

Mitigating Metric Bias in Minimum Bayes Risk Decoding

Geza Kovacs and Dan Deutsch and Markus Freitag

Google

{geza, danddeutsch, freitag}@google.com

Abstract

While Minimum Bayes Risk (MBR) decoding using metrics such as COMET or MetricX has outperformed traditional decoding methods such as greedy or beam search, it introduces a challenge we refer to as *metric bias*. As MBR decoding aims to produce translations that score highly according to a specific utility metric, this very process makes it impossible to use the same metric for both decoding and evaluation, as improvements might simply be due to *reward hacking* rather than reflecting real quality improvements. In this work we find that compared to human ratings, neural metrics not only overestimate the quality of MBR decoding when the same metric is used as the utility metric, but they also overestimate the quality of MBR/QE decoding with other neural utility metrics as well. We also show that the metric bias issue can be mitigated by using an ensemble of utility metrics during MBR decoding: human evaluations show that MBR decoding using an ensemble of utility metrics outperforms a single utility metric.

1 Introduction

Minimum bayes risk (MBR) decoding is a decoding approach where n candidate translations are sampled from the MT system, and they are used as pseudoreferences for a reference-based *utility metric*. MBR decoding computes the utility metric for all $O(n^2)$ pairs of candidates and pseudoreferences, selecting the candidate that achieves the best average score across all pseudoreferences. Quality Estimation (QE) decoding¹ selects the candidate that scores best according to a QE utility metric. Previous work on MBR decoding has shown that it results in improvements on the utility metric (Amrhein and Sennrich, 2022; Cheng and Vlachos, 2023; Eikema and Aziz, 2022), however other metrics do not improve as much as the utility metric (Guttmann et al., 2024; Vamvas and Sennrich,

2024). This issue of MBR/QE decoding exhibiting bias towards the utility metric complicates our ability to use automatic metrics to compare the quality of MBR/QE-based MT systems, as we cannot tell whether improvements in automatic metrics from MBR/QE decoding correspond to actual improvements in quality, or if it simply reward hacking. Prior work has assumed that this issue can be avoided by using a different metric for evaluating MBR decoding outputs (Tomani et al., 2023), though this assumption has never been tested.

In this work we compare the results of human vs metric-based evaluation of MBR/QE decoding with a wide variety of metrics to show that the quality of MBR/QE decoding is overestimated by not only the utility metric, but also other similar metrics. While MBR/QE decoding with a single utility metric results in significant gains in automatic metrics, it does not perform better than greedy decoding in our human evaluations. This may be due to MBR decoding preferring fluent yet inaccurate candidates. Using an ensemble of metrics as the utility helps us mitigate the metric bias issue, with human evaluations showing that MBR decoding with an ensemble utility metric results in significantly better translations than greedy decoding or MBR/QE decoding with a single utility metric.

In this paper we contribute:

1. A large-scale analysis of metric bias in MBR and QE decoding with metrics commonly used in MT, showing that this metric bias issue holds across many different metrics and language pairs, and is not resolved by simply using a different metric for evaluation.
2. Mitigation strategies for MBR bias using QE filtering followed by MBR decoding, as well as MBR decoding using an ensemble of metrics as the utility function.
3. A human evaluation showing that MBR decoding with ensembles outperforms MBR decoding with a single metric.

¹Also known as QE reranking or QE filtering.

2 Related Work

Cheng and Vlachos (2023); Eikema and Aziz (2022); Guttmann et al. (2024) find that MBR decoding improves automated metrics on various high, medium, and low resource language pairs. Freitag et al. (2023a, 2022); Tomani et al. (2023) find that human raters prefer the outputs of MBR/QE decoding over greedy decoding.

MBR variants achieve speedups via heuristics (Trabelsi et al., 2024; Jinnai and Ariu, 2024), filtering pseudoreferences via a QE metric (Deguchi et al., 2024, 2023) or filtering via another reference-based metric (Vamvas and Sennrich, 2024; Eikema and Aziz, 2022). Quality-aware translation, which incorporates quality estimation into the training process, has been found to improve translation quality over standard MBR (Tomani et al., 2023).

Other techniques for aligning translation models with human preferences include direct preference optimization (Rafailov et al., 2024; Yang et al., 2024), reinforcement learning from human feedback (Christiano et al., 2017), and reinforcement learning from AI feedback (Bai et al., 2022).

Guttmann et al. (2024); Vamvas and Sennrich (2024) show evidence of metric bias in MBR decoding, as they find that neural evaluation metrics favor models using MBR on the metric used as the utility function. However, these papers only cover only 2 metrics, and neither have human evaluations.

Sellam et al. (2020b); Freitag et al. (2023b); Glushkova et al. (2023) find that ensembling metrics can improve their ability to detect critical errors and improve agreement with human preferences, though they do not investigate the effects of ensembling utility metrics on MBR decoding.

Reward hacking (Skalse et al., 2022) is an issue in reinforcement learning where the reward function improves but the system’s behavior is not aligned with human preferences. The metric bias problem in MBR decoding can be viewed as an instance of reward hacking, as the utility function improves while not necessarily improving quality.

3 Study 1: Metric Bias in MBR Decoding

3.1 Methodology

To investigate metric bias in MBR/QE decoding, we perform MBR/QE decoding via various utility metrics and compare how they perform on various evaluation metrics. We investigate MBR decoding using these reference-based utility metrics:

1. MetricX-23 (Juraska et al., 2023)
2. XCOMET-XXL (Guerreiro et al., 2023)
3. XCOMET-XL (Guerreiro et al., 2023)
4. COMET22 (Rei et al., 2022a)
5. AfriCOMET (Wang et al., 2024)
6. IndicCOMET (Sai B et al., 2023)
7. BLEURT (Sellam et al., 2020a)
8. YiSi-1 (Lo, 2019)
9. sentBLEU (Papineni et al., 2002)
10. chrF (Popović, 2015)
11. chrF++ (Popović, 2017)
12. TER (Snover et al., 2006)

We also investigate QE decoding (Fernandes et al., 2022) using the following QE metrics:

1. MetricX-QE (Juraska et al., 2023)
2. CometKiwi23-XXL (Rei et al., 2023)
3. CometKiwi23-XL (Rei et al., 2023)
4. CometKiwi22 (Rei et al., 2022b)
5. AfriCOMET-QE (Wang et al., 2024)

We used a dev set for selecting ensembles, and a test set for reporting final results and human evaluation. The dev datasets and language pairs are:

1. FLORES-200 dev set (Costa-jussà et al., 2022): English-Swahili (en-sw), Igbo (en-ig), Hindi (en-hi), Tamil (en-ta), Somali (en-so), Hausa (en-ha), Malayalam (en-ml), Gujarati (en-gu), Hungarian (en-hu), Vietnamese (en-vi)
2. WMT2022 (Kocmi et al., 2022): English-Chinese (en-zh), Chinese-English (zh-en), English-German (en-de), German-English (de-en)

The test set datasets and language pairs are:

1. FLORES-200 test set: en-sw, en-ig, en-hi, en-ta, en-so, en-ha, en-ml, en-gu, en-hu, en-vi
2. WMT2023 (Kocmi et al., 2023): en-zh, zh-en², en-de, de-en

We produced translations using Gemini 1.0 Pro (Gemini Team Google, 2023) with prompts including 5-shot examples. We used epsilon sampling as recommended by Freitag et al. (2023a) with a sample size of 128. See Appendix A for prompts used for generating translations and instructions for computing scores from metrics.

²Due to errors in the WMT2023 zh-en reference translations, we use the references from Liu et al. (2024) for zh-en.

MBR/QE Method	Evaluated Metric													
	MetricX	MetricX-QE	XCOMET-XXL	XCOMET-XL	CometKiw23-XXL	CometKiw23-XL	CometKiw22	COMET22	BLEURT	YiSi	chrF	chrF++	sentBLEU	TER
Greedy	1.58	1.16	82.3	77.8	76.8	68.2	77.5	85.2	77.3	84.3	57.2	54.2	26.4	63.4
MetricX	0.656‡	0.557‡	85.5‡	79.6‡	79.0‡	69.4‡	77.7‡	84.9‡	76.6‡	81.2‡	50.3‡	46.9‡	18.1‡	75.7‡
MetricX-QE	0.899‡	0.349‡	84.4‡	78.2‡	78.3‡	68.8‡	77.6‡	84.4‡	75.6‡	81.1‡	49.3‡	45.9‡	17.6‡	75.3‡
XCOMET-XXL	1.25‡	0.868‡	89.9‡	80.4‡	80.8‡	69.9‡	78.1‡	85.0‡	76.6‡	81.5‡	50.4‡	47.0‡	18.5‡	73.6‡
XCOMET-XL	1.38‡	1.00‡	86.4‡	85.0‡	80.2‡	71.5‡	78.7‡	85.3‡	77.6‡	82.2‡	51.9‡	48.7‡	20.1‡	71.5‡
CometKiw23-XXL	1.43‡	0.940‡	86.6‡	80.4‡	85.5‡	71.4‡	78.7‡	85.2‡	76.7‡	82.2‡	51.7‡	48.4‡	19.9‡	71.7‡
CometKiw23-XL	1.46‡	0.978‡	85.0‡	81.5‡	81.3‡	74.8‡	78.8‡	85.2‡	76.8‡	82.1‡	51.7‡	48.4‡	19.8‡	72.6‡
CometKiw22	1.57‡	1.07‡	84.0‡	79.6‡	79.7‡	70.5‡	81.9‡	85.4‡	76.8‡	82.3‡	51.9‡	48.6‡	20.1‡	71.0‡
COMET22	1.40‡	1.02‡	84.7‡	80.0‡	79.3‡	70.0‡	78.7‡	87.4‡	78.1‡	83.5‡	55.3‡	52.0‡	23.2‡	67.0‡
BLEURT	1.35‡	0.986‡	83.8‡	79.1‡	78.6‡	69.4‡	78.1‡	85.5‡	82.3‡	82.6‡	53.2‡	49.8‡	21.0‡	71.3‡
YiSi	1.57	1.14†	82.6‡	78.0*	77.3‡	68.7‡	77.7‡	85.6‡	77.7‡	85.0‡	57.7‡	54.5‡	26.1*	62.6
chrF	1.54‡	1.13†	82.6‡	78.0*	77.6‡	68.9‡	77.7‡	85.7‡	77.8‡	84.5‡	58.6‡	55.3‡	25.8‡	65.1‡
chrF++	1.54‡	1.13†	82.6‡	78.0†	77.5‡	68.9‡	77.7‡	85.6‡	77.9‡	84.6‡	58.6‡	55.4‡	26.2	64.6†
sentBLEU	1.61	1.18*	82.2*	77.8*	76.8	68.2	77.5	85.2	77.3*	84.3	57.0‡	54.1*	27.1‡	62.3
TER	1.74‡	1.27‡	81.9‡	77.2‡	75.9‡	67.5‡	77.2‡	84.7‡	76.7‡	83.9‡	55.7‡	52.7‡	25.6‡	59.7‡
rankAvg:all	1.08‡	0.739‡	86.5‡	81.7‡	81.2‡	71.4‡	79.3‡	86.5‡	79.3‡	84.3	57.1	53.9	25.3‡	63.7
rankAvg:qe	1.04‡	0.580‡	86.6‡	81.8‡	83.2‡	73.0‡	80.3‡	85.9‡	77.7‡	82.6‡	52.8‡	49.5‡	20.8‡	70.7‡
rankAvg:top	0.899‡	0.566‡	88.2‡	83.0‡	83.0‡	72.7‡	78.9‡	85.8‡	78.1‡	82.5‡	52.8‡	49.5‡	20.7‡	71.0‡
rankAvg:topQe	1.00‡	0.527‡	86.8‡	81.7‡	83.7‡	73.3‡	78.9‡	85.6‡	77.5	82.4‡	52.3‡	48.9‡	20.2‡	71.7‡
rankAvg:mxmxqe	0.700‡	0.417‡	85.6‡	79.7‡	79.2‡	69.6‡	77.8‡	84.9‡	76.7‡	81.3‡	50.4‡	47.0‡	18.2‡	75.1‡
rankAvg:noLex	0.993‡	0.657‡	87.3‡	82.4‡	82.0‡	72.0‡	79.6‡	86.6‡	79.5‡	83.8‡	55.6‡	52.3‡	23.4‡	66.7‡
rankAvg:noNC	1.09‡	0.734‡	85.2‡	80.4‡	79.5‡	70.1‡	78.5‡	86.4‡	79.2‡	84.4‡	57.4‡	54.1*	25.7‡	63.0*
rankAvg:noNCnoLex	0.968‡	0.636‡	85.8‡	80.8‡	80.0‡	70.4‡	78.6‡	86.6‡	79.7‡	84.0‡	56.1‡	52.8‡	24.0‡	66.0‡
allQE(32)allMBR	1.06‡	0.733‡	86.7‡	81.9‡	81.3‡	71.4‡	79.2‡	86.5‡	79.2‡	84.1‡	56.6‡	53.4‡	24.9‡	64.5
allQE(32)nolexMBR	0.978‡	0.680‡	87.5‡	82.6‡	81.6‡	71.7‡	79.2‡	86.6‡	79.5‡	83.7‡	55.6‡	52.3‡	23.6‡	66.6‡
topQE(32)topMBR	0.861‡	0.599‡	88.4‡	83.3‡	82.0‡	71.9‡	78.8‡	85.7‡	78.1‡	82.4‡	52.7‡	49.4‡	20.7‡	70.9‡
noncQE(32)noncMBR	0.992‡	0.629‡	85.6‡	80.6‡	79.8‡	70.2‡	78.5‡	86.3‡	78.9‡	83.9‡	56.1‡	52.8‡	24.2‡	65.2‡
noncQE(32)noncnolexMBR	0.911‡	0.596‡	86.0‡	81.0‡	80.1‡	70.4‡	78.7‡	86.5‡	79.4‡	83.6‡	55.1‡	51.7‡	22.9‡	67.5‡
mxQE(32)mxMBR	0.662‡	0.475‡	85.6‡	79.8‡	79.2‡	69.5‡	77.8‡	85.0‡	76.8‡	81.5‡	50.7‡	47.3‡	18.5‡	74.9‡
ckQE(32)xcMBR	1.24‡	0.847‡	89.6‡	80.8‡	82.8‡	70.7‡	78.4‡	85.2	77.0‡	81.9‡	51.3‡	48.0‡	19.5‡	72.2‡
mxQE(32)xcMBR	1.03‡	0.593‡	89.5‡	80.6‡	80.9‡	70.1‡	78.2‡	85.1	76.9‡	81.7‡	50.7‡	47.4‡	18.8‡	73.1‡
ckQE(32)mxMBR	0.728‡	0.557‡	86.5‡	80.6‡	82.2‡	70.7‡	78.3‡	85.4‡	77.3	81.9‡	51.7‡	48.3‡	19.5‡	73.3‡

Table 1: Reference-based and QE evaluation scores for greedy and MBR/QE decoding (1st block), and ensembles (2nd block), averaged across all languages (test datasets). Higher scores are better, except MetricX, MetricX-QE, and TER, where lower is better. Green is better than greedy, red is worse. Ensembles are defined in Table 2. Significant differences from greedy (pairwise t-test) indicated by * for p<0.05, † for p<0.01, ‡ for p<0.001. The green diagonal in the 1st block shows metrics prefer outputs from MBR/QE decoding using the same utility metric.

	MetricX-QE	CometKiw23-XXL	CometKiw23-XL	CometKiw22	AfriCOMET-QE (African only)	MetricX	XCOMET-XXL	XCOMET-XL	COMET22	AfriCOMET (African only)	IndicCOMET (Indic only)	BLEURT	YiSi	chrF	chrF++	sentBLEU	TER
all																	
qe																	
top																	
topQe																	
mxmxqe																	
noLex																	
noNC																	
noNCnoLex																	
noNCQe																	
allQE(N)allMBR	1	1	1	1	1	2	2	2	2	2	2	2	2	2	2	2	2
allQE(N)nolexMBR	1	1	1	1	1	2	2	2	2	2	2	2	2				
topQE(N)topMBR	1	1	1			2	2	2									
noncQE(N)noncMBR	1				1	2			2	2	2	2	2	2	2	2	2
noncQE(N)noncnolexMBR	1				1	2			2	2	2	2	2				
mxQE(N)xcMBR	1						2										
ckQE(N)xcMBR		1					2										
mxQE(N)mxMBR	1					2											
ckQE(N)mxMBR		1				2											

Table 2: Metrics included in each ensemble. Rows are ensembles, columns are metrics. Black cells indicate that the metric is included in a single-step ensemble. Green cells indicate the metric is used for the 1st step (QE filtering) in a 2-step ensemble. Red cells indicate the metric is used for the 2nd step (MBR decoding) in a 2-step ensemble.

3.2 Results

Results are shown in Table 1 for average scores across all language pairs on the test datasets. We observe that for all reference-based metrics, the best-performing system is MBR decoding using the same utility metric. This result also holds for all QE metrics, but that is by definition, because QE decoding picks the sample with the best QE score. These results also hold on individual languages and the dev set (Appendix G and E).

We can also see that MBR decoding outputs for utility metrics which are similar to the evaluation metric tend to score better than when the MBR utility metric is dissimilar to the evaluation metric. For example, MBR/QE decoding with neural metrics (MetricX and COMET families) performs better than greedy when evaluated with other neural metrics, but worse than greedy if evaluated via lexical metrics. Likewise, MBR decoding with lexical metrics (sentBLEU, chrF, chrF++, and TER) and semantic metrics (YiSi) perform highly when evaluated by lexical and semantic metrics, but poorly when evaluated via neural metrics. The pattern also holds for similar metrics within the same family – XCOMET-XXL prefers MBR/QE decoding using CometKiwi23-XXL and XCOMET-XL, and MetricX prefers outputs from MetricX-QE.

These results suggest the existence of metric bias in MBR decoding – that is, they suggest that MBR decoding will result in a disproportionately large improvement in the utility metric and metrics similar to the utility metric, relative to the actual improvement in quality. In order to address this issue, in the next section we will investigate ensembling metrics during MBR decoding as a means of avoiding overfitting to a particular utility metric.

4 Study 2: MBR Decoding using Ensembles of Metrics

4.1 Methodology

As a mitigation strategy for utility metric bias in MBR decoding, we investigate how using an ensemble of metrics performs for MBR decoding. We explore the following ensembling techniques (see Appendix C for pseudocode for these techniques):

1. rankAvg: For each metric, assigns a rank to each of the 128 samples (where 0 is best and 127 is worst). Select the sample where the average rank across metrics is minimized.
2. rankMed: Select the sample where the median

rank across metrics is minimized.

