based models. In the table, we contrast four variants of the model fine-tuned in two stages (DA then MQM data) with a model fine-tuned on DA data only (i.e., the first stage only). We present the results on a subset of four language pairs, two of which are present in our MQM training data (en-de and zh-en) and two which are not (en-cs and de-en).

The experiments with mixing DA and MQM data in the second stage of fine-tuning were motivated by the large differences in performance on the WMT23 DA evaluation set observed between a model trained on DA ratings only (row 1 in Table 5) and the same model further fine-tuned on MQM ratings (row 2). As already discussed in §2.2, this can be partly explained by the discrepancy in DA and MQM rating distributions. This discrepancy understandably affects Pearson correlations, however, it should not have a significant effect on how the metric ranks segments or systems. Nevertheless, while we observed large drops in Pearson’s  $r$ , the pairwise accuracy also dropped substantially for most of the language pairs, both at the segment and the system level. For example, on en-cs the segment-level accuracy drops from 59.54 to 57.43, and the system-level accuracy from 87.62 to 82.86.

Considering the fact that the performance difference between the models in rows 1 and 2 on en-de and zh-en (i.e., the language pairs with a good amount of MQM training data), are relatively

---

**Gibberish (zh-en example)**

*Created from: corpus hypothesis vocabulary*

src 我希望你们能准时，不是想要你们的优惠券！！  
hyp filter two that to also in allegations train 800 city, continuous the  
ref I hope you can be on time, and it's not that I want your coupons! !  
label 25

---

**Fluent but unrelated translation (de-en example)**

*Created from: corpus references*

src Damit können doppelt so viele Studierende ausgebildet werden wie bisher.  
hyp She booked a return flight and went home the next day.  
ref In that way, twice as many students can be educated as before.  
label 25

---

**Undertranslation (cs-en example)**

*Created from: hypothesis*

src Dlouhodobě napjaté vztahy mezi oběma zeměmi se vyostřily v roce 2018 poté, co Washington odstoupil od jaderné dohody z roku 2015 mezi Íránem a světovými mocnostmi a zavedl vůči Íránu sankce, které mají tvrdý dopad na jeho ekonomiku.  
hyp Long-tense relations between the two countries sharpened in 2018 after Washington withdrew from the 2015 nuclear deal between Iran and world powers and imposed sanctions.  
ref Long-term tense relations between both countries escalated in 2018 after that Washington withdrew from the nuclear deal closed in 2015 between Iran and the world powers and imposed sanctions against Iran, which have had hard impacts on its economy.  
label 12.75

---

**Duplication (fi-en example)**

*Created from: hypothesis*

src Ensi vuoden vaje on yli 2,4 prosenttia kansantuotteesta.  
hyp Next year's deficit will be over 2.4 per cent of national product. Next year's deficit will be over 2.4 per cent of national product.  
ref Next year's deficit is over 2.4 per cent of GDP.  
label 15

---

**Missing punctuation (ru-en example)**

*Created from: reference*

src Последний альбом Асе вышел в 2016 году.  
hyp Their last album, "Ace", came out in 2016  
ref Their last album, "Ace", came out in 2016.  
label 1

---

**Reference-matching translation (ja-en example)**

*Created from: reference*

src グレタさんは、27日の金曜日にも行うことを呼びかけていた。  
hyp Now, Greta is calling for further strikes to be held on Friday the 27th.  
ref Now, Greta is calling for further strikes to be held on Friday the 27th.  
label 0

---

Table 4: Synthetic examples for the different failure mode categories (except for the trivial empty translation case), along with the MQM scores we label the training examples with. Each category also has an indication of how the hypothesis was created/generated in order to produce a synthetic example (e.g., by modifying the original hypothesis or reference).

MetricX variant	+DA	+Synth	Segment-level pairwise accuracy				System-level pairwise accuracy			
			en-de	zh-en	en-cs	de-en	en-de	zh-en	en-cs	de-en
DA only	N/A	N/A	61.77	56.33	59.54	61.14	<b>95.45</b>	79.05	<b>87.62</b>	92.31
DA then MQM	–	–	61.59	55.99	57.43	61.65	93.94	81.90	82.86	85.90
	✓	–	61.88	<b>56.67</b>	<b>60.16</b>	62.29	<b>95.45</b>	80.95	86.67	88.46
	–	✓	<b>62.60</b>	56.35	59.02	61.92	<b>95.45</b>	<b>84.76</b>	81.90	<b>93.59</b>
	✓	✓	61.89	56.64	60.04	<b>62.32</b>	93.94	83.81	86.67	<b>93.59</b>

Table 5: Meta-evaluation scores of reference-based MetricX variants on a subset of the language pairs of the WMT23 DA evaluation set. “DA only” is a model after just the first stage of fine-tuning (i.e., on DA data only), whereas the “DA then MQM” section contains models fine-tuned in full two stages. The last row thus corresponds to the “24” row in Tables 2 and 3, i.e., our secondary submission “MetricX-24”.

MetricX variant	+DA	+Synth	Segment-level Pearson’s $r$				System-level Pearson’s $r$			
			en-de	zh-en	en-cs	de-en	en-de	zh-en	en-cs	de-en
DA only	N/A	N/A	<b>60.10</b>	<b>41.52</b>	<b>43.47</b>	52.59	98.48	<b>89.21</b>	<b>92.49</b>	97.20
DA then MQM	–	–	48.18	34.66	39.77	44.09	93.41	87.58	90.60	87.40
	✓	–	53.67	36.29	43.11	52.79	93.67	87.75	90.40	91.53
	–	✓	56.03	35.21	37.24	45.26	<b>99.48</b>	88.40	89.32	96.79
	✓	✓	57.92	37.17	42.55	<b>55.23</b>	98.56	88.94	91.58	<b>97.97</b>

Table 6: Same as Table 5, but showing Pearson correlations instead of pairwise accuracies.

small, we conjecture that further fine-tuning on MQM data alone causes the model to partially “forget” other languages from the first stage of fine-tuning. We attempt to prevent the model from this sort of forgetting by mixing some DA ratings into the training set in the second stage.