3. rankMax: Select the sample where the maximum rank across metrics is minimized.
4. rank75q: Select the sample where the 0.75th quartile rank across metrics is minimized.

For each of these ensembling techniques, we compute ensembles with the following groups of metrics (see Table 2 and Appendix B for the complete list of metrics included in each ensemble):

1. all: All metrics
2. qe: All QE metrics
3. top: Top-performing metrics in WMT2023 metrics shared task (Freitag et al., 2023b)
4. topQe: Top-performing QE metrics
5. mxmxqe: MetricX + MetricX-QE ensemble
6. noLex: Non-lexical metrics
7. noNC: Metrics that permit commercial use
8. noNCnoLex: Non-lexical metrics that permit commercial use
9. noNCQe: QE metrics that permit commercial use

In addition to the ensembles above, we also investigate QE filtering followed by MBR decoding (QE filtering selects the top N candidates according to a QE metric, where N can be either 4, 8, 16, 32, 64). This two-step approach is faster than standard MBR decoding, as QE filtering is linear-time whereas MBR decoding is quadratic time. We include the following two-step ensembles:

1. allQE(N)allMBR: QE filter with all QE metrics, then MBR decode with all reference-based metrics
2. allQE(N)nolexMBR: QE filter with all QE metrics, then MBR decode with non-lexical reference-based metrics
3. topQE(N)topMBR: QE filter with top QE metrics, then MBR decode with top reference-based metrics
4. noncQE(N)noncMBR: QE filter with QE metrics that permit commercial use, then MBR decode with reference-based metrics that permit commercial use
5. noncQE(N)noncnolexMBR: QE filter with QE metrics that permit commercial use, then MBR decode with non-lexical reference-based metrics that permit commercial use
6. mxQE(N)xcMBR: QE filter with MetricX-QE, then MBR decode with XCOMET-XXL
7. ckQE(N)xcMBR: QE filter with CometKiwi23-XXL, then MBR decode

with XCOMET-XXL

8. mxQE(N)mxMBR: QE filter with MetricX-QE, then MBR decode with MetricX
9. ckQE(N)mxMBR: QE filter with CometKiwi23-XXL, then MBR decode with MetricX

The metrics included in each ensemble is shown in Table 2 and Appendix B.

4.2 Results

Results for a subset of ensembles averaged across all language pairs on the test sets are at Table 1 with additional ensembles shown in Appendix F. Results on the dev sets are shown in Appendix E. Breakdowns per language pair can be found in Appendix G. As expected, ensembles tend to perform better if judged by metrics that are better represented in the ensemble; for example, if judging by MetricX, the best ensembles are mxQE(32)mxMBR and rankAvg:mxmxqe, both of which are ensembles consisting of MetricX and MetricX-QE.

That said, observe that compared to MBR/QE decoding with a single utility metric, ensembles often improve on automated evaluations even according to metrics not included in the ensemble. For example, if we use the XCOMET or CometKiwi families of metrics to evaluate rankAvg:noNCnoLex and noncQE(32)noncnolexMBR (which do not include any metrics from the XCOMET or CometKiwi families), they outperform MBR/QE decoding with any single metric outside the XCOMET or CometKiwi families. Similarly, if lexical metrics are used to evaluate the rankAvg:noLex and alQE(32)nolexMBR ensembles, which do not include any lexical metrics, they still outperform MBR/QE decoding with any single neural metric. This suggests that ensembles help reduce metric bias towards a single metric, which results in improved automated evaluation scores according to other metrics not included in the ensemble.

5 Study 3: Human Evaluation

5.1 Methodology

For the human evaluation, we chose the following baselines and ensembles to evaluate:

1. Greedy decoding
2. Reference translation
3. MetricX (MBR decoding)
4. MetricX-QE (QE decoding)

5. AfriCOMET for African languages (MBR decoding)
6. AfriCOMET-QE for African languages (QE decoding)
7. IndicCOMET for Indic languages (MBR decoding)
8. rankAvg:noNC (single-step ensemble)
9. rankAvg:noNCnoLex (single-step ensemble)
10. mxQE(32)mxMBR (multi-step ensemble)
11. noncQE(32)noncnolexMBR (multi-step ensemble)

We evaluated the following conditions only on en-de and zh-en due to budget constraints:

1. XCOMET-XXL (MBR decoding)
2. CometKiwi23-XXL (QE decoding)
3. COMET22 (MBR decoding)
4. rankAvg:all (single-step ensemble)

We chose MetricX, MetricX-QE, AfriCOMET, AfriCOMET-QE, and IndicCOMET because they had shown good performance in previously-published evaluations (Tomani et al., 2023; Wang et al., 2024; Sai B et al., 2023; Freitag et al., 2023b), had good performance in automated evaluations on the dev set (Appendix E), and lacked restrictions on commercial use. In our en-de and zh-en evaluations we also included metrics and ensembles with restrictions on commercial use (XCOMET, CometKiwi, rankAvg:all) for comparison. The 6 language pairs and datasets we evaluate are en-ha en-sw en-ml en-hi (from FLORES200 test) and en-de zh-en (from WMT2023). We chose these languages to have a wide distribution in resource level. For each language pair, we sampled 400 source segments to evaluate. WMT2023 was evaluated with document context, whereas FLORES200 segments were evaluated in isolation. We asked each rater to provide MQM annotations for all translation candidates for each source segment (we evaluated 15 systems on en-de and zh-en and 11 systems on others), and compute scores as described in Freitag et al. (2021). Scores range from 0 to 25, lower is better. To control for variance between raters, the same rater was used to score all candidate translations resulting from each source segment.

5.2 Results

Results are shown in Table 3. We observe that overall the best-performing system is rankAvg:noNC, which significantly outperforms greedy ($p < 0.001$ on pairwise t-test). rankAvg:noNC also performs

	Greedy	Reference	MetricX	MetricX-QE	XCOMET-XXL	CometKiwI23-XXL	COMET22	AfriCOMET	AfriCOMET-QE	IndicCOMET	rankAvg:all	rankAvg:noNC	rankAvg:noNCnoLex	mxQE32mxMBR	noncQE32noncnoLexMBR
all:total	1.52	1.80†	1.59	1.77†								1.27‡	1.28‡	1.53	1.27‡
en-de:total	2.22	2.52	2.38	2.32	2.74	2.96*	2.07				2.07	1.89	1.83	2.13	1.69*
zh-en:total	2.56	2.42	3.15†	3.05	3.04*	2.98	2.65				2.43	2.49	2.53	2.81	2.55
en-sw:total	1.03	1.41	1.08	0.95				0.97	1.44*			0.75*	0.82	0.99	0.86
en-ha:total	1.02	1.25	1.07	1.04				1.17	1.29			0.85	0.95	0.98	0.87
en-hi:total	0.95	1.50‡	0.70	1.09						0.93		0.78	0.71	0.86	0.70*
en-ml:total	1.74	1.94	1.70	2.60‡						2.29†		1.31*	1.28*	1.84	1.39*
all:fluency	0.29	0.38†	0.30	0.33†								0.30‡	0.32‡	0.26	0.26‡
en-de:fluency	0.46	0.45	0.50	0.39	0.46	0.75*	0.38				0.45	0.47	0.45	0.29	0.37*
zh-en:fluency	0.42	0.43	0.28†	0.24	0.27*	0.39	0.32				0.35	0.37	0.39	0.19	0.30
en-sw:fluency	0.14	0.18	0.17	0.27				0.21	0.26*			0.13*	0.12	0.19	0.13
en-ha:fluency	0.37	0.49	0.38	0.36				0.48	0.47			0.33	0.37	0.32	0.33
en-hi:fluency	0.17	0.32‡	0.24	0.30						0.20		0.26	0.26	0.24	0.16*
en-ml:fluency	0.24	0.42	0.26	0.37‡						0.30†		0.26*	0.33*	0.31	0.29*
all:accuracy	0.80	0.94†	0.98	1.06†								0.70‡	0.70‡	0.95	0.74‡
en-de:accuracy	1.06	1.45	1.24	1.42	1.62	1.53*	1.12				1.11	0.86	0.90	1.14	0.85*
zh-en:accuracy	1.72	1.67	2.57†	2.54	2.44*	2.25	2.00				1.74	1.80	1.79	2.34	1.96
en-sw:accuracy	0.58	0.48	0.59	0.44				0.52	0.76*			0.40*	0.44	0.51	0.47
en-ha:accuracy	0.50	0.62	0.59	0.44				0.54	0.70			0.45	0.45	0.58	0.46
en-hi:accuracy	0.32	0.65‡	0.32	0.46						0.44		0.25	0.22	0.41	0.32*
en-ml:accuracy	0.94	1.07	1.11	1.56‡						1.65†		0.80*	0.75*	1.19	0.77*
all:other	0.43	0.48†	0.30	0.38†								0.28‡	0.26‡	0.32	0.27‡
en-de:other	0.69	0.62	0.64	0.51	0.66	0.68*	0.58				0.51	0.56	0.47	0.71	0.46*
zh-en:other	0.42	0.32	0.30†	0.27	0.33*	0.35	0.32				0.35	0.32	0.35	0.27	0.30
en-sw:other	0.31	0.74	0.32	0.24				0.24	0.42*			0.22*	0.25	0.29	0.25
en-ha:other	0.15	0.13	0.10	0.23				0.15	0.12			0.06	0.13	0.08	0.08
en-hi:other	0.46	0.52‡	0.14	0.33						0.29		0.26	0.22	0.21	0.22*
en-ml:other	0.56	0.46	0.33	0.67‡						0.33†		0.25*	0.20*	0.34	0.32*

Table 3: Human evaluation results broken down by language and MQM error type. Columns indicate the system used for MBR/QE decoding; ensembles are defined in Table 2. Rows starting with “all” shows results across all languages. 1st block is total error scores, 2nd is fluency error scores, 3rd is accuracy error scores, 4th is other error scores. For each system, average human evaluation scores across the evaluated segments are shown. Lower scores are better. Colors are relative to greedy, green is better than greedy, red is worse. Black cells were not evaluated. Significant differences from greedy (pairwise t-test) indicated by * for $p < 0.05$, † for $p < 0.01$, ‡ for $p < 0.001$.

the best on each language pair except en-hi. Interestingly, rankAvg:noNC and greedy decoding beat the reference translation in all language pairs, suggesting either that the reference translations in WMT2023 and FLORES200 are of poor quality, or that Gemini’s translation quality has achieved human parity for these language pairs.

A surprising result from our human evaluation was that although MBR decoding with an ensemble of metrics was judged as having superior quality to greedy decoding, MBR/QE decoding with a single metric (MetricX, MetricX-QE, XCOMET-XXL, CometKiwI23-XXL, COMET22, AfriCOMET, AfriCOMET-QE, IndicCOMET) did

not generally improve over greedy decoding (Table 3). In fact, translations from MetricX MBR decoding for zh-en, MetricX-QE decoding for en-ml, AfriCOMET-QE decoding for en-sw, and IndicCOMET MBR decoding for en-ml were rated by humans as significantly worse than greedy decoding (Table 3), even though automatic evaluation with other neural metrics such as MetricX and XCOMET-XXL estimated those translations as being significantly better than greedy (Appendix G). This suggests that evaluation with neural metrics overestimates the quality of MBR/QE decoding, even if different metrics are used for decoding and evaluation. Our findings contrast with previous

System	Translation	Fluency MQM	Accuracy MQM	Other MQM	MetricX	MetricX-QE	XCOMET-XXL	CometKiwi23-XXL	COMET22
Greedy	The seller said not yet, and it will be shipped in the afternoon.	1.0	0.0	0.0	0.659	0.88	0.999	0.83	0.74
MetricX /XCOMET-XXL	The seller said that they don't have it in stock yet , and will be able to ship it out this afternoon .	1.0	10.0	0.0	0.259	0.94	1.000	0.70	0.68
MetricX-QE	The seller said he hadn't shipped it, but could ship it that afternoon.	0.0	0.0	0.0	0.438	0.49	0.997	0.78	0.68
CometKiwi23-XXL	The seller said that it was not ready yet and that it would be shipped that afternoon.	0.0	0.0	0.0	0.264	0.67	0.998	0.87	0.73
COMET22	The seller said not yet, it will be sent in the afternoon.	1.0	0.0	0.0	0.981	1.06	0.998	0.86	0.76
noncQE32noncnolexMBR /rankAvg:noNCnoLex	The seller said no, it won't be shipped until this afternoon.	1.0	0.0	1.0	0.552	0.60	0.998	0.76	0.77
rankAvg:noNC /rankAvg:all	The seller said not yet, it will be shipped in the afternoon.	2.0	0.0	0.0	0.608	0.90	0.998	0.84	0.71
mxQE32mxMBR	The seller said that it is not yet ready, and it will be shipped in the afternoon.	0.0	5.0	0.0	0.432	0.75	0.998	0.84	0.73

Table 4: An example where MetricX and XCOMET-XXL MBR decoding result in an inaccurate translation. The source text is 卖家说还没，下午才能发。（“Seller says not yet, can ship in the afternoon.”） The preceding sentence is 结果，第二天打电话问，发货了吗？（“So the next day I called to ask, has it shipped?”）. MetricX and XCOMET-XXL MBR decoding, as well as the reference-based MetricX and XCOMET-XXL evaluations, all prefer a translation which inaccurately states the item is out of stock. The other metrics assign a lower score to the inaccurate translation. Lower scores are better for MQM, MetricX, and MetricX-QE, for other metrics higher is better. Green is better than greedy, red is worse. Spans marked as errors by the rater are bolded.

studies which find that MBR decoding with a single metric outperforms greedy decoding in human evaluations (Freitag et al., 2022, 2023a; Tomani et al., 2023).

We hypothesize a few potential causes of the failure of single-metric MBR/QE decoding to outperform greedy decoding: firstly, machine translation quality has improved considerably in recent years. This is reflected by how in our study the greedy decoding outputs achieved better human evaluation results compared to the references generated by professional human translators, especially when looking at fluency scores (Table 3), in contrast with previous work where reference translations were rated as better (Freitag et al., 2022, 2023a). Therefore, it is possible that improvements in greedy translation quality have reduced the quality gains from MBR/QE decoding, and have resulted in the adverse effects of metric bias from MBR/QE decoding with a single utility metric outweighing the benefits to translation quality. For example, in Table 3 we can see that single-metric MBR/QE decoding generally improves fluency on high-resource languages, and reduces errors in style, terminology, and locale convention (labeled “other”). How-

ever, accuracy suffers with single-metric MBR/QE decoding for most language pairs (Table 3). We show an example in Table 4, where MetricX and XCOMET-XXL MBR decoding favor a fluent yet inaccurate translation. Perhaps part of the reason for this decrease in accuracy is that MBR decoding with metrics such as MetricX considers only similarity to the pseudoreferences and does not consider the source sentence, so fluent hallucinations that occur in a large number of pseudoreferences will be favored by MBR decoding. Therefore, we hypothesize that past gains from single-metric MBR/QE decoding might have been driven by improvements in fluency and style, but modern LLMs have become good at producing fluent outputs (as indicated by the low fluency error scores for the greedy condition in Table 3), so we are no longer seeing overall quality improvements from single-metric MQM/QE decoding.

We also considered the effects of domain on the quality of single-metric MBR/QE decoding. Since the WMT2023 datasets which were used include novel domains such as speech transcripts and mastodon posts which are not well-represented in the data that metrics such as MetricX and

	Greedy	Reference	MetricX	MetricX-QE	XCOMET-XXL	CometKiw123-XXL	COMET22	rankAvg:all	rankAvg:noNC	rankAvg:noCnoLex	mxQE32mxMBR	noncQE32noncnoLexMBR
en-de@news:total	1.95	2.97	3.47	3.28	2.52	3.74	1.99	1.99	2.16	1.91	2.05	1.78
en-de@user-review:total	3.66	2.79	3.30	2.80	3.11	4.07	3.90	3.71	2.81	3.12	3.09	2.68
en-de@mastodon:total	1.29	1.70	1.17	1.60	1.60	1.87*	1.13	1.19	1.04	1.03	1.65	0.98
en-de@speech:total	3.59	3.78	3.37	2.60	5.43*	3.83	2.97	2.97	2.88	2.65	2.61	2.48
zh-en@news:total	3.56	4.51	3.90	3.90	3.83	4.16	3.36	2.98	3.90	3.15	3.83	4.11
zh-en@user-review:total	2.28	1.73	2.93*	2.71	2.83*	2.62	2.45	2.22	2.06	2.42	2.39	2.01
zh-en@manuals:total	1.70	1.32	2.60*	2.98	2.28	2.21	2.01	2.35	1.58	1.55	2.76	1.89
en-de@news:fluency	0.38	0.69	1.33	0.77	0.31	1.49	0.44	0.46	0.84	0.66	0.36	0.42
en-de@user-review:fluency	0.57	0.65	0.37	0.70	0.52	0.89	0.88	0.53	0.49	0.79	0.18	0.60
en-de@mastodon:fluency	0.15	0.21	0.17	0.13	0.25	0.30*	0.15	0.18	0.12	0.15	0.22	0.21
en-de@speech:fluency	1.21	0.63	0.48	0.34	1.05*	0.90	0.47	1.03	0.87	0.70	0.48	0.55
zh-en@news:fluency	0.29	1.02	0.42	0.36	0.38	0.46	0.29	0.32	0.37	0.42	0.25	0.33
zh-en@user-review:fluency	0.51	0.18	0.22*	0.18	0.23*	0.31	0.32	0.33	0.34	0.37	0.10	0.22
zh-en@manuals:fluency	0.20	0.47	0.29*	0.37	0.28	0.71	0.43	0.58	0.51	0.43	0.57	0.70
en-de@news:accuracy	0.65	1.55	1.63	2.14	1.59	1.63	0.96	0.97	0.59	0.78	1.12	0.95
en-de@user-review:accuracy	2.32	1.25	0.89	1.47	1.16	1.96	1.79	2.37	1.25	1.32	1.00	0.82
en-de@mastodon:accuracy	0.54	1.06	0.68	0.97	0.95	1.10*	0.63	0.57	0.57	0.53	1.01	0.47
en-de@speech:accuracy	1.77	2.41	2.44	1.66	3.65*	2.13	1.94	1.56	1.56	1.61	1.56	1.66
zh-en@news:accuracy	3.03	3.10	3.21	3.39	3.14	3.36	2.79	2.40	3.21	2.39	3.39	3.56
zh-en@user-review:accuracy	1.23	1.21	2.36*	2.20	2.24*	1.92	1.75	1.48	1.35	1.66	1.99	1.46
zh-en@manuals:accuracy	1.38	0.81	2.19*	2.46	1.88	1.38	1.50	1.62	0.96	1.04	1.85	0.88
en-de@news:other	0.92	0.73	0.51	0.37	0.62	0.62	0.59	0.55	0.73	0.46	0.58	0.41
en-de@user-review:other	0.77	0.89	2.04	0.63	1.44	1.21	1.23	0.81	1.07	1.02	1.91	1.26
en-de@mastodon:other	0.60	0.43	0.31	0.49	0.39	0.47*	0.36	0.44	0.36	0.35	0.42	0.30
en-de@speech:other	0.61	0.75	0.45	0.61	0.73*	0.80	0.55	0.38	0.45	0.34	0.56	0.27
zh-en@news:other	0.24	0.39	0.27	0.16	0.31	0.34	0.29	0.26	0.31	0.34	0.19	0.23
zh-en@user-review:other	0.54	0.34	0.35*	0.33	0.36*	0.38	0.38	0.41	0.36	0.39	0.30	0.32
zh-en@manuals:other	0.12	0.04	0.12*	0.15	0.12	0.12	0.08	0.15	0.12	0.08	0.35	0.31

Table 5: Human evaluation results broken down by domain and MQM error type for en-de and zh-en. Columns indicate the system used for MBR/QE decoding; ensembles are defined in Table 2. 1st block is total error scores, 2nd is fluency error scores, 3rd is accuracy error scores, 4th is other error scores. For each system, average human evaluation scores across the evaluated segments are shown. Lower scores are better. Colors are relative to greedy, green is better than greedy, red is worse. Significant differences from greedy (pairwise t-test) indicated by * for $p < 0.05$, † for $p < 0.01$, ‡ for $p < 0.001$.

XCOMET-XXL were trained on, we hypothesized that this may adversely impacting MBR quality. However, contrary to our expectations, as we can observe in Table 5 there is no clear effect of the domain on the quality of MBR decoding results. Thus, we do not believe effects of domain to be the primary factor behind our findings.

We also considered whether MBR decoding with other metrics we did not evaluate with human raters, such as BLEURT, would have performed better than the metrics we evaluated. To do so, we looked at the correlation between the MQM scores from our human evaluation, compared to the scores assigned by metrics. We include scores from QE

metrics (to simulate QE decoding), scores from reference-based metrics based on the 128 pseudoreferences (to simulate MBR decoding), as well as scores from reference-based metrics using the actual references (to simulate a reference-based metric oracle). Table 6 shows Kendall-Tau correlation and Table 7 shows Pearson correlation. Note that this an imperfect simulation of what would happen if we actually performed human evaluation with the MBR/QE decoding outputs for these metrics, as we are considering correlations with human judgments only the subset of candidates which were evaluated (which is a biased sample, as they are the results of MBR/QE decoding), not all 128 samples.

We observe that among the individual metrics that we did not evaluate, simulated XCOMET-XL MBR decoding seems to correlate the best with human judgements, and the other metrics are generally worse than MetricX/XCOMET-XXL MBR decoding. We also include some ensembles, finding that they are generally better correlated with human judgements than individual metrics in our simulation. Therefore, we do not expect that changing to another metric for MBR/QE decoding would have resulted in significantly better translation quality.

6 Discussion

While previous work has sometimes assumed that MBR decoding outputs can be evaluated by automated metrics so long as a non-utility metric is used (Tomani et al., 2023), we find MBR/QE decoding outputs are often preferred by automated metrics despite the fact that human raters believe they are worse quality. For example, while MetricX-QE decoding outputs are considered by human raters to be of worse quality than greedy decoding (Table 3), they still achieve higher scores when evaluated by XCOMET-XXL, XCOMET-XL, MetricX, CometKiwi22, CometKiwi23-XL, and CometKiwi23-XXL (Table 1 and Appendix G). Thus, the metric bias issue that results from MBR/QE decoding complicates evaluation with automated metrics.

That said, while we have shown that MBR/QE decoding generated translations with higher automated evaluation scores are not always judged as having better quality by humans, this does not mean that automated metrics are no longer useful. In our study, automatic reference-based metrics, QE metrics, and ensembles of metrics are still somewhat correlated with MQM scores, as shown in Table 6. Therefore, while it is advisable to perform a human evaluation when feasible if evaluating systems that make use of MBR/QE decoding, existing metrics still correlate with human preferences. Additionally, using an ensemble of metrics for MBR decoding results in improved translation quality compared to greedy decoding and MBR/QE decoding with a single metric (Table 3).

Why is it that using an ensemble of metrics for MBR decoding improves translation quality compared to just using a single metric (Table 3)? We hypothesize that each metric has its own biases that lead it to prefer bad translations, but different metrics have different biases, so using an ensemble

	zh-en	en-de	en-ha	en-sw	en-hi	en-ml
XCOMET-XXL	0.278	0.110	0.114	0.201	0.073	0.152
XCOMET-XXL:mbr	0.275	0.111	0.125	0.212	0.094	0.152
XCOMET-XL	0.335	0.126	0.123	0.187	0.087	0.179
XCOMET-XL:mbr	0.336	0.134	0.137	0.201	0.093	0.168
MetricX	0.252	0.065	0.077	0.192	0.087	0.154
MetricX:mbr	0.289	0.089	0.111	0.211	0.097	0.149
MetricX-QE	0.291	0.046	0.093	0.166	0.065	0.130
CometKiwi23-XXL	0.264	0.080	0.115	0.160	0.085	0.140
CometKiwi23-XL	0.281	0.094	0.113	0.138	0.101	0.165
CometKiwi22	0.274	0.107	0.032	0.173	0.087	0.179
COMET22	0.271	0.100	0.062	0.179	0.076	0.166
COMET22:mbr	0.290	0.125	0.067	0.183	0.088	0.159
BLEURT	0.279	0.128	0.098	0.173	0.083	0.146
BLEURT:mbr	0.271	0.134	0.119	0.187	0.108	0.132
YiSi	0.178	0.049	0.072	0.105	0.061	0.138
YiSi:mbr	0.183	0.068	0.096	0.119	0.065	0.154
chrF	0.044	0.040	0.083	0.115	0.067	0.129
chrF:mbr	0.091	0.049	0.098	0.135	0.056	0.146
chrF++	0.057	0.045	0.084	0.118	0.064	0.123
chrF++:mbr	0.103	0.052	0.098	0.135	0.057	0.141
sentBLEU	0.102	0.059	0.072	0.106	0.052	0.083
sentBLEU:mbr	0.155	0.058	0.082	0.121	0.058	0.103
TER	0.129	0.061	0.084	0.086	0.077	0.087
TER:mbr	0.114	0.060	0.088	0.097	0.067	0.116
MetricX	0.287	0.055	0.084	0.196	0.088	0.155
+MetricX-QE						
MetricX	0.304	0.070	0.107	0.203	0.097	0.151
+MetricX-QE						
XCOMET-XXL	0.326	0.124	0.121	0.210	0.088	0.186
+XCOMET-XL						
XCOMET-XXL:mbr	0.324	0.131	0.136	0.216	0.098	0.177
+XCOMET-XL:mbr						
XCOMET-XXL	0.346	0.127	0.116	0.213	0.090	0.193
+XCOMET-XL						
+COMET22						
XCOMET-XXL:mbr	0.348	0.140	0.129	0.220	0.100	0.184
+XCOMET-XL:mbr						
+COMET22:mbr						

Table 6: Kendall-Tau correlation between MQM evaluation scores and automated evaluation scores. For reference-based metrics, rows with “:mbr” indicate pseudoreference-based evaluation. Bottom rows are ensembles that take the average between the listed metrics. Higher scores indicate better agreement with human raters. See Table 7 for Pearson correlation.

reduces metric bias. We see an example of this in Table 4 where MetricX and XCOMET-XXL assign high scores to an inaccurate translation, but this translation is rated poorly by CometKiwi23-XXL and COMET22, so the ensemble ends up picking a good translation that is preferred by all metrics.

Techniques other than MBR/QE decoding for making use of human preferences to improve translation quality, such as DPO (direct preference optimization) (Rafailov et al., 2024; Yang et al., 2024) and RLHF (reinforcement learning from human feedback) (Christiano et al., 2017), might be more resilient to this metric bias issue, as they do not

directly make use of the evaluation metric. However, given that the data used for DPO/RLHF is similar to the data used to train evaluation metrics, and given that the reward hacking issue is prevalent throughout reinforcement learning (Skalse et al., 2022), issues similar to metric bias may still occur with these techniques.

An open question that remains is how to develop new evaluation techniques that are resilient to metric bias in MBR/QE decoding. One potential way is to develop metrics specialized for evaluating MBR/QE decoding outputs from a particular system, by generating MBR/QE decoding outputs from a translation model, obtaining human annotations for those, and training a metric with them. This process is unfortunately costly and time-intensive, and the learned metric might not be able to generalize beyond translations generated by the particular utility metric and translation model it was trained on. Perhaps a better approach would be to view the metric bias problem as an adversarial learning problem, and apply techniques such as generative adversarial training (Yang et al., 2018) to help train metrics resilient to MBR bias.

7 Conclusion

In this paper we have explored the problem of metric bias, where MBR or QE decoding with a single utility metric shows improvements on automated evaluation with the utility metric and related metrics, but does not actually improve quality when judged by a human rater. We find that the metric bias issue is most severe when using a single utility metric, and using an ensemble of metrics to perform MBR decoding can help improve quality as judged by a human rater. While we have shown that metric bias can result in overly-optimistic automatic evaluations of systems that make use of MBR/QE decoding, the question of how to resolve this issue and automatically evaluate systems that make use of MBR/QE decoding is still an open problem which we leave to future work.

Dataset

Dataset is at <https://mbrbias.github.io/>

Limitations

In this work we compare to only full MBR decoding and QE filtering as baselines, but there are many alternative approaches, such as MBR approximation heuristics (Trabelsi et al., 2024; Jinnai and

Ariu, 2024; Deguchi et al., 2024, 2023; Vamvas and Sennrich, 2024; Eikema and Aziz, 2022), direct preference optimization training (Yang et al., 2024), quality-aware training (Tomani et al., 2023), or training on MBR decoding outputs (Finkelstein and Freitag, 2023), that are more practical to use if translation latency is important. In this work we only look at translations coming from Gemini 1.0 Pro with 5-shot sample prompts and epsilon sampling, and it is possible that results may differ if using a different translation system, different prompts, or a different sampling technique. In this work we only look at using 128 samples due to the computationally expensive $O(n^2)$ cost of running full MBR decoding, but it is possible that using additional samples can achieve further quality improvements. In this work we only looked at segment-level translation, and it is possible that results may differ if performing document-level translation. However, MetricX and the COMET families of models have input token limits – 1024 tokens for MetricX, 512 tokens for COMET – which make it difficult to use them for document-level MBR decoding. Our human evaluation used only a single rater for each translation, which introduces the question of how reliable and consistent the ratings are – using multiple raters and looking at inter-rater agreement is preferable, but was beyond our budget constraints.

Ethics Statement

MBR decoding is resource-intensive, and using ensembles of multiple metrics increases computational complexity compared to a single utility metric. To mitigate this issue, we presented two-step ensembles that use QE filtering followed by MBR decoding, which reduce the computational cost below the cost of standard MBR decoding with a single metric.

References

- 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.
- Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini,

- Cameron McKinnon, et al. 2022. Constitutional ai: Harmlessness from ai feedback. *arXiv preprint arXiv:2212.08073*.
- Julius Cheng and Andreas Vlachos. 2023. **Faster minimum Bayes risk decoding with confidence-based pruning**. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12473–12480, Singapore. Association for Computational Linguistics.
- Paul F Christiano, Jan Leike, Tom Brown, Miljan Martić, Shane Legg, and Dario Amodei. 2017. Deep reinforcement learning from human preferences. *Advances in neural information processing systems*, 30.
- Marta R Costa-jussà, James Cross, Onur Çelebi, Maha Elbayad, Kenneth Heafield, Kevin Heffernan, Elahe Kalbassi, Janice Lam, Daniel Licht, Jean Maillard, et al. 2022. No language left behind: Scaling human-centered machine translation. *arXiv preprint arXiv:2207.04672*.
- Hiroyuki Deguchi, Kenji Imamura, Yuto Nishida, Yusuke Sakai, Justin Vasselli, and Taro Watanabe. 2023. **NAIST-NICT WMT’23 general MT task submission**. In *Proceedings of the Eighth Conference on Machine Translation*, pages 110–118, Singapore. Association for Computational Linguistics.
- Hiroyuki Deguchi, Yusuke Sakai, Hidetaka Kamigaito, Taro Watanabe, Hideki Tanaka, and Masao Utiyama. 2024. Centroid-based efficient minimum bayes risk decoding. *arXiv preprint arXiv:2402.11197*.
- Bryan Eikema and Wilker Aziz. 2022. **Sampling-based approximations to minimum Bayes risk decoding for neural machine translation**. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 10978–10993, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Patrick Fernandes, António Farinhas, Ricardo Rei, José G. C. de Souza, Perez Ogayo, Graham Neubig, and Andre Martins. 2022. **Quality-aware decoding for neural machine translation**. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1396–1412, Seattle, United States. Association for Computational Linguistics.
- Mara Finkelstein and Markus Freitag. 2023. Mbr and qe finetuning: Training-time distillation of the best and most expensive decoding methods. *arXiv preprint arXiv:2309.10966*.
- 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, Behrooz Ghorbani, and Patrick Fernandes. 2023a. **Epsilon sampling rocks: Investigating sampling strategies for minimum Bayes risk decoding for machine translation**. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 9198–9209, Singapore. Association for Computational Linguistics.
- Markus Freitag, David Grangier, Qijun Tan, and Bowen Liang. 2022. **High quality rather than high model probability: Minimum Bayes risk decoding with neural metrics**. *Transactions of the Association for Computational Linguistics*, 10:811–825.
- 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. 2023b. **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.
- Gemini Team Google. 2023. Gemini: A family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.
- Taisiya Glushkova, Chrysoula Zerva, and André F. T. Martins. 2023. **BLEU meets COMET: Combining lexical and neural metrics towards robust machine translation evaluation**. In *Proceedings of the 24th Annual Conference of the European Association for Machine Translation*, pages 47–58, Tampere, Finland. European Association for Machine Translation.
- 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*.
- Kamil Guttman, Mikołaj Pokrywka, Adrian Charkiewicz, and Artur Nowakowski. 2024. **Chasing comet: Leveraging minimum bayes risk decoding for self-improving machine translation**. *arXiv preprint arXiv:2405.11937*.
- Yuu Jinnai and Kaito Ariu. 2024. **Hyperparameter-free approach for faster minimum bayes risk decoding**. *arXiv preprint arXiv:2401.02749*.
- 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.
- 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.
- Zhongtao Liu, Parker Riley, Alison Lui, Daniel Deutsch, Mengmeng Niu, Apurva Shah, and Markus Freitag. 2024. [Beyond human-only: Evaluating human-machine collaboration for collecting high-quality translation data](#). *arXiv preprint*.
- Chi-kiu Lo. 2019. [YiSi - a unified semantic MT quality evaluation and estimation metric for languages with different levels of available resources](#). In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 507–513, Florence, Italy. Association for Computational Linguistics.
- 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.
- 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.
- Maja Popović. 2017. [chrF++: words helping character n-grams](#). In *Proceedings of the Second Conference on Machine Translation*, pages 612–618, Copenhagen, Denmark. Association for Computational Linguistics.
- Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Brussels, Belgium. Association for Computational Linguistics.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. 2024. [Direct preference optimization: Your language model is secretly a reward model](#). *Advances in Neural Information Processing Systems*, 36.
- 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, Josã© Pombal, Daan van Stigt, Marcos Treviso, Luisa Coheur, José G. C. de Souza, and André Martins. 2023. [Scaling up CometKiwi: Unbabel-IST 2023 submission for the quality estimation shared task](#). In *Proceedings of the Eighth Conference on Machine Translation*, pages 841–848, Singapore. 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. 2020a. [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.
- Thibault Sellam, Amy Pu, Hyung Won Chung, Sebastian Gehrmann, Qijun Tan, Markus Freitag, Dipanjan Das, and Ankur Parikh. 2020b. [Learning to evaluate translation beyond English: BLEURT submissions to the WMT metrics 2020 shared task](#). In *Proceedings of the Fifth Conference on Machine Translation*, pages 921–927, Online. Association for Computational Linguistics.
- Joar Skalse, Nikolaus Howe, Dmitrii Krasheninnikov, and David Krueger. 2022. [Defining and characterizing reward hacking](#). *Advances in Neural Information Processing Systems*, 35:9460–9471.

- Matthew Snover, Bonnie Dorr, Rich Schwartz, Linea Micciulla, and John Makhoul. 2006. [A study of translation edit rate with targeted human annotation](#). In *Proceedings of the 7th Conference of the Association for Machine Translation in the Americas: Technical Papers*, pages 223–231, Cambridge, Massachusetts, USA. Association for Machine Translation in the Americas.
- Christian Tomani, David Vilar, Markus Freitag, Colin Cherry, Subhajit Naskar, Mara Finkelstein, and Daniel Cremers. 2023. [Quality control at your fingertips: Quality-aware translation models](#). *arXiv preprint arXiv:2310.06707*.
- Firas Trabelsi, David Vilar, Mara Finkelstein, and Markus Freitag. 2024. [Efficient minimum bayes risk decoding using low-rank matrix completion algorithms](#). *arXiv preprint arXiv:2406.02832*.
- Jannis Vamvas and Rico Sennrich. 2024. [Linear-time minimum bayes risk decoding with reference aggregation](#). *arXiv preprint arXiv:2402.04251*.
- Jiayi Wang, David Adelani, Sweta Agrawal, Marek Masiak, Ricardo Rei, Eleftheria Briakou, Marine Carpuat, Xuanli He, Sofia Bourhim, Andiswa Bukula, Muhidin Mohamed, Temitayo Olatoye, Tosin Adewumi, Hamam Mokayed, Christine Mwase, Wangui Kimotho, Foutse Yuehgoh, Anuoluwapo Aremu, Jessica Ojo, Shamsuddeen Muhammad, Salomey Osei, Abdul-Hakeem Omotayo, Chiamaka Chukwunke, Perez Ogayo, Oumaima Hourrane, Salma El Anigri, Lolwethu Ndolela, Thabiso Mangwana, Shafie Mohamed, Hassan Ayinde, Oluwabusayo Awoyomi, Lama Alkhaled, Sana Al-azzawi, Naome Etori, Millicent Ochieng, Clemencia Siro, Njoroge Kiragu, Eric Muchiri, Wangari Kimotho, Toadoum Sari Sakayo, Lyse Naomi Wamba, Daud Abolade, Simbiat Ajao, Iyanuoluwa Shode, Ricky Macharm, Ruqayya Iro, Saheed Abdullahi, Stephen Moore, Bernard Opoku, Zainab Akinjobi, Abeeb Afolabi, Nnaemeka Obiefuna, Onyekachi Ogbu, Sam Ochieng', Verrah Otiende, Chinedu Mbonu, Yao Lu, and Pontus Stenetorp. 2024. [AfriMTE and AfriCOMET: Enhancing COMET to embrace under-resourced African languages](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 5997–6023, Mexico City, Mexico. Association for Computational Linguistics.
- Guangyu Yang, Jinghong Chen, Weizhe Lin, and Bill Byrne. 2024. [Direct preference optimization for neural machine translation with minimum Bayes risk decoding](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 2: Short Papers)*, pages 391–398, Mexico City, Mexico. Association for Computational Linguistics.
- Zhen Yang, Wei Chen, Feng Wang, and Bo Xu. 2018. [Improving neural machine translation with conditional sequence generative adversarial nets](#). 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 1346–1355, New Orleans, Louisiana. Association for Computational Linguistics.