As the scores of the model in row 3 in the table demonstrate, we are able, for the most part, to restore the performance observed in the first stage of fine-tuning by adding a small proportion of DA training data in the second stage too. Adding not only the DA data, but also the synthetic data, in the second stage (row 5) sometimes boosts the performance further, significantly improving even over the first-stage performance (row 1). Most importantly, the gains over fine-tuning on MQM data alone (row 2) are achieved not at the expense of the model’s performance on the MQM or the synthetic test set, as evidenced by the results in Tables 2 and 3.

Finally, Table 6 shows the expected big drops in Pearson correlation with the DA ratings after fine-tuning on MQM data (see rows 1 and 2), especially at the segment level. Adding DA data in the second stage helps recover most of the performance (compare rows 3 and 5 with row 1), but as expected, the correlations remain lower particularly for language pairs present in the MQM data the model is fine-tuned on in the second stage (en-de and zh-en).

## D.2 QE Models

In Tables 7 and 8, we present the meta-evaluation results for our QE models. These are analogous to those presented in §5, only the hybrid model is evaluated in a reference-free mode, and the non-hybrid models are ones trained on the source and hypothesis segments only. Note that the hybrid model is the same checkpoint as the one for which we reported the reference-based results in Tables 2 and 3, i.e., not one optimized for QE performance.

Examining first the results on the synthetic test set, summarized in Table 7, we see similar trends to those observed with reference-based models (Table 2). The main difference is that the QE models achieve significantly lower performance in the missing punctuation and the reference-matching translation categories. This, however, is expected because both the types of synthetic examples were created from references. In case of the missing punctuation examples, the synthetic translation is simply the reference with the end punctuation removed. Comparing such a hypothesis with the corresponding reference is arguably a significantly easier task than comparing it to the source segment and identifying a missing punctuation symbol. Moreover, there may be a mismatch in the presence of punctuation between the source and the reference in the training examples, making it even more difficult for a QE model to reliably identify missing punctuation. As for the reference-matching

MetricX variant	+DA	+Synth	Empty transl.	Gibberish	Unrelated	Under-transl.	Duplication	Missing punct.	Ref-match
23	–	~	<b>100.00</b>	<b>99.86</b>	96.43	63.25	88.29	69.93	63.00
24	✓	✓	97.86	<b>99.86</b>	<b>99.43</b>	<b>98.50</b>	<b>98.14</b>	65.36	63.43
24-Hybrid	–	–	69.86	<b>99.86</b>	82.43	81.25	63.00	<b>77.78</b>	63.00
	✓	–	66.14	99.57	95.29	93.50	97.86	73.86	62.57
	–	✓	93.57	99.71	99.29	96.50	84.43	69.28	62.14
	✓	✓	93.71	<b>99.86</b>	<b>99.43</b>	97.25	<b>98.14</b>	69.28	<b>64.14</b>

Table 7: Accuracy of reference-free (QE) MetricX variants in all 7 categories of our synthetic test set. “23” is the baseline, the last row of “24-Hybrid” corresponds to our primary submission, and “24” is our secondary submission. The hybrid model is the same as in Table 2, only evaluated without references provided as input.

MetricX variant	+DA	+Synth	Segment-level pairwise accuracy				System-level pairwise accuracy			
			en-de	zh-en	zh-en <sup>†</sup>	en-zh	en-de	zh-en	zh-en <sup>†</sup>	en-zh
23	–	~	59.57	52.64	52.89	54.47	92.42	86.67	85.83	74.36
24	✓	✓	59.70	<b>54.30</b>	<b>54.48</b>	56.00	98.48	<b>92.38</b>	90.83	<b>87.18</b>
24-Hybrid	–	–	60.11	53.80	54.00	<b>56.27</b>	<b>100.00</b>	89.52	89.17	84.62
	✓	–	59.18	54.08	54.30	56.14	<b>100.00</b>	<b>92.38</b>	90.00	84.62
	–	✓	<b>60.27</b>	53.76	53.99	55.88	98.48	89.52	90.00	83.33
	✓	✓	59.52	54.15	54.41	55.94	98.48	90.48	<b>91.67</b>	83.33

Table 8: Meta-evaluation scores of reference-free (QE) MetricX variants on the WMT23 MQM evaluation set. “23” is the baseline, the last row of “24-Hybrid” corresponds to our primary submission, and “24” is our secondary submission. The hybrid model is the same as in Table 3, only evaluated without references provided as input. <sup>†</sup>Alternate references.

translation category, a QE model does not have access to the reference, so it makes perfect sense for it to score a candidate translation better than the reference translation if the reference is of low quality.

Switching over to Table 8, which shows the pairwise accuracy of the QE model scores, the trends are also in line with those of the reference-based models in Table 3. In contrast to the reference-based results, however, the hybrid model (row 6) does not outperform the standalone model (row 2), although most of the differences are within the expected variance. An astute reader might notice that the accuracy scores on the zh-en test set with the original references and the one with the alternate references do not match (despite the QE models not using the references), and that is because the latter has the original references included as an additional “human system”.

Finally, we note that our QE models do not fall far behind their reference-based counterparts. In fact, both our primary and secondary QE submissions of MetricX-24 outperform our reference-based MetricX-23 submission from last year, according to the WMT23 MQM evaluation set.