A Methodology Details

A.1 Prompts Used for Generating Samples

For each language pair, we obtained 5-shot examples for our prompts from the dev split of FLORES-200 by randomly sampling among those reference pairs that had perfect MetricX QE scores (scores of 0). We used MetricX QE filtering to ensure we used high-quality examples as our 5-shot examples. The sampled examples and prompt text for each language pair is included in our dataset release.

A.2 Instructions for Computing Metrics

sentBLEU, chrF, chrF++, and TER scores were computed with sacreBLEU 2.4.2 (Post, 2018) on python 3.11.8 with the following parameters:

chrF: -m chrF

chrF++: -m chrF --chrF-word-order 2

sentBLEU: -m bleu --sentence-level

TER: -m ter

For other metrics, we used the publicly released models on HuggingFace, running with the unbabel-comet package version 2.2.1 available on pip, on Python 3.10.14. We ran on an NVIDIA A100 GPU for all metrics except XCOMET-XXL and CometKiwi23-XXL, which required an NVIDIA A100 80GB GPU.

B Metrics Included in Each Ensemble

This section presents the same information that is present in Table 2, but in textual format. The following are the groups of metrics included in the single-step ensembles that we include in our study. For each of these metric groups the rankAvg, rankMed, rankMax, and rank75q ensembling techniques are used to generate an ensemble.

1. all: All metrics, both reference-based and QE (MetricX, MetricX-QE, XCOMET-XXL, XCOMET-XL, CometKiwi23-XXL, CometKiwi23-XL, CometKiwi22, COMET22, BLEURT, YiSi, chrF, chrF++, sentBLEU, TER, AfriCOMET and

- AfriCOMET-QE for African languages, IndicCOMET for Indic languages)
2. qe: All QE metrics (MetricX-QE, CometKiwi23-XXL, CometKiwi23-XL, CometKiwi22, and AfriCOMET-QE for African languages)
 3. top: MetricX, MetricX-QE, XCOMET-XXL, XCOMET-XL, CometKiwi23-XXL, CometKiwi23-XL
 4. topQe: MetricX-QE, CometKiwi23-XXL, CometKiwi23-XL
 5. mxmxqe: MetricX, MetricX-QE
 6. noLex: All non-lexical metrics (MetricX, MetricX-QE, XCOMET-XXL, XCOMET-XL, CometKiwi23-XXL, CometKiwi23-XL, CometKiwi22, COMET22, BLEURT, YiSi, AfriCOMET and AfriCOMET-QE for African languages, IndicCOMET for Indic languages)
 7. noNC: All metrics that permit commercial use (MetricX, MetricX-QE, CometKiwi22, COMET22, BLEURT, YiSi, chrF, chrF++, sentBLEU, TER, AfriCOMET and AfriCOMET-QE for African languages, IndicCOMET for Indic languages)
 8. noNCnoLex: All non-lexical metrics that permit commercial use (MetricX, MetricX-QE, COMET22, BLEURT, YiSi, AfriCOMET and AfriCOMET-QE for African languages, IndicCOMET for Indic languages)
 9. noNCqe: All QE metrics that permit commercial use (MetricX-QE, and AfriCOMET-QE for African languages)

In addition, we also investigate QE filtering followed by MBR decoding (here we define QE filtering as selecting the top N candidates according to a QE metric, where N can be either 4, 8, 16, 32, 64). We include the following ensembles of this form:

1. allQE(N)allMBR: Use QE filtering with an ensemble of all QE metrics (MetricX-QE, CometKiwi23-XXL, CometKiwi23-XL, CometKiwi22, AfriCOMET-QE for African languages), then perform MBR decoding on the N resulting candidates with all reference-based metrics (MetricX, XCOMET-XXL, XCOMET-XL, COMET22, BLEURT, YiSi, chrF, chrF++, sentBLEU, TER, AfriCOMET for African languages, IndicCOMET for Indic languages).
2. allQE(N)nolexMBR: Use QE filtering with an ensemble of all QE metrics (MetricX-QE, CometKiwi23-XXL, CometKiwi23-XL,

- CometKiwi22, AfriCOMET-QE for African languages), then perform MBR decoding on the N resulting candidates with all non-lexical reference-based metrics (MetricX, XCOMET-XXL, XCOMET-XL, COMET22, BLEURT, YiSi, AfriCOMET for African languages, IndicCOMET for Indic languages).
3. topQE(N)topMBR: Use QE filtering with an ensemble of top-performing QE metrics (MetricX QE, CometKiwi23-XXL, CometKiwi23-XL), then perform MBR decoding on the N resulting candidates with an ensemble of top-performing reference-based metrics (MetricX, XCOMET-XXL, XCOMET-XL).
 4. noncQE(N)noncMBR: Use QE filtering with an ensemble of QE metrics that permit commercial use (MetricX-QE, AfriCOMET-QE for African languages), then perform MBR decoding with an ensemble of reference-based metrics that permit commercial use (MetricX, COMET22, BLEURT, YiSi, chrF, chrF++, sentBLEU, TER, AfriCOMET for African languages, IndicCOMET for Indic languages).
 5. noncQE(N)noncnolexMBR: Use QE filtering with an ensemble of QE metrics that permit commercial use (MetricX-QE, AfriCOMET-QE for African languages), then perform MBR decoding with an ensemble of non-lexical reference-based metrics that permit commercial use (MetricX, COMET22, BLEURT, YiSi, AfriCOMET for African languages, IndicCOMET for Indic languages).
 6. mxQE(N)xcMBR: Use QE filtering with MetricX-QE, then perform MBR decoding with XCOMET-XXL
 7. ckQE(N)xcMBR: Use QE filtering with CometKiwi23-XXL, then perform MBR decoding with XCOMET-XXL
 8. mxQE(N)mxMBR: Use QE filtering with MetricX-QE, then perform MBR decoding with MetricX
 9. ckQE(N)mxMBR: Use QE filtering with CometKiwi23-XXL, then perform MBR decoding with MetricX

C Pseudocode for Ensembles

rankAvg ensembling strategy:

```
def rankAvg(
    sample_list: List[str], metric_list: List[str]
):
    sample_ranks =
        get_ranks_for_samples_by_ensemble(sample_list,
```

```

metric_list)
score_list = [np.mean(x) for x in
sample_ranks]
return select_samples_by_score(sample_list,
score_list)

```

rankMed ensembling strategy:

```

def rankMed(
sample_list: List[str], metric_list: List[str]
):
sample_ranks =
get_ranks_for_samples_by_ensemble(sample_list,
metric_list)
score_list = [np.median(x) for x in
sample_ranks]
return select_samples_by_score(sample_list,
score_list)

```

rankMax ensembling strategy:

```

def rankMax(
sample_list: List[str], metric_list: List[str]
):
sample_ranks =
get_ranks_for_samples_by_ensemble(sample_list,
metric_list)
score_list = [np.max(x) for x in sample_ranks]
return select_samples_by_score(sample_list,
score_list)

```

rank75q ensembling strategy:

```

def rank75q(
sample_list: List[str], metric_list: List[str]
):
sample_ranks =
get_ranks_for_samples_by_ensemble(sample_list,
metric_list)
score_list = [np.quantile(x, q=[0.75])[0] for
x in sample_ranks]
return select_samples_by_score(sample_list,
score_list)

```

Here are helper functions that were used:

```

def get_ranks_for_samples_by_ensemble(
sample_list: List[str], metric_list: List[str]
):
output = [[None for y in metric_list] for x
in sample_list]
for metric_idx, metric in
enumerate(metric_list):
sample_to_rank =
rank_samples_by_metric(sample_list,
metric)
for sample_idx, sample in
enumerate(sample_list):
output[sample_idx][metric_idx] =
sample_to_rank[sample]
return output

def select_samples_by_score(
sample_list: List[str],
score_list: List[float]
):
sample_with_score = zip(sample_list,
score_list)

```

```

top_candidate, top_score =
min(sample_with_score, key=lambda x: x[1])
return top_candidate

```

D Correlation Between Human Evaluation MQM Scores and Metrics

	zh-en	en-de	en-ha	en-sw	en-hi	en-ml
XCOMET-XXL	0.391	0.084	0.146	0.139	0.111	0.202
XCOMET-XXL:mbr	0.389	0.076	0.178	0.145	0.141	0.198
XCOMET-XL	0.543	0.126	0.154	0.160	0.141	0.208
XCOMET-XL:mbr	0.550	0.124	0.170	0.174	0.156	0.194
MetricX	0.391	0.105	0.077	0.146	0.100	0.216
MetricX:mbr	0.431	0.127	0.173	0.150	0.153	0.200
MetricX-QE	0.485	0.120	0.115	0.132	0.074	0.170
CometKiwi23-XXL	0.241	0.088	0.116	0.121	0.128	0.208
CometKiwi23-XL	0.284	0.092	0.098	0.118	0.124	0.202
CometKiwi22	0.277	0.148	0.050	0.156	0.116	0.235
COMET22	0.298	0.170	0.058	0.146	0.099	0.195
COMET22:mbr	0.312	0.209	0.069	0.149	0.118	0.190
BLEURT	0.308	0.152	0.134	0.143	0.115	0.205
BLEURT:mbr	0.322	0.170	0.143	0.150	0.149	0.191
YiSi	0.211	0.088	0.105	0.100	0.092	0.187
YiSi:mbr	0.214	0.124	0.138	0.105	0.100	0.202
chrF	0.054	0.055	0.106	0.106	0.086	0.164
chrF:mbr	0.083	0.069	0.111	0.115	0.084	0.189
chrF++	0.062	0.058	0.109	0.108	0.087	0.157
chrF++:mbr	0.091	0.069	0.110	0.113	0.088	0.183
sentBLEU	0.128	0.072	0.095	0.094	0.073	0.091
sentBLEU:mbr	0.160	0.072	0.098	0.111	0.088	0.113
TER	0.101	0.071	0.104	0.061	0.105	0.118
TER:mbr	0.096	0.085	0.105	0.067	0.107	0.149
MetricX	0.463	0.124	0.101	0.152	0.101	0.229
+MetricX-QE	0.483	0.130	0.160	0.151	0.131	0.209
XCOMET-XXL	0.532	0.110	0.159	0.161	0.144	0.228
+XCOMET-XL	0.537	0.105	0.183	0.171	0.166	0.215
XCOMET-XXL:mbr	0.521	0.136	0.150	0.169	0.144	0.235
+XCOMET-XL	0.529	0.136	0.173	0.176	0.167	0.223
XCOMET-XXL:mbr	0.529	0.136	0.173	0.176	0.167	0.223
+XCOMET-XL:mbr						
+COMET22:mbr						

Table 7: Pearson correlation between MQM evaluation scores and automated evaluation scores. For reference-based metrics, rows with “:mbr” indicate pseudoreference-based evaluation. Bottom rows are ensembles that take the average between the listed metrics. Higher scores indicate better agreement with human raters. See Table 6 for Kendall-Tau correlation.

E Results on Dev Datasets (WMT2022 and FLORES200 dev)

MBR/QE Method	Evaluated Metric													
	MetricX	MetricX-QE	XCOMET-XXL	XCOMET-XL	CometKiwi23-XXL	CometKiwi23-XL	CometKiwi22	COMET22	BLEURT	YiSi	chrF	chrF++	sentBLEU	TER
Greedy	1.58	1.16	82.3	77.8	76.8	68.2	77.5	85.2	77.3	84.3	57.2	54.2	26.4	63.4
rankAvg:all	1.08‡	0.739‡	86.5‡	81.7‡	81.2‡	71.4‡	79.3‡	86.5‡	79.3‡	84.3	57.1	53.9	25.3‡	63.7
rankAvg:qe	1.04‡	0.580‡	86.6‡	81.8‡	83.2‡	73.0‡	80.3‡	85.9‡	77.7‡	82.6‡	52.8‡	49.5‡	20.8‡	70.7‡
rankAvg:top	0.899‡	0.566‡	88.2‡	83.0‡	83.0‡	72.7‡	78.9‡	85.8‡	78.1‡	82.5‡	52.8‡	49.5‡	20.7‡	71.0‡
rankAvg:topQe	1.00‡	0.527‡	86.8‡	81.7‡	83.7‡	73.3‡	78.9‡	85.6‡	77.5‡	82.4‡	52.3‡	48.9‡	20.2‡	71.7‡
rankAvg:mxmxqe	0.700‡	0.417‡	85.6‡	79.7‡	79.2‡	69.6‡	77.8‡	84.9‡	76.7‡	81.3‡	50.4‡	47.0‡	18.2‡	75.1‡
rankAvg:noLex	0.993‡	0.657‡	87.3‡	82.4‡	82.0‡	72.0‡	79.6‡	86.6‡	79.5‡	83.8‡	55.6‡	52.3‡	23.4‡	66.7‡
rankAvg:noNC	1.09‡	0.734‡	85.2‡	80.4‡	79.5‡	70.1‡	78.5‡	86.4‡	79.2‡	84.4‡	57.4‡	54.1*	25.7‡	63.0*
rankAvg:noNCnoLex	0.968‡	0.636‡	85.8‡	80.8‡	80.0‡	70.4‡	78.6‡	86.6‡	79.7‡	84.0‡	56.1‡	52.8‡	24.0‡	66.0‡
rankAvg:noNCQe	0.934‡	0.400‡	84.5‡	78.3‡	78.5‡	69.0‡	77.7‡	84.6‡	75.6‡	81.1‡	49.5‡	46.1‡	17.6‡	75.5‡
rankMax:all	1.16‡	0.776‡	86.1‡	81.0‡	80.8‡	71.1‡	79.2‡	86.3‡	78.9‡	83.9‡	56.1‡	52.8‡	24.3‡	64.1
rankMax:qe	1.06‡	0.595‡	86.3‡	81.5‡	82.8‡	72.6‡	80.2‡	85.9‡	77.7‡	82.7‡	53.0‡	49.6‡	20.9‡	70.5‡
rankMax:top	0.929‡	0.586‡	88.0‡	82.7‡	82.7‡	71.4‡	78.8‡	85.7‡	78.0‡	82.5‡	52.8‡	49.5‡	20.8‡	70.6‡
rankMax:topQe	0.964‡	0.480‡	86.7‡	80.7‡	84.0‡	71.2‡	78.6‡	85.4‡	77.0‡	82.1‡	51.7‡	48.3‡	19.7‡	72.0‡
rankMax:mxmxqe	0.704‡	0.420‡	85.6‡	79.7‡	79.3‡	69.6‡	77.8‡	84.9‡	76.7‡	81.3‡	50.5‡	47.1‡	18.2‡	75.0‡
rankMax:noLex	1.11‡	0.739‡	86.6‡	81.5‡	81.3‡	71.4‡	79.4‡	86.4‡	79.1‡	83.8‡	55.5‡	52.2‡	23.4‡	66.5‡
rankMax:noNC	1.11‡	0.733‡	85.1‡	80.1‡	79.3‡	69.9‡	78.4‡	86.3‡	79.1‡	84.0‡	56.3‡	53.1‡	24.7‡	63.6
rankMax:noNCnoLex	1.05‡	0.685‡	85.4‡	80.4‡	79.6‡	70.2‡	78.5‡	86.4‡	79.5‡	83.9‡	55.9‡	52.6‡	23.8‡	66.0‡
rankMax:noNCQe	0.937‡	0.405‡	84.5‡	78.3‡	78.5‡	69.0‡	77.6‡	84.6‡	75.6‡	81.1‡	49.4‡	46.0‡	17.6‡	75.5‡
rankMed:all	1.06‡	0.733‡	86.5‡	81.9‡	81.0‡	71.3‡	79.1‡	86.5‡	79.2‡	84.1‡	56.8‡	53.6‡	25.1‡	64.5*
rankMed:qe	1.14‡	0.679‡	86.5‡	81.7‡	83.3‡	73.0‡	79.9‡	85.7‡	77.5‡	82.4‡	52.3‡	49.0‡	20.3‡	71.6‡
rankMed:top	0.895‡	0.573‡	88.2‡	83.1‡	82.8‡	72.5‡	78.9‡	85.6‡	77.9‡	82.2‡	52.3‡	48.9‡	20.1‡	71.9‡
rankMed:topQe	1.21‡	0.726‡	86.5‡	81.4‡	83.8‡	73.2‡	78.9‡	85.3‡	77.1‡	82.1‡	51.7‡	48.3‡	19.7‡	72.4‡
rankMed:mxmxqe	0.700‡	0.417‡	85.6‡	79.7‡	79.2‡	69.6‡	77.8‡	84.9‡	76.7‡	81.3‡	50.4‡	47.0‡	18.2‡	75.1‡
rankMed:noLex	0.935‡	0.611‡	87.6‡	82.8‡	82.2‡	72.2‡	79.4‡	86.4‡	79.1‡	83.1‡	54.3‡	51.0‡	22.1‡	69.0‡
rankMed:noNC	1.28‡	0.927‡	84.2‡	79.6‡	78.6‡	69.5‡	78.2‡	86.2‡	78.7‡	84.6‡	57.9‡	54.7‡	26.3	62.6*
rankMed:noNCnoLex	0.910‡	0.607‡	85.8‡	80.9‡	80.0‡	70.4‡	78.6‡	86.5‡	79.3‡	83.5‡	55.1‡	51.8‡	23.0‡	67.9‡
rankMed:noNCQe	0.934‡	0.400‡	84.5‡	78.3‡	78.5‡	69.0‡	77.7‡	84.6‡	75.6‡	81.1‡	49.5‡	46.1‡	17.6‡	75.5‡
rank75q:all	1.09‡	0.743‡	86.5‡	81.7‡	81.1‡	71.3‡	79.1‡	86.5‡	79.2‡	84.2	56.9‡	53.6‡	25.0‡	64.2
rank75q:qe	1.06‡	0.600‡	86.5‡	81.7‡	83.2‡	72.9‡	80.0‡	85.9‡	77.6‡	82.6‡	52.7‡	49.4‡	20.7‡	70.9‡
rank75q:top	0.892‡	0.564‡	88.0‡	82.9‡	82.8‡	72.6‡	78.9‡	85.7‡	78.0‡	82.4‡	52.7‡	49.4‡	20.6‡	71.2‡
rank75q:topQe	1.00‡	0.526‡	86.7‡	81.7‡	83.6‡	73.3‡	78.9‡	85.6‡	77.5‡	82.4‡	52.3‡	49.0‡	20.3‡	71.6‡
rank75q:mxmxqe	0.705‡	0.419‡	85.6‡	79.7‡	79.2‡	69.6‡	77.8‡	84.9‡	76.7‡	81.3‡	50.5‡	47.0‡	18.2‡	75.1‡
rank75q:noLex	0.990‡	0.651‡	87.3‡	82.5‡	82.0‡	72.0‡	79.5‡	86.5‡	79.4‡	83.5‡	55.2‡	51.9‡	23.0‡	67.4‡
rank75q:noNC	1.13‡	0.780‡	85.0‡	80.2‡	79.3‡	69.9‡	78.4‡	86.4‡	79.0‡	84.3‡	57.3*	54.1	25.6‡	63.3
rank75q:noNCnoLex	0.955‡	0.628‡	85.8‡	80.9‡	80.0‡	70.4‡	78.6‡	86.6‡	79.6‡	83.7‡	55.6‡	52.2‡	23.4‡	67.0‡
rank75q:noNCQe	0.937‡	0.403‡	84.5‡	78.3‡	78.5‡	69.0‡	77.6‡	84.6‡	75.6‡	81.1‡	49.4‡	46.1‡	17.6‡	75.5‡

Table 8: Reference-based and QE evaluation scores for greedy, MBR, and QE decoding using a single-step ensemble utility metric, averaged across all languages (test datasets). Higher scores are better, except MetricX, MetricX-QE, and TER, where lower is better. Green is better than greedy, red is worse. Ensembles are defined in Table 2. Significant differences from greedy (pairwise t-test) indicated by * for $p < 0.05$, † for $p < 0.01$, ‡ for $p < 0.001$.

F Results for Additional Ensembles

F.1 Additional Single-Step Ensembles on Test Datasets

MBR/QE Method	Evaluated Metric													
	MetricX	MetricX-QE	XCOMET-XXL	XCOMET-XL	CometKiwi23-XXL	CometKiwi23-XL	CometKiwi22	COMET22	BLEURT	YiSi	chrF	chrF++	sentBLEU	TER
Greedy	1.58	1.16	82.3	77.8	76.8	68.2	77.5	85.2	77.3	84.3	57.2	54.2	26.4	63.4
rankAvg:all	1.08‡	0.739‡	86.5‡	81.7‡	81.2‡	71.4‡	79.3‡	86.5‡	79.3‡	84.3	57.1	53.9	25.3‡	63.7
rankAvg:qe	1.04‡	0.580‡	86.6‡	81.8‡	83.2‡	73.0‡	80.3‡	85.9‡	77.7‡	82.6‡	52.8‡	49.5‡	20.8‡	70.7‡
rankAvg:top	0.899‡	0.566‡	88.2‡	83.0‡	83.0‡	72.7‡	78.9‡	85.8‡	78.1‡	82.5‡	52.8‡	49.5‡	20.7‡	71.0‡
rankAvg:topQe	1.00‡	0.527‡	86.8‡	81.7‡	83.7‡	73.3‡	78.9‡	85.6‡	77.5	82.4‡	52.3‡	48.9‡	20.2‡	71.7‡
rankAvg:mxmxqe	0.700‡	0.417‡	85.6‡	79.7‡	79.2‡	69.6‡	77.8‡	84.9‡	76.7‡	81.3‡	50.4‡	47.0‡	18.2‡	75.1‡
rankAvg:noLex	0.993‡	0.657‡	87.3‡	82.4‡	82.0‡	72.0‡	79.6‡	86.6‡	79.5‡	83.8‡	55.6‡	52.3‡	23.4‡	66.7‡
rankAvg:noNC	1.09‡	0.734‡	85.2‡	80.4‡	79.5‡	70.1‡	78.5‡	86.4‡	79.2‡	84.4‡	57.4‡	54.1*	25.7‡	63.0*
rankAvg:noNCnoLex	0.968‡	0.636‡	85.8‡	80.8‡	80.0‡	70.4‡	78.6‡	86.6‡	79.7‡	84.0‡	56.1‡	52.8‡	24.0‡	66.0‡
rankAvg:noNCQe	0.934‡	0.400‡	84.5‡	78.3‡	78.5‡	69.0‡	77.7‡	84.6‡	75.6‡	81.1‡	49.5‡	46.1‡	17.6‡	75.5‡
rankMax:all	1.16‡	0.776‡	86.1‡	81.0‡	80.8‡	71.1‡	79.2‡	86.3‡	78.9‡	83.9‡	56.1‡	52.8‡	24.3‡	64.1
rankMax:qe	1.06‡	0.595‡	86.3‡	81.5‡	82.8‡	72.6‡	80.2‡	85.9‡	77.7‡	82.7‡	53.0‡	49.6‡	20.9‡	70.5‡
rankMax:top	0.929‡	0.586‡	88.0‡	82.7‡	82.7‡	71.4‡	78.8‡	85.7‡	78.0‡	82.5‡	52.8‡	49.5‡	20.8‡	70.6‡
rankMax:topQe	0.964‡	0.480‡	86.7‡	80.7‡	84.0‡	71.2‡	78.6‡	85.4‡	77.0‡	82.1‡	51.7‡	48.3‡	19.7‡	72.0‡
rankMax:mxmxqe	0.704‡	0.420‡	85.6‡	79.7‡	79.3‡	69.6‡	77.8‡	84.9‡	76.7‡	81.3‡	50.5‡	47.1‡	18.2‡	75.0‡
rankMax:noLex	1.11‡	0.739‡	86.6‡	81.5‡	81.3‡	71.4‡	79.4‡	86.4‡	79.1‡	83.8‡	55.5‡	52.2‡	23.4‡	66.5‡
rankMax:noNC	1.11‡	0.733‡	85.1‡	80.1‡	79.3‡	69.9‡	78.4‡	86.3‡	79.1‡	84.0‡	56.3‡	53.1‡	24.7‡	63.6
rankMax:noNCnoLex	1.05‡	0.685‡	85.4‡	80.4‡	79.6‡	70.2‡	78.5‡	86.4‡	79.5‡	83.9‡	55.9‡	52.6‡	23.8‡	66.0‡
rankMax:noNCQe	0.937‡	0.405‡	84.5‡	78.3‡	78.5‡	69.0‡	77.6‡	84.6‡	75.6‡	81.1‡	49.4‡	46.0‡	17.6‡	75.5‡
rankMed:all	1.06‡	0.733‡	86.5‡	81.9‡	81.0‡	71.3‡	79.1‡	86.5‡	79.2‡	84.1‡	56.8‡	53.6‡	25.1‡	64.5*
rankMed:qe	1.14‡	0.679‡	86.5‡	81.7‡	83.3‡	73.0‡	79.9‡	85.7‡	77.5	82.4‡	52.3‡	49.0‡	20.3‡	71.6‡
rankMed:top	0.895‡	0.573‡	88.2‡	83.1‡	82.8‡	72.5‡	78.9‡	85.6‡	77.9‡	82.2‡	52.3‡	48.9‡	20.1‡	71.9‡
rankMed:topQe	1.21‡	0.726‡	86.5‡	81.4‡	83.8‡	73.2‡	78.9‡	85.3‡	77.1‡	82.1‡	51.7‡	48.3‡	19.7‡	72.4‡
rankMed:mxmxqe	0.700‡	0.417‡	85.6‡	79.7‡	79.2‡	69.6‡	77.8‡	84.9‡	76.7‡	81.3‡	50.4‡	47.0‡	18.2‡	75.1‡
rankMed:noLex	0.935‡	0.611‡	87.6‡	82.8‡	82.2‡	72.2‡	79.4‡	86.4‡	79.1‡	83.1‡	54.3‡	51.0‡	22.1‡	69.0‡
rankMed:noNC	1.28‡	0.927‡	84.2‡	79.6‡	78.6‡	69.5‡	78.2‡	86.2‡	78.7‡	84.6‡	57.9‡	54.7‡	26.3	62.6*
rankMed:noNCnoLex	0.910‡	0.607‡	85.8‡	80.9‡	80.0‡	70.4‡	78.6‡	86.5‡	79.3‡	83.5‡	55.1‡	51.8‡	23.0‡	67.9‡
rankMed:noNCQe	0.934‡	0.400‡	84.5‡	78.3‡	78.5‡	69.0‡	77.7‡	84.6‡	75.6‡	81.1‡	49.5‡	46.1‡	17.6‡	75.5‡
rank75q:all	1.09‡	0.743‡	86.5‡	81.7‡	81.1‡	71.3‡	79.1‡	86.5‡	79.2‡	84.2	56.9‡	53.6‡	25.0‡	64.2
rank75q:qe	1.06‡	0.600‡	86.5‡	81.7‡	83.2‡	72.9‡	80.0‡	85.9‡	77.6‡	82.6‡	52.7‡	49.4‡	20.7‡	70.9‡
rank75q:top	0.892‡	0.564‡	88.0‡	82.9‡	82.8‡	72.6‡	78.9‡	85.7‡	78.0‡	82.4‡	52.7‡	49.4‡	20.6‡	71.2‡
rank75q:topQe	1.00‡	0.526‡	86.7‡	81.7‡	83.6‡	73.3‡	78.9‡	85.6‡	77.5	82.4‡	52.3‡	49.0‡	20.3‡	71.6‡
rank75q:mxmxqe	0.705‡	0.419‡	85.6‡	79.7‡	79.2‡	69.6‡	77.8‡	84.9‡	76.7‡	81.3‡	50.5‡	47.0‡	18.2‡	75.1‡
rank75q:noLex	0.990‡	0.651‡	87.3‡	82.5‡	82.0‡	72.0‡	79.5‡	86.5‡	79.4‡	83.5‡	55.2‡	51.9‡	23.0‡	67.4‡
rank75q:noNC	1.13‡	0.780‡	85.0‡	80.2‡	79.3‡	69.9‡	78.4‡	86.4‡	79.0‡	84.3‡	57.3*	54.1	25.6‡	63.3
rank75q:noNCnoLex	0.955‡	0.628‡	85.8‡	80.9‡	80.0‡	70.4‡	78.6‡	86.6‡	79.6‡	83.7‡	55.6‡	52.2‡	23.4‡	67.0‡
rank75q:noNCQe	0.937‡	0.403‡	84.5‡	78.3‡	78.5‡	69.0‡	77.6‡	84.6‡	75.6‡	81.1‡	49.4‡	46.1‡	17.6‡	75.5‡

Table 9: Reference-based and QE evaluation scores for greedy, MBR, and QE decoding using a single-step ensemble utility metric, averaged across all languages (test datasets). Higher scores are better, except MetricX, MetricX-QE, and TER, where lower is better. Green is better than greedy, red is worse. Ensembles are defined in Table 2. Significant differences from greedy (pairwise t-test) indicated by * for $p < 0.05$, † for $p < 0.01$, ‡ for $p < 0.001$.

F.2 Additional Two-Step Ensembles on Test Datasets

MBR/QE Method	Evaluated Metric														
	MetricX	MetricX-QE	XCOMET-XXL	XCOMET-XL	CometKiwii23-XXL	CometKiwii23-XL	CometKiwii22	COMET22	BLEURT	YiSi	chrF	chrF++	sentBLEU	TER	
Greedy	1.58	1.16	82.3	77.8	76.8	68.2	77.5	85.2	77.3	84.3	57.2	54.2	26.4	63.4	
allQE(64)allMBR	1.09‡	0.781‡	86.5‡	81.7‡	80.6‡	71.0‡	78.9‡	86.5‡	79.3‡	84.3	57.1	53.9	25.4‡	63.6	
allQE(32)allMBR	1.06‡	0.733‡	86.7‡	81.9‡	81.3‡	71.4‡	79.2‡	86.5‡	79.2‡	84.1‡	56.6‡	53.4‡	24.9‡	64.5	
allQE(16)allMBR	1.04‡	0.688‡	86.8‡	82.0‡	81.8‡	71.9‡	79.4‡	86.4‡	79.1‡	83.9‡	56.2‡	52.9‡	24.3‡	65.4‡	
allQE(8)allMBR	1.04‡	0.654‡	86.8‡	82.0‡	82.2‡	72.2‡	79.7‡	86.3‡	78.7‡	83.6‡	55.4‡	52.1‡	23.4‡	66.8‡	
allQE(4)allMBR	1.04‡	0.629‡	86.8‡	82.0‡	82.7‡	72.5‡	79.9‡	86.2‡	78.4‡	83.2‡	54.5‡	51.2‡	22.3‡	68.3‡	
allQE(64)nolexMBR	0.991‡	0.708‡	87.4‡	82.4‡	81.1‡	71.3‡	79.0‡	86.6‡	79.7‡	83.9‡	55.9‡	52.6‡	23.9‡	66.0‡	
allQE(32)nolexMBR	0.978‡	0.680‡	87.5‡	82.6‡	81.6‡	71.7‡	79.2‡	86.6‡	79.5‡	83.7‡	55.6‡	52.3‡	23.6‡	66.6‡	
allQE(16)nolexMBR	0.972‡	0.647‡	87.5‡	82.6‡	82.1‡	72.0‡	79.4‡	86.5‡	79.3‡	83.5‡	55.2‡	51.9‡	23.2‡	67.2‡	
allQE(8)nolexMBR	0.977‡	0.625‡	87.3‡	82.5‡	82.4‡	72.3‡	79.7‡	86.4‡	79.0‡	83.3‡	54.6‡	51.3‡	22.5‡	68.3‡	
allQE(4)nolexMBR	0.988‡	0.608‡	87.2‡	82.4‡	82.8‡	72.6‡	79.9‡	86.2‡	78.6‡	83.0‡	53.9‡	50.5‡	21.7‡	69.4‡	
topQE(64)topMBR	0.868‡	0.621‡	88.5‡	83.3‡	81.5‡	71.5‡	78.7‡	85.6‡	78.1‡	82.3‡	52.4‡	49.1‡	20.4‡	71.2‡	
topQE(32)topMBR	0.861‡	0.599‡	88.4‡	83.3‡	82.0‡	71.9‡	78.8‡	85.7‡	78.1‡	82.4‡	52.7‡	49.4‡	20.7‡	70.9‡	
topQE(16)topMBR	0.879‡	0.585‡	88.3‡	83.2‡	82.4‡	72.2‡	78.9‡	85.7‡	78.1‡	82.5‡	52.8‡	49.4‡	20.8‡	70.8‡	
topQE(8)topMBR	0.897‡	0.567‡	88.1‡	82.9‡	82.8‡	72.6‡	78.9‡	85.7‡	78.0‡	82.5‡	52.8‡	49.5‡	20.7‡	71.0‡	
topQE(4)topMBR	0.925‡	0.548‡	87.7‡	82.6‡	83.2‡	72.9‡	78.9‡	85.7‡	77.8‡	82.4‡	52.6‡	49.2‡	20.5‡	71.3‡	
noncQE(64)noncnolexMBR	0.955‡	0.668‡	85.9‡	81.0‡	80.0‡	70.4‡	78.7‡	86.6‡	79.8‡	83.9‡	55.9‡	52.6‡	23.8‡	66.3‡	
noncQE(32)noncnolexMBR	0.911‡	0.596‡	86.0‡	81.0‡	80.1‡	70.4‡	78.7‡	86.5‡	79.4‡	83.6‡	55.1‡	51.7‡	22.9‡	67.5‡	
noncQE(16)noncnolexMBR	0.883‡	0.533‡	86.0‡	80.8‡	80.0‡	70.3‡	78.6‡	86.2‡	78.8‡	83.1‡	54.1‡	50.7‡	21.8‡	69.2‡	
noncQE(8)noncnolexMBR	0.877‡	0.487‡	85.8‡	80.5‡	79.5‡	70.2‡	78.4‡	85.9‡	78.2‡	82.6‡	53.0‡	49.6‡	20.7‡	70.7‡	
noncQE(4)noncnolexMBR	0.890‡	0.450‡	85.4‡	79.8‡	79.5‡	69.9‡	78.2‡	85.5‡	77.4‡	82.0‡	51.5‡	48.1‡	19.4‡	72.9‡	
noncQE(64)noncMBR	1.06‡	0.728‡	85.3‡	80.6‡	79.6‡	70.1‡	78.5‡	86.4‡	79.2‡	84.3	57.0	53.8‡	25.3‡	63.7	
noncQE(32)noncMBR	0.992‡	0.629‡	85.6‡	80.6‡	79.8‡	70.2‡	78.5‡	86.3‡	78.9‡	83.9‡	56.1‡	52.8‡	24.2‡	65.2‡	
noncQE(16)noncMBR	0.960‡	0.559‡	85.6‡	80.4‡	79.7‡	70.1‡	78.5‡	86.0‡	78.4‡	83.5‡	54.9‡	51.7‡	22.9‡	67.1‡	
noncQE(8)noncMBR	0.942‡	0.506‡	85.5‡	80.1‡	79.6‡	69.9‡	78.3‡	85.8‡	77.8‡	82.9‡	53.7‡	50.3‡	21.5‡	69.2‡	
noncQE(4)noncMBR	0.931‡	0.461‡	85.2‡	79.5‡	79.3‡	69.7‡	78.1‡	85.4‡	77.0‡	82.2‡	52.1‡	48.7‡	19.9‡	71.8‡	
mxQE(64)xcMBR	1.11‡	0.690‡	89.8‡	80.6‡	80.9‡	70.1‡	78.2‡	85.1	76.9‡	81.7‡	50.7‡	47.3‡	18.8‡	73.1‡	
mxQE(32)xcMBR	1.03‡	0.593‡	89.5‡	80.6‡	80.9‡	70.1‡	78.2‡	85.1	76.9‡	81.7‡	50.7‡	47.4‡	18.8‡	73.1‡	
mxQE(16)xcMBR	0.965‡	0.517‡	89.1‡	80.5‡	80.7‡	70.0‡	78.2‡	85.1*	76.9‡	81.6‡	50.6‡	47.2‡	18.7‡	73.4‡	
mxQE(8)xcMBR	0.924‡	0.459‡	88.4‡	80.3‡	80.3‡	69.9‡	78.1‡	85.0‡	76.7‡	81.6‡	50.4‡	47.1‡	18.6‡	73.4‡	
mxQE(4)xcMBR	0.904‡	0.411‡	87.5‡	79.8‡	79.9‡	69.8‡	78.0‡	84.9‡	76.4‡	81.4‡	50.1‡	46.8‡	18.3‡	73.9‡	
ckQE(64)xcMBR	1.23‡	0.851‡	89.8‡	80.7‡	81.9‡	70.4‡	78.3‡	85.1	76.9‡	81.8‡	50.9‡	47.6‡	19.1‡	72.8‡	
ckQE(32)xcMBR	1.24‡	0.847‡	89.6‡	80.8‡	82.8‡	70.7‡	78.4‡	85.2	77.0‡	81.9‡	51.3‡	48.0‡	19.5‡	72.2‡	
ckQE(16)xcMBR	1.25‡	0.850‡	89.3‡	81.0‡	83.5‡	71.0‡	78.6‡	85.3‡	77.1‡	82.1‡	51.6‡	48.3‡	19.9‡	71.6‡	
ckQE(8)xcMBR	1.30‡	0.870‡	88.9‡	80.9‡	84.1‡	71.2‡	78.7‡	85.3‡	77.0‡	82.2‡	51.8‡	48.5‡	19.9‡	71.5‡	
ckQE(4)xcMBR	1.33‡	0.883‡	88.3‡	80.8‡	84.7‡	71.4‡	78.7‡	85.3‡	76.9‡	82.2‡	51.8‡	48.5‡	20.0‡	71.5‡	
mxQE(64)mxMBR	0.653‡	0.508‡	85.6‡	79.8‡	79.2‡	69.5‡	77.8‡	85.0‡	76.8‡	81.4‡	50.6‡	47.2‡	18.4‡	75.2‡	
mxQE(32)mxMBR	0.662‡	0.475‡	85.6‡	79.8‡	79.2‡	69.5‡	77.8‡	85.0‡	76.8‡	81.5‡	50.7‡	47.3‡	18.5‡	74.9‡	
mxQE(16)mxMBR	0.681‡	0.450‡	85.5‡	79.6‡	79.1‡	69.4‡	77.8‡	85.0‡	76.7‡	81.5‡	50.5‡	47.1‡	18.4‡	74.9‡	
mxQE(8)mxMBR	0.712‡	0.421‡	85.3‡	79.5‡	79.0‡	69.4‡	77.8‡	84.9‡	76.4‡	81.4‡	50.3‡	46.9‡	18.3‡	74.9‡	
mxQE(4)mxMBR	0.762‡	0.395‡	85.0‡	79.0‡	78.8‡	69.3‡	77.7‡	84.7‡	76.2‡	81.3‡	50.1‡	46.7‡	18.1‡	75.1‡	
ckQE(64)mxMBR	0.687‡	0.553‡	86.1‡	80.3‡	81.0‡	70.2‡	78.1‡	85.2	77.1‡	81.7‡	51.2‡	47.7‡	19.0‡	74.2‡	
ckQE(32)mxMBR	0.728‡	0.557‡	86.5‡	80.6‡	82.2‡	70.7‡	78.3‡	85.4‡	77.3	81.9‡	51.7‡	48.3‡	19.5‡	73.3‡	
ckQE(16)mxMBR	0.798‡	0.594‡	86.8‡	80.9‡	83.2‡	71.1‡	78.5‡	85.4‡	77.4	82.1‡	51.9‡	48.5‡	19.8‡	72.7‡	
ckQE(8)mxMBR	0.892‡	0.644‡	87.0‡	81.0‡	84.0‡	71.3‡	78.7‡	85.5‡	77.4	82.2‡	52.1‡	48.8‡	20.1‡	72.0‡	
ckQE(4)mxMBR	1.01‡	0.714‡	86.9‡	80.9‡	84.6‡	71.4‡	78.7‡	85.4‡	77.2‡	82.2‡	52.0‡	48.7‡	20.0‡	71.9‡	

Table 10: Reference-based and QE evaluation scores for greedy, MBR, and QE decoding using a two-step ensemble (QE filtering followed by MBR) utility metric, averaged across all languages (test datasets). Higher scores are better, except MetricX, MetricX-QE, and TER, where lower is better. Green is better than greedy, red is worse. Ensembles are defined in Table 2. Significant differences from greedy (pairwise t-test) indicated by * for $p < 0.05$, † for $p < 0.01$, ‡ for $p < 0.001$.

G Breakdown of Results on Individual Language Pairs

G.1 Results for English-Swahili (en-sw) on FLORES200 test dataset

MBR/QE Method	Evaluated Metric															
	MetricX	MetricX-QE	XCOMET-XXL	XCOMET-XL	CometKiwi23-XXL	CometKiwi23-XL	CometKiwi22	COMET22	AfriCOMET	AfriCOMET-QE	BLEURT	YiSi	chrF	chrF++	sentBLEU	TER
Greedy	1.70	1.28	83.7	85.0	84.6	73.2	83.8	86.0	85.7	76.9	77.5	86.3	65.0	62.6	34.9	51.7
MetricX	0.598‡	0.477‡	88.9‡	87.7‡	89.0‡	75.5‡	84.9‡	85.9	87.4‡	79.7‡	76.3‡	83.0‡	58.1‡	55.1‡	24.4‡	61.1*
MetricX-QE	0.811‡	0.293‡	87.5‡	86.6‡	88.4‡	75.0‡	84.7‡	85.2‡	86.6‡	79.3‡	75.3‡	82.7‡	57.2‡	54.1‡	23.8‡	61.4*
XCOMET-XXL	1.03‡	0.698‡	94.2‡	89.1‡	91.3‡	76.4‡	85.3‡	86.3*	87.6‡	79.5‡	76.9*	83.5‡	58.8‡	56.0‡	25.6‡	59.3
XCOMET-XL	1.08‡	0.788‡	89.9‡	92.4‡	89.9‡	77.9‡	85.4‡	86.4‡	87.9‡	79.5‡	77.8	83.7‡	59.3‡	56.4‡	26.1‡	58.5
CometKiwi23-XXL	1.23‡	0.784‡	90.6‡	88.3‡	93.6‡	76.9‡	85.5‡	86.0	87.1‡	79.4‡	76.6‡	83.5‡	58.7‡	55.8‡	25.6‡	58.8
CometKiwi23-XL	1.26‡	0.816‡	88.3‡	90.2‡	90.1‡	79.9‡	85.5‡	86.2	87.3‡	79.3‡	77.1	83.7‡	59.5‡	56.7‡	26.4‡	58.4
CometKiwi22	1.26‡	0.852‡	87.6‡	87.8‡	89.1‡	76.4‡	87.3‡	86.7‡	87.9‡	79.8‡	77.0	83.7‡	59.4‡	56.4‡	25.8‡	58.3
COMET22	1.25‡	0.927‡	87.5‡	88.0‡	88.4‡	75.9‡	85.4‡	88.3‡	87.8‡	79.7‡	78.5‡	85.2‡	62.8‡	60.0‡	30.2‡	52.7
AfriCOMET	1.10‡	0.769‡	88.7‡	88.4‡	89.2‡	75.9‡	85.7‡	86.7‡	90.0‡	80.7‡	77.5	84.2‡	60.7‡	57.8‡	27.5‡	56.3
AfriCOMET-QE	1.42‡	0.964‡	85.1‡	85.3	87.1‡	74.5‡	84.6‡	85.8	87.6‡	83.3‡	74.6‡	82.9‡	57.6‡	54.5‡	23.5‡	61.8*
BLEURT	1.37‡	1.05‡	86.3‡	86.9‡	87.3‡	75.0‡	84.8‡	86.2	86.7‡	78.5‡	82.9‡	84.0‡	60.0‡	57.0‡	25.8‡	58.0
YiSi	1.62	1.26	84.2	85.3	84.9	73.4	83.9	86.2*	85.7	76.9	77.9	86.9‡	65.7‡	63.2‡	35.1	46.7
chrF	1.57‡	1.23	84.7‡	85.7‡	85.5‡	74.0‡	84.1*	86.5‡	86.2‡	77.3‡	78.3‡	86.4	66.4‡	63.7‡	34.3*	49.3
chrF++	1.57‡	1.22	84.8‡	85.8‡	85.5‡	74.0‡	84.1*	86.5‡	86.2‡	77.2‡	78.4‡	86.5‡	66.4‡	63.9‡	34.7	48.8
sentBLEU	1.64	1.29	84.1	85.4	84.7	73.4	83.9	86.1	85.7	76.8	77.6	86.5*	65.3	63.0	35.8‡	46.3
TER	1.73	1.36*	83.2	84.2‡	83.7‡	72.7*	83.6*	85.8	85.0‡	76.3‡	77.3	86.2	64.4‡	62.0‡	34.6	45.1
rankAvg:all	1.01‡	0.711‡	89.8‡	89.7‡	90.1‡	77.1‡	85.9‡	87.5‡	88.5‡	79.8‡	79.4‡	85.9‡	64.5*	61.9‡	33.0‡	49.7
rankAvg:qe	0.893‡	0.506‡	90.3‡	89.8‡	91.9‡	78.3‡	86.5‡	86.9‡	88.5‡	81.3‡	77.4	83.8‡	59.7‡	56.8‡	26.2‡	58.6
rankAvg:top	0.781‡	0.484‡	92.2‡	90.9‡	92.0‡	78.4‡	85.8‡	86.7‡	88.2‡	80.0‡	77.8	83.8‡	59.8‡	57.0‡	26.7‡	58.3
rankAvg:topQe	0.900‡	0.455‡	90.7‡	89.8‡	92.3‡	78.7‡	85.7‡	86.4‡	87.9‡	79.9‡	77.4	83.5‡	59.0‡	56.0‡	25.7‡	59.4
rankAvg:mxmxqe	0.638‡	0.347‡	88.8‡	87.4‡	89.0‡	75.7‡	85.0‡	85.8	87.3‡	79.7‡	76.2‡	83.0‡	58.2‡	55.1‡	24.4‡	60.8*
rankAvg:noLex	0.899‡	0.606‡	91.2‡	90.5‡	91.0‡	77.7‡	86.2‡	87.5‡	88.9‡	80.4‡	79.7‡	85.2‡	62.7‡	60.0‡	30.2‡	53.4
rankAvg:noNC	1.06‡	0.724‡	88.0‡	88.4‡	88.4‡	75.8‡	85.3‡	87.4‡	88.3‡	79.7‡	79.5‡	86.2	64.8	62.3	33.7‡	48.6
rankAvg:noNCnoLex	0.919‡	0.597‡	89.2‡	89.0‡	89.4‡	76.5‡	85.7‡	87.6‡	88.7‡	80.6‡	80.1‡	85.7‡	63.6‡	60.9‡	31.4‡	51.8
allQE(32)allMBR	0.992‡	0.705‡	90.2‡	89.8‡	90.2‡	77.1‡	85.8‡	87.4‡	88.5‡	79.8‡	79.4‡	85.8‡	64.1‡	61.5‡	32.6‡	50.4
allQE(32)nolexMBR	0.904‡	0.636‡	91.2‡	90.6‡	90.7‡	77.5‡	86.0‡	87.5‡	88.8‡	80.1‡	79.5‡	85.2‡	62.7‡	60.0‡	30.6‡	53.3
topQE(32)topMBR	0.761‡	0.552‡	92.4‡	90.9‡	91.2‡	77.7‡	85.7‡	86.6‡	88.1‡	79.9‡	78.0	83.7‡	59.6‡	56.8‡	26.9‡	58.1
noncQE(32)noncMBR	0.968‡	0.648‡	88.7‡	88.7‡	89.0‡	76.1‡	85.5‡	87.3‡	88.5‡	80.1‡	79.4‡	85.8‡	64.3‡	61.7‡	32.8‡	49.9
noncQE(32)noncnolexMBR	0.885‡	0.604‡	89.3‡	89.1‡	89.5‡	76.5‡	85.7‡	87.5‡	88.8‡	80.4‡	79.7‡	85.4‡	63.1‡	60.3‡	30.7‡	52.7
mxQE(32)mxMBR	0.628‡	0.434‡	89.2‡	87.7‡	88.9‡	75.6‡	85.0‡	86.0	87.4‡	79.7‡	76.5‡	83.2‡	58.5‡	55.5‡	25.0‡	60.3*
ckQE(32)xcMBR	1.05‡	0.696‡	94.0‡	89.1‡	91.8‡	76.6‡	85.4‡	86.3*	87.6‡	79.6‡	77.2	83.6‡	59.1‡	56.2‡	25.9‡	59.1
mxQE(32)xcMBR	0.928‡	0.538‡	93.6‡	89.0‡	91.2‡	76.5‡	85.4‡	86.2	87.6‡	79.7‡	77.0*	83.5‡	58.7‡	55.8‡	25.4‡	59.4
ckQE(32)mxMBR	0.657‡	0.499‡	90.0‡	88.3‡	90.9‡	76.2‡	85.3‡	86.2	87.5‡	79.8‡	76.8*	83.4‡	58.7‡	55.8‡	25.4‡	59.6

Table 11: Reference-based and QE evaluation scores for greedy and MBR/QE decoding (1st block), and ensembles (2nd block), on en-sw (FLORES200 test dataset). Higher scores are better, except MetricX, MetricX-QE, and TER, where lower is better. Green is better than greedy, red is worse. Ensembles are defined in Table 2. Significant differences from greedy (pairwise t-test) indicated by * for p<0.05, † for p<0.01, ‡ for p<0.001. The green diagonal in the 1st block shows metrics prefer outputs from MBR/QE decoding using the same utility metric.

G.2 Results for English-Hausa (en-ha) on FLORES200 test dataset

MBR/QE Method	Evaluated Metric															
	MetricX	MetricX-QE	XCOMET-XXL	XCOMET-XL	CometKiwi23-XXL	CometKiwi23-XL	CometKiwi22	COMET22	AfriCOMET	AfriCOMET-QE	BLEURT	YiSi	chrF	chrF++	sentBLEU	TER
Greedy	2.31	1.46	74.7	70.3	79.9	66.2	60.1	81.5	79.3	72.3	84.1	79.5	52.5	49.8	21.3	65.9
MetricX	0.818 ‡	0.515 ‡	76.7 ‡	70.5	82.8 ‡	67.5 ‡	60.0	81.5	81.3 ‡	76.2 ‡	81.2 ‡	77.0 ‡	47.5 ‡	44.3 ‡	14.3 ‡	80.2 ‡
MetricX-QE	1.19 ‡	0.278 ‡	76.5 ‡	69.6	82.4 ‡	67.1 ‡	59.5	81.0 ‡	80.7 ‡	75.6 ‡	80.0 ‡	76.9 ‡	46.6 ‡	43.3 ‡	13.7 ‡	78.7 ‡
XCOMET-XXL	1.65 ‡	0.909 ‡	87.5 ‡	73.9 ‡	86.9 ‡	69.8 ‡	61.4 ‡	82.5 ‡	81.5 ‡	75.4 ‡	82.6 ‡	77.7 ‡	48.6 ‡	45.5 ‡	15.3 ‡	75.9 ‡
XCOMET-XL	1.75 ‡	0.990 ‡	80.9 ‡	80.6 ‡	85.1 ‡	71.8 ‡	62.8 ‡	82.5 ‡	81.7 ‡	75.4 ‡	83.2 ‡	77.8 ‡	48.8 ‡	45.8 ‡	15.9 ‡	74.2 ‡
CometKiwi23-XXL	1.97 ‡	0.989 ‡	82.2 ‡	72.9 ‡	90.2 ‡	70.5 ‡	62.2 ‡	82.2 ‡	81.1 ‡	75.6 ‡	81.8 ‡	77.6 ‡	48.0 ‡	45.0 ‡	15.4 ‡	75.4 ‡
CometKiwi23-XL	1.99 ‡	1.12 ‡	78.5 ‡	75.4 ‡	85.4 ‡	74.5 ‡	62.6 ‡	82.3 ‡	80.6 ‡	74.9 ‡	82.0 ‡	77.7 ‡	48.3 ‡	45.3 ‡	15.5 ‡	76.1 ‡
CometKiwi22	2.55 †	1.51	73.5 *	69.0 †	81.2 ‡	67.8 ‡	74.8 ‡	82.0 ‡	78.8 *	73.1 ‡	81.6 ‡	77.5 ‡	47.5 ‡	44.4 ‡	14.8 ‡	75.6 ‡
COMET22	1.94 ‡	1.14 ‡	77.7 ‡	72.4 ‡	84.0 ‡	69.0 ‡	63.4 ‡	84.9 ‡	81.6 ‡	75.9 ‡	83.6 ‡	78.4 ‡	50.0 ‡	46.9 ‡	16.4 ‡	74.3 ‡
AfriCOMET	1.71 ‡	0.956 ‡	77.8 ‡	71.3 *	83.2 ‡	68.0 ‡	60.2	82.0 *	85.4 ‡	77.5 ‡	81.8 ‡	77.8 ‡	48.9 ‡	45.7 ‡	15.7 ‡	75.6 ‡
AfriCOMET-QE	1.93 ‡	1.07 ‡	73.4 *	67.6 ‡	81.3 ‡	66.5	59.0 ‡	81.3	82.2 ‡	80.3 ‡	78.9 ‡	76.8 ‡	46.6 ‡	43.3 ‡	13.4 ‡	81.1 ‡
BLEURT	2.10 ‡	1.34 *	75.8 *	71.2 *	81.3 ‡	67.5 ‡	61.4 ‡	82.2 ‡	80.2 ‡	73.7 ‡	89.8 ‡	78.3 ‡	49.9 ‡	46.9 ‡	16.8 ‡	73.3 ‡
YiSi	2.37 ‡	1.46	73.7 *	69.2 ‡	79.4	66.2	60.5	81.7	79.1	72.5	83.6	80.4 ‡	52.3	49.6	20.2 ‡	66.6
chrF	2.30	1.40	74.7	69.7	80.3	67.0 ‡	60.6 *	82.1 ‡	80.0 ‡	73.5 ‡	84.3	79.4	53.7 ‡	50.6 ‡	19.4‡	71.9‡
chrF++	2.34	1.43	74.4	69.9	80.2	66.9 ‡	60.7 *	82.1 ‡	79.9 ‡	73.4 ‡	84.4	79.5	53.5 ‡	50.6 ‡	19.8‡	71.4‡
sentBLEU	2.36	1.50	73.6 †	69.8	78.8 ‡	65.7	60.1	81.5	79.0	72.3	83.7	79.6	52.3	49.7	21.2	65.1 *
TER	2.66 ‡	1.69 ‡	72.8 ‡	68.8 ‡	77.3 ‡	64.4 ‡	59.6 *	80.8 ‡	77.9 ‡	71.3 ‡	82.9 ‡	79.6	51.3 ‡	48.8 ‡	21.1	61.3 ‡
rankAvg:all	1.47 ‡	0.782 ‡	81.8 ‡	75.8 ‡	85.9 ‡	70.8 ‡	64.8 ‡	83.6 ‡	82.9 ‡	76.1 ‡	85.7 ‡	79.3	51.9 *	49.0 ‡	19.3 ‡	68.8 ‡
rankAvg:qe	1.40 ‡	0.581 ‡	81.1 ‡	75.1 ‡	87.5 ‡	71.9 ‡	68.1 ‡	83.0 ‡	82.7 ‡	77.5 ‡	82.8 ‡	77.7 ‡	48.9 ‡	45.7 ‡	15.5 ‡	76.7 ‡
rankAvg:top	1.15 ‡	0.524 ‡	84.1 ‡	77.5 ‡	88.0 ‡	72.4 ‡	62.4 ‡	82.6 ‡	82.4 ‡	76.3 ‡	83.3 *	77.8 ‡	49.1 ‡	46.0 ‡	16.1 ‡	76.0 ‡
rankAvg:topQe	1.36 ‡	0.491 ‡	81.8 ‡	75.8 ‡	88.4 ‡	72.8 ‡	62.7 ‡	82.5 ‡	82.0 ‡	76.0 ‡	82.5 ‡	77.6 ‡	48.4 ‡	45.4 ‡	15.7 ‡	75.9 ‡
rankAvg:mxmxqe	0.891 ‡	0.336 ‡	77.4 ‡	70.9	82.8 ‡	67.6 ‡	60.1	81.4	81.5 ‡	76.1 ‡	81.1 ‡	77.0 ‡	47.3 ‡	44.1 ‡	14.2 ‡	79.4 ‡
rankAvg:noLex	1.34 ‡	0.657 ‡	82.8 ‡	76.4 ‡	86.9 ‡	71.6 ‡	65.8 ‡	83.7 ‡	83.3 ‡	76.9 ‡	85.1 ‡	78.7 ‡	50.5 ‡	47.5 ‡	17.5 ‡	72.8 ‡
rankAvg:noNC	1.43 ‡	0.736 ‡	79.4 ‡	73.4 ‡	83.6 ‡	68.9 ‡	62.0 ‡	83.4 ‡	82.9 ‡	76.2 ‡	85.5 ‡	79.5	52.1	49.2 †	19.8 ‡	68.0 ‡
rankAvg:noNCnoLex	1.27 ‡	0.625 ‡	80.1 ‡	73.9 ‡	84.4 ‡	69.3 ‡	62.3 ‡	83.7 ‡	83.5 ‡	77.2 ‡	85.5 ‡	79.0 ‡	51.0 ‡	48.0 ‡	17.9 ‡	72.4 ‡
allQE(32)allMBR	1.45 ‡	0.802 ‡	82.0 ‡	76.1 ‡	85.8 ‡	70.6 ‡	63.7 ‡	83.6 ‡	83.0 ‡	76.1 ‡	85.6 ‡	79.1 ‡	51.6 ‡	48.7 ‡	19.1 ‡	70.0 ‡
allQE(32)noLexMBR	1.31 ‡	0.715 ‡	83.4 ‡	76.9 ‡	86.4 ‡	71.2 ‡	63.7 ‡	83.7 ‡	83.4 ‡	76.6 ‡	85.3 ‡	78.6 ‡	50.5 ‡	47.6 ‡	17.7 ‡	73.1 ‡
topQE(32)topMBR	1.12 ‡	0.595 ‡	84.6 ‡	78.0 ‡	87.0 ‡	71.6 ‡	62.4 ‡	82.6 ‡	82.4 ‡	76.1 ‡	83.5	77.8 ‡	49.1 ‡	46.1 ‡	16.2 ‡	75.6 ‡
noncQE(32)noncMBR	1.35 ‡	0.704 ‡	79.4 ‡	73.6 ‡	83.7 ‡	69.0 ‡	61.9 ‡	83.4 ‡	83.0 ‡	76.5 ‡	84.9 *	79.0 ‡	51.5 ‡	48.5 ‡	18.7 ‡	70.2 ‡
noncQE(32)noLexMBR	1.22 ‡	0.653 ‡	80.4 ‡	74.3 ‡	84.6 ‡	69.3 ‡	62.3 ‡	83.7 ‡	83.6 ‡	76.9 ‡	85.1 †	78.8 ‡	50.6 ‡	47.6 ‡	17.7 ‡	73.0 ‡
mxQE(32)mxMBR	0.859 ‡	0.444 ‡	76.7 ‡	70.7	82.8 ‡	67.6 ‡	60.3	81.6	81.3 ‡	76.1 ‡	81.1 ‡	77.1 ‡	47.7 ‡	44.5 ‡	14.6 ‡	79.5 ‡
ckQE(32)xcMBR	1.60 ‡	0.865 ‡	87.2 ‡	74.2 ‡	87.7 ‡	70.1 ‡	61.3 ‡	82.4 ‡	81.6 ‡	75.7 ‡	82.4 ‡	77.7 ‡	48.6 ‡	45.5 ‡	15.3 ‡	76.6 ‡
mxQE(32)xcMBR	1.39 ‡	0.613 ‡	86.4 ‡	73.9 ‡	86.9 ‡	69.9 ‡	61.5 ‡	82.3 ‡	81.9 ‡	75.8 ‡	82.2 ‡	77.6 ‡	48.5 ‡	45.3 ‡	15.3 ‡	76.8 ‡
ckQE(32)mxMBR	0.908 ‡	0.516 ‡	79.6 ‡	72.8 ‡	86.6 ‡	69.3 ‡	61.1 ‡	82.1 ‡	82.0 ‡	76.3 ‡	81.7 ‡	77.3 ‡	48.0 ‡	44.9 ‡	14.7 ‡	78.9 ‡

Table 12: Reference-based and QE evaluation scores for greedy and MBR/QE decoding (1st block), and ensembles (2nd block), on en-ha (FLORES200 test dataset). Higher scores are better, except MetricX, MetricX-QE, and TER, where lower is better. Green is better than greedy, red is worse. Ensembles are defined in Table 2. Significant differences from greedy (pairwise t-test) indicated by * for p<0.05, † for p<0.01, ‡ for p<0.001. The green diagonal in the 1st block shows metrics prefer outputs from MBR/QE decoding using the same utility metric.

G.3 Results for English-Igbo (en-ig) on FLORES200 test dataset

MBR/QE Method	Evaluated Metric															
	MetricX	MetricX-QE	XCOMET-XXL	XCOMET-XL	CometKiwi23-XXL	CometKiwi23-XL	CometKiwi22	COMET22	AfriCOMET	AfriCOMET-QE	BLEURT	YiSi	chrF	chrF++	sentBLEU	TER
Greedy	3.73	2.67	27.3	19.8	16.9	16.8	30.2	72.1	68.9	66.4	33.5	78.3	42.7	40.7	15.2	74.2
MetricX	1.68‡	1.31‡	26.7‡	18.9‡	18.0‡	17.8‡	30.3	72.1	71.7‡	69.7‡	28.8‡	77.2‡	40.7‡	38.6‡	12.8‡	79.6‡
MetricX-QE	2.15‡	0.937‡	26.7‡	19.2*	17.4	17.7‡	30.6*	72.1	71.2‡	69.3‡	29.0‡	77.3‡	40.7‡	38.5‡	12.8‡	78.8‡
XCOMET-XXL	5.12‡	3.55‡	31.4‡	21.0‡	20.3‡	18.9‡	32.1‡	72.2	66.5‡	65.6‡	31.2‡	77.5‡	40.1‡	38.0‡	12.8‡	74.1
XCOMET-XL	6.81‡	5.12‡	27.2	26.7‡	20.0‡	19.6‡	33.1‡	69.1‡	62.6‡	64.3‡	29.2‡	75.9‡	37.2‡	35.2‡	11.8‡	79.5‡
CometKiwi23-XXL	6.08‡	4.17‡	26.4‡	20.1	36.6‡	21.7‡	32.3‡	70.2‡	64.5‡	64.9‡	27.1‡	76.0‡	37.9‡	35.8‡	11.4‡	83.2‡
CometKiwi23-XL	5.94‡	3.99‡	26.2‡	20.7‡	24.8‡	29.3‡	32.6‡	70.2‡	64.5‡	65.5‡	27.4‡	75.8‡	38.0‡	35.7‡	10.9‡	83.6‡
CometKiwi22	6.13‡	4.33‡	26.7‡	21.9‡	21.0‡	20.1‡	40.6‡	69.6‡	63.9‡	65.3‡	28.8‡	75.9‡	37.0‡	35.0‡	11.3‡	77.6‡
COMET22	5.28‡	3.80‡	27.2	18.5‡	17.4	17.3‡	30.7‡	76.3‡	66.7‡	66.1	29.1‡	77.7‡	41.0‡	38.9‡	13.0‡	78.9‡
AfriCOMET	3.26‡	2.26‡	26.8‡	19.0‡	18.4‡	17.6‡	30.6‡	72.3	77.3‡	72.6‡	28.8‡	77.4‡	40.9‡	38.6‡	12.7‡	79.7‡
AfriCOMET-QE	4.16‡	2.87*	26.4‡	19.0‡	17.9‡	18.0‡	31.0‡	71.8	73.4‡	75.3‡	24.7‡	76.6‡	38.8‡	36.6‡	11.4‡	81.8‡
BLEURT	4.20‡	3.05‡	27.2	19.9	19.0‡	17.6‡	30.7‡	71.9	69.0	66.3	40.9‡	77.8‡	41.8‡	39.8‡	14.1‡	74.9
YiSi	4.03‡	2.87‡	27.9‡	19.2*	17.1	17.3‡	30.0	73.2‡	69.6‡	66.9‡	33.9	79.2‡	43.5‡	41.5‡	15.3	73.7
chrF	4.01‡	2.86‡	27.3	18.8‡	17.7‡	17.8‡	29.9‡	73.0‡	69.9‡	66.9‡	33.1	78.5‡	44.5‡	42.3‡	15.3	77.2‡
chrF++	3.95‡	2.84‡	27.3	18.9‡	17.6‡	17.7‡	30.0*	72.8‡	69.7‡	66.8*	33.7	78.5‡	44.4‡	42.3‡	15.4	76.8‡
sentBLEU	4.05‡	2.94‡	27.7*	20.0	18.2‡	17.5‡	30.9‡	72.2	68.8	66.1	34.6‡	78.3	42.9	41.1*	15.8‡	72.4‡
TER	4.30‡	3.07‡	28.2‡	20.7‡	18.1‡	17.2*	31.5‡	71.7‡	68.2‡	65.9‡	34.8‡	78.2	42.1‡	40.3*	15.3	69.5‡
rankAvg:all	3.05‡	2.05‡	28.6‡	20.8‡	22.1‡	20.6‡	32.5‡	73.6‡	72.4‡	69.3‡	34.8‡	78.5*	43.2‡	41.1*	15.4	73.0
rankAvg:qe	3.34‡	1.94‡	27.5	21.1‡	27.2‡	24.1‡	35.1‡	71.6‡	71.7‡	71.2‡	29.5‡	76.9‡	40.2‡	37.9‡	12.3‡	78.2‡
rankAvg:top	2.86‡	1.77‡	29.1‡	22.0‡	26.2‡	23.1‡	32.5‡	72.0	70.4‡	68.4‡	31.6‡	77.4‡	41.1‡	38.9‡	13.4‡	75.5*
rankAvg:topQe	3.17‡	1.70‡	27.2	20.4*	28.9‡	24.8‡	32.0‡	71.6‡	69.8‡	68.4‡	29.4‡	76.8‡	40.4‡	38.1‡	12.4‡	80.1‡
rankAvg:mxmxqe	1.83‡	1.03‡	26.8‡	19.2*	17.9‡	17.9‡	30.4	72.1	71.9‡	69.6‡	29.5‡	77.3‡	41.0‡	38.8‡	12.9‡	79.0‡
rankAvg:noLex	2.98‡	1.93‡	28.8‡	21.2‡	23.4‡	21.3‡	33.1‡	73.8‡	72.9‡	70.0‡	34.1	78.3	42.7	40.5	14.8	73.7
rankAvg:noNC	2.74‡	1.85‡	28.0‡	19.7	18.5‡	18.1‡	30.8‡	73.6‡	73.1‡	69.7‡	34.6‡	78.6‡	43.4‡	41.3‡	15.4	73.4
rankAvg:noNCnoLex	2.60‡	1.67‡	27.7*	19.5	17.9‡	18.0‡	30.7‡	74.1‡	73.9‡	70.8‡	34.6*	78.6‡	42.9	40.7	14.7*	74.9
allQE(32)allMBR	3.06‡	2.06‡	28.6‡	20.9‡	21.9‡	20.4‡	32.5‡	73.2‡	72.3‡	69.3‡	34.3	78.3	42.9	40.8	15.1	73.3
allQE(32)nolexMBR	2.90‡	1.97‡	28.8‡	21.3‡	22.1‡	20.3‡	32.5‡	73.5‡	72.9‡	69.7‡	34.6*	78.3	42.7	40.5	14.7*	73.6
topQE(32)topMBR	2.71‡	1.80‡	29.3‡	22.3‡	23.5‡	21.5‡	32.6‡	71.9	70.6‡	68.4‡	32.3‡	77.6‡	41.3‡	39.1‡	13.7‡	74.3
noncQE(32)noncMBR	2.64‡	1.76‡	27.8‡	19.7	18.4‡	18.0‡	30.7‡	73.3‡	73.6‡	70.1‡	34.2	78.6‡	43.4‡	41.3‡	15.1	73.9
noncQE(32)noncnolexMBR	2.51‡	1.69‡	27.9‡	19.6	18.1‡	17.9‡	30.7‡	73.9‡	74.2‡	70.6‡	34.3	78.5‡	43.0	40.8	14.6*	75.0
mxQE(32)mxMBR	1.73‡	1.24‡	26.7‡	19.0‡	18.0‡	17.6‡	30.4	72.1	71.8‡	69.6‡	29.4‡	77.3‡	40.9‡	38.8‡	13.0‡	78.9‡
ckQE(32)xcMBR	5.07‡	3.53‡	30.4‡	21.3‡	26.6‡	20.4‡	32.4‡	71.9	66.6‡	66.0*	30.2‡	77.2‡	39.8‡	37.7‡	12.7‡	75.8*
mxQE(32)xcMBR	3.41‡	1.90‡	30.5‡	21.0‡	20.6‡	18.7‡	32.0‡	72.5*	69.3	67.6‡	31.6‡	77.7‡	41.0‡	38.8‡	13.3‡	74.2
ckQE(32)mxMBR	2.11‡	1.56‡	27.0	19.4	25.4‡	19.6‡	30.9‡	72.0	71.2‡	69.2‡	29.0‡	77.1‡	40.7‡	38.5‡	12.7‡	80.2‡

Table 13: Reference-based and QE evaluation scores for greedy and MBR/QE decoding (1st block), and ensembles (2nd block), on en-ig (FLORES200 test dataset). Higher scores are better, except MetricX, MetricX-QE, and TER, where lower is better. Green is better than greedy, red is worse. Ensembles are defined in Table 2. Significant differences from greedy (pairwise t-test) indicated by * for p<0.05, † for p<0.01, ‡ for p<0.001. The green diagonal in the 1st block shows metrics prefer outputs from MBR/QE decoding using the same utility metric.

G.4 Results for English-Somali (en-so) on FLORES200 test dataset

MBR/QE Method	Evaluated Metric															
	MetricX	MetricX-QE	XCOMET-XXL	XCOMET-XL	CometKiwi23-XXL	CometKiwi23-XL	CometKiwi22	COMET22	AfriCOMET	AfriCOMET-QE	BLEURT	YiSi	chrF	chrF++	sentBLEU	TER
Greedy	2.66	1.89	67.9	66.0	78.3	65.8	70.0	80.7	75.9	73.4	108.	77.7	46.7	42.4	11.3	86.7
MetricX	0.996‡	0.635‡	71.5‡	69.2‡	83.8‡	69.5‡	70.5*	81.9‡	79.8‡	78.4‡	110.‡	76.2‡	44.2‡	39.5‡	8.89‡	91.9
MetricX-QE	1.34‡	0.396‡	70.6‡	68.3‡	83.0‡	68.6‡	70.3	81.5‡	79.0‡	77.8‡	109.‡	76.3‡	43.8‡	39.2‡	8.94‡	89.2
XCOMET-XXL	1.69‡	0.998‡	81.0‡	72.3‡	87.2‡	70.5‡	71.5‡	82.1‡	79.8‡	78.0‡	110.‡	76.5‡	44.2‡	39.7‡	9.14‡	89.0
XCOMET-XL	1.82‡	1.08‡	74.7‡	78.7‡	85.4‡	72.7‡	72.6‡	82.5‡	80.1‡	77.7‡	111.‡	76.8‡	45.0‡	40.5‡	9.85‡	88.2
CometKiwi23-XXL	1.89‡	1.09‡	75.6‡	71.6‡	90.4‡	71.4‡	72.3‡	82.1‡	79.6‡	78.0‡	110.‡	76.4‡	44.1‡	39.6‡	9.34‡	89.0
CometKiwi23-XL	2.00‡	1.17‡	72.8‡	74.6‡	85.9‡	75.1‡	72.7‡	82.2‡	79.5‡	77.6‡	110.‡	76.5‡	44.4‡	39.9‡	9.64‡	89.8
CometKiwi22	2.56	1.65†	68.0	66.9*	81.4	68.4‡	80.3‡	81.7‡	77.1‡	75.7‡	109.*	76.7‡	44.3‡	39.7‡	9.49‡	87.8
COMET22	2.01‡	1.30‡	72.1‡	69.9‡	83.5‡	69.6‡	73.0‡	84.6‡	79.8‡	78.3‡	111.‡	77.2‡	45.8‡	41.2‡	9.93‡	87.7
AfriCOMET	1.78‡	1.08‡	73.7‡	71.2‡	84.7‡	70.0‡	71.5‡	82.4‡	83.9‡	80.2‡	110.‡	76.8‡	45.1‡	40.5‡	9.66‡	88.5
AfriCOMET-QE	2.01‡	1.25‡	69.3‡	67.3‡	82.1‡	68.6‡	70.2	82.2‡	81.1‡	82.6‡	108.	76.1‡	43.9‡	39.1‡	8.55‡	92.1
BLEURT	2.23‡	1.52‡	68.4	67.5‡	81.7‡	68.2‡	71.0‡	81.5‡	77.6‡	76.6‡	120.‡	76.0‡	43.9‡	39.2‡	8.29‡	96.9‡
YiSi	2.66	1.71*	68.0	65.8	79.4‡	66.4*	70.5*	81.3‡	76.6‡	74.3‡	108.	78.6‡	47.2*	42.7	11.2	82.9
chrF	2.48*	1.66†	68.2	66.4	80.1‡	67.1‡	70.3	81.5‡	77.4‡	74.9‡	109.‡	77.9	48.3‡	43.6‡	11.2	87.1
chrF++	2.52	1.67†	68.5	66.5	79.9‡	66.9‡	70.4*	81.4‡	77.1‡	74.6‡	109.‡	77.9‡	48.2‡	43.7‡	11.2	87.1
sentBLEU	2.65	1.76	67.9	66.4	79.0	66.2	69.8	80.8	76.3	73.7	108.	77.8	46.9	42.6	11.8†	81.9
TER	2.85*	1.95	67.9	65.2	77.4*	65.1†	69.8	80.5	75.3*	73.0*	106.‡	77.7	45.9‡	41.7†	11.6	77.1‡
rankAvg:all	1.63‡	0.909‡	75.8‡	73.5‡	85.9‡	71.4‡	74.2‡	83.1‡	81.1‡	78.4‡	112.‡	77.8	47.0	42.5	11.0	84.1
rankAvg:qe	1.55‡	0.724‡	74.9‡	73.5‡	87.8‡	72.7‡	76.4‡	82.9‡	81.1‡	80.0‡	111.‡	76.7‡	44.7‡	40.1‡	9.39‡	89.5
rankAvg:top	1.34‡	0.658‡	77.7‡	75.6‡	88.1‡	73.1‡	72.5‡	82.7‡	81.0‡	78.7‡	112.‡	76.6‡	44.8‡	40.3‡	9.54‡	89.5
rankAvg:topQe	1.42‡	0.623‡	75.8‡	74.0‡	88.6‡	73.4‡	72.3‡	82.6‡	80.6‡	78.6‡	111.‡	76.6‡	44.7‡	40.2‡	9.60‡	89.6
rankAvg:mxmxqe	1.08‡	0.458‡	71.8‡	69.4‡	83.9‡	69.4‡	70.7†	81.9‡	79.7‡	78.3‡	110.‡	76.3‡	44.1‡	39.4‡	8.78‡	90.9
rankAvg:noLex	1.45‡	0.786‡	76.8‡	74.7‡	87.1‡	72.3‡	75.0‡	83.3‡	81.7‡	79.3‡	113.‡	77.4*	46.0†	41.4‡	10.1‡	87.7
rankAvg:noNC	1.66‡	0.928‡	73.5‡	71.1‡	83.9‡	69.8‡	72.3‡	82.9‡	80.8‡	78.3‡	112.‡	77.9	47.2*	42.7	11.2	83.8
rankAvg:noNCnoLex	1.41‡	0.764‡	74.4‡	71.9‡	85.1‡	70.5‡	72.5‡	83.3‡	81.7‡	79.5‡	114.‡	77.6	46.4	41.8*	10.2‡	87.9
allQE(32)allMBR	1.55‡	0.903‡	76.1‡	74.0‡	85.8‡	71.4‡	73.2‡	83.1‡	81.2‡	78.4‡	112.‡	77.5	46.6	42.1	10.7†	85.7
allQE(32)noLexMBR	1.43‡	0.841‡	77.4‡	75.0‡	86.8‡	71.9‡	73.5‡	83.3‡	81.9‡	79.1‡	114.‡	77.3‡	46.1†	41.5‡	10.1‡	87.9
topQE(32)topMBR	1.25‡	0.740‡	78.4‡	76.1‡	87.2‡	72.3‡	72.2‡	82.5‡	80.9‡	78.5‡	111.‡	76.6‡	44.7‡	40.2‡	9.68‡	89.4
noncQE(32)noncMBR	1.54‡	0.887‡	74.4‡	71.7‡	84.5‡	70.1‡	72.3‡	83.0‡	81.2‡	78.8‡	112.‡	77.6	46.7	42.2	10.7†	85.5
noncQE(32)noLexMBR	1.41‡	0.809‡	74.7‡	72.1‡	85.2‡	70.7‡	72.7‡	83.3‡	81.9‡	79.3‡	114.‡	77.4†	46.2*	41.6‡	9.95‡	88.4
mxQE(32)mxMBR	1.03‡	0.590‡	71.6‡	69.1‡	83.9‡	69.3‡	70.5*	81.9‡	79.9‡	78.3‡	110.‡	76.3‡	44.3‡	39.6‡	9.03‡	91.2
ckQE(32)xcMBR	1.68‡	0.974‡	80.4‡	72.2‡	88.0‡	70.9‡	71.8‡	82.2‡	79.9‡	78.1‡	110.‡	76.3‡	44.0‡	39.6‡	9.26‡	89.4
mxQE(32)xcMBR	1.51‡	0.770‡	80.3‡	72.4‡	87.1‡	70.7‡	71.5‡	82.2‡	80.1‡	78.2‡	110.‡	76.5‡	44.4‡	39.9‡	9.40‡	89.3
ckQE(32)mxMBR	1.09‡	0.650‡	73.8‡	70.6‡	86.6‡	70.3‡	71.1‡	82.1‡	80.2‡	78.6‡	111.‡	76.4‡	44.4‡	39.8‡	8.90‡	91.4

Table 14: Reference-based and QE evaluation scores for greedy and MBR/QE decoding (1st block), and ensembles (2nd block), on en-so (FLORES200 test dataset). Higher scores are better, except MetricX, MetricX-QE, and TER, where lower is better. Green is better than greedy, red is worse. Ensembles are defined in Table 2. Significant differences from greedy (pairwise t-test) indicated by * for p<0.05, † for p<0.01, ‡ for p<0.001. The green diagonal in the 1st block shows metrics prefer outputs from MBR/QE decoding using the same utility metric.

G.5 Results for English-Hindi (en-hi) on FLORES200 test dataset

MBR/QE Method	Evaluated Metric														
	MetricX	MetricX-QE	XCOMET-XXL	XCOMET-XL	CometKiwi23-XXL	CometKiwi23-XL	CometKiwi22	COMET22	IndicCOMET	BLEURT	YiSi	chrF	chrF++	sentBLEU	TER
Greedy	0.815	0.455	84.9	75.9	77.5	68.7	85.5	82.4	80.5	74.3	86.6	59.8	57.8	32.9	51.4
MetricX	0.257‡	0.0928‡	92.1‡	81.4‡	78.7‡	68.9	86.0‡	82.0†	80.8	73.1‡	83.1‡	50.9‡	48.7‡	21.5‡	66.8‡
MetricX-QE	0.463‡	0.0413‡	89.6‡	77.3‡	76.9*	67.6‡	85.6	81.1‡	79.8‡	72.1‡	82.9‡	49.4‡	47.2‡	20.7‡	65.9‡
XCOMET-XXL	0.462‡	0.158‡	96.4‡	80.5‡	79.5‡	68.5	85.9‡	81.6‡	80.2	72.8‡	83.0‡	49.8‡	47.6‡	20.8‡	65.5‡
XCOMET-XL	0.455‡	0.176‡	91.9‡	87.7‡	79.4‡	70.1‡	86.3‡	82.6	81.2‡	73.9†	83.8‡	52.3‡	50.2‡	23.5‡	62.8‡
CometKiwi23-XXL	0.531‡	0.195‡	91.7‡	80.7‡	84.3‡	70.4‡	86.4‡	82.3	80.6	73.7‡	84.2‡	53.0‡	50.8‡	23.8‡	61.7‡
CometKiwi23-XL	0.590‡	0.238‡	89.4‡	81.4‡	80.3‡	73.3‡	86.4‡	82.4	80.4	73.5‡	84.3‡	53.1‡	50.9‡	23.8‡	61.7‡
CometKiwi22	0.581‡	0.202‡	90.1‡	80.5‡	79.7‡	69.9‡	87.6‡	82.8†	81.1‡	74.0*	83.9‡	52.5‡	50.4‡	23.2‡	62.8‡
COMET22	0.563‡	0.244‡	89.3‡	80.4‡	79.2‡	69.8‡	86.5‡	84.7‡	81.9‡	75.0‡	85.6‡	57.1‡	55.0‡	28.9‡	55.6‡
IndicCOMET	0.641‡	0.306‡	88.7‡	78.4‡	77.8	68.4*	86.0‡	82.6	85.4‡	73.7‡	84.4‡	57.7‡	51.5‡	24.8‡	59.7‡
BLEURT	0.579‡	0.259‡	89.1‡	80.1‡	79.0‡	69.4‡	86.3‡	83.1‡	81.2‡	76.6‡	85.3‡	56.5‡	54.5‡	28.3‡	55.7‡
YiSi	0.772*	0.403†	85.7*	76.4	77.9*	69.0*	85.8‡	82.8‡	80.6	74.8‡	86.9‡	59.8	57.8	32.7	51.4
chrF	0.746‡	0.397‡	85.9†	76.6*	78.0‡	69.1‡	85.8‡	82.9‡	80.8*	74.8‡	86.8‡	60.6‡	58.5‡	32.6	52.4†
chrF++	0.752†	0.404†	85.9†	76.6†	78.0‡	69.0†	85.8‡	82.9‡	80.7	74.9‡	86.8‡	60.6‡	58.7‡	33.1	52.0
sentBLEU	0.779	0.419*	85.2	76.1	77.4	68.6	85.5	82.6*	80.5	74.5	86.6	59.8	57.9	33.1	50.9
TER	0.803	0.431	85.4	75.8	77.0†	68.1‡	85.3*	82.4	80.4	74.3	86.5	58.8‡	56.9‡	32.1†	50.0‡
rankAvg:all	0.477‡	0.175‡	91.4‡	82.2‡	80.7‡	70.7‡	86.7‡	83.9‡	82.4‡	75.5‡	86.3‡	58.7‡	56.7‡	30.7‡	53.1‡
rankAvg:qe	0.442‡	0.0972‡	91.8‡	83.1‡	82.5‡	72.0‡	87.1‡	83.0‡	81.4‡	74.3	84.4‡	53.8‡	51.5‡	24.3‡	61.5‡
rankAvg:top	0.349‡	0.0980‡	94.2‡	85.0‡	82.3‡	71.7‡	86.7‡	82.9‡	81.5‡	74.2	84.2‡	53.4‡	51.1‡	23.9‡	62.4‡
rankAvg:topQe	0.442‡	0.0875‡	92.1‡	82.6‡	82.9‡	72.2‡	86.6‡	82.7*	81.2‡	73.9*	84.3‡	53.3‡	51.0‡	24.0‡	62.0‡
rankAvg:mxmxqe	0.257‡	0.0579‡	92.1‡	81.4‡	78.8‡	68.9	86.0‡	82.0†	80.8	73.1‡	83.0‡	50.7‡	48.4‡	21.5‡	66.7‡
rankAvg:noLex	0.420‡	0.141‡	92.8‡	83.7‡	81.5‡	71.2‡	86.8‡	83.9‡	82.8‡	75.5‡	85.8‡	57.3‡	55.2‡	28.7‡	55.8‡
rankAvg:noNC	0.491‡	0.186‡	89.8‡	80.4‡	79.5‡	69.8‡	86.3‡	83.8‡	82.4‡	75.5‡	86.4*	59.2*	57.2†	31.5‡	52.3†
rankAvg:noNCnoLex	0.431‡	0.144‡	90.7‡	81.7‡	79.9‡	70.1‡	86.5‡	84.0‡	83.0‡	75.6‡	86.1‡	57.9‡	55.8‡	29.7‡	54.6‡
allQE(32)allMBR	0.464‡	0.171‡	91.6‡	82.7‡	80.8‡	70.8‡	86.7‡	83.7‡	82.3‡	75.4‡	86.1‡	58.1‡	56.0‡	29.9‡	54.0‡
allQE(32)nolexMBR	0.421‡	0.146‡	92.6‡	83.8‡	81.1‡	70.9‡	86.8‡	83.8‡	82.8‡	75.4‡	85.6‡	56.9‡	54.8‡	28.4‡	56.2‡
topQE(32)topMBR	0.323‡	0.110‡	94.7‡	85.7‡	81.3‡	70.8‡	86.5‡	82.8†	81.5‡	74.2	84.0‡	53.0‡	50.8‡	23.5‡	62.9‡
noncQE(32)noncMBR	0.446‡	0.121‡	90.4‡	81.2‡	79.4‡	69.8‡	86.4‡	83.5‡	82.1‡	75.1‡	85.8‡	57.3‡	55.2‡	29.2‡	54.8‡
noncQE(32)noncnolexMBR	0.398‡	0.109‡	91.0‡	82.2‡	79.6‡	69.9‡	86.5‡	83.6‡	82.7‡	75.1‡	85.4‡	56.2‡	54.1‡	27.6‡	56.9‡
mxQE(32)mxMBR	0.266‡	0.0795‡	92.1‡	81.4‡	78.9‡	68.9	86.1‡	82.2	80.8	73.3‡	83.4‡	51.5‡	49.3‡	22.2‡	65.7‡
ckQE(32)xcMBR	0.445‡	0.157‡	95.8‡	81.4‡	81.9‡	69.6‡	86.3‡	82.1*	80.6	73.5‡	83.6‡	51.3‡	49.1‡	22.2‡	63.6‡
mxQE(32)xcMBR	0.412‡	0.104‡	95.9‡	80.8‡	79.5‡	68.5	85.9‡	81.7‡	80.5	73.0‡	83.2‡	50.0‡	47.9‡	21.1‡	65.0‡
ckQE(32)mxMBR	0.282‡	0.0967‡	92.8‡	82.7‡	81.7‡	70.0‡	86.5‡	82.5	81.0†	73.9†	83.7‡	52.5‡	50.2‡	22.9‡	64.3‡

Table 15: Reference-based and QE evaluation scores for greedy and MBR/QE decoding (1st block), and ensembles (2nd block), on en-hi (FLORES200 test dataset). Higher scores are better, except MetricX, MetricX-QE, and TER, where lower is better. Green is better than greedy, red is worse. Ensembles are defined in Table 2. Significant differences from greedy (pairwise t-test) indicated by * for p<0.05, † for p<0.01, ‡ for p<0.001. The green diagonal in the 1st block shows metrics prefer outputs from MBR/QE decoding using the same utility metric.

G.6 Results for English-Tamil (en-ta) on FLORES200 test dataset

MBR/QE Method	Evaluated Metric														
	MetricX	MetricX-QE	XCOMET-XXL	XCOMET-XL	CometKiwi23-XXL	CometKiwi23-XL	CometKiwi22	COMET22	IndicCOMET	BLEURT	YiSi	chrF	chrF++	sentBLEU	TER
Greedy	0.789	0.491	80.6	65.9	83.4	76.6	86.7	90.2	71.9	80.9	83.7	60.2	54.4	17.6	73.4
MetricX	0.215 [‡]	0.100 [‡]	87.5 [‡]	67.9 [‡]	85.5 [‡]	77.6 [‡]	86.8	90.2	72.1	82.1 [‡]	80.8 [‡]	53.8 [‡]	47.3 [‡]	11.0 [‡]	87.1 [‡]
MetricX-QE	0.424 [‡]	0.0432 [‡]	84.1 [‡]	65.3	84.4 [‡]	76.7	86.7	89.7 [‡]	71.1 [‡]	81.0	80.7 [‡]	53.0 [‡]	46.7 [‡]	10.9 [‡]	85.6 [‡]
XCOMET-XXL	0.442 [‡]	0.185 [‡]	93.6 [‡]	69.1 [‡]	86.7 [‡]	77.9 [‡]	87.0 [†]	90.3	72.0	82.1 [‡]	81.0 [‡]	53.4 [‡]	47.1 [‡]	11.4 [‡]	84.9 [‡]
XCOMET-XL	0.471 [‡]	0.225 [‡]	87.4 [‡]	76.7 [‡]	86.1 [‡]	79.1 [‡]	87.3 [‡]	90.6 [‡]	72.6 [‡]	82.5 [‡]	81.7 [‡]	55.3 [‡]	49.0 [‡]	12.5 [‡]	82.4 [‡]
CometKiwi23-XXL	0.492 [‡]	0.215 [‡]	88.4 [‡]	69.4 [‡]	90.1 [‡]	79.4 [‡]	87.5 [‡]	90.6 [‡]	72.3 [*]	82.4 [‡]	81.7 [‡]	55.0 [‡]	48.6 [‡]	12.0 [‡]	82.0 [‡]
CometKiwi23-XL	0.549 [‡]	0.256 [‡]	85.4 [‡]	70.1 [‡]	87.2 [‡]	82.3 [‡]	87.5 [‡]	90.6 [‡]	72.1	82.8 [‡]	81.5 [‡]	54.9 [‡]	48.4 [‡]	11.5 [‡]	84.6 [‡]
CometKiwi22	0.578 [‡]	0.263 [‡]	83.9 [‡]	68.0 [‡]	86.2 [‡]	78.9 [‡]	88.5 [‡]	90.8 [‡]	72.4 [†]	81.6 [‡]	81.9 [‡]	55.1 [‡]	48.7 [‡]	12.0 [‡]	80.8 [‡]
COMET22	0.509 [‡]	0.274 [‡]	84.8 [‡]	69.3 [‡]	86.0 [‡]	78.9 [‡]	87.6 [‡]	92.1 [‡]	73.7 [‡]	82.7 [‡]	82.9 [‡]	58.2 [‡]	52.0 [‡]	14.5 [‡]	78.8 [‡]
IndicCOMET	0.628 [‡]	0.356 [‡]	82.3 [‡]	66.4	83.9 [*]	77.1 [*]	87.0 [†]	90.7 [‡]	76.9 [‡]	82.1 [‡]	82.0 [‡]	55.7 [‡]	49.3 [‡]	12.5 [‡]	81.5 [‡]
BLEURT	0.636 [‡]	0.347 [‡]	82.2 [‡]	65.0 [*]	84.6 [‡]	77.9 [‡]	86.7	90.0	71.4 [†]	87.0 [‡]	80.7 [‡]	53.6 [‡]	46.9 [‡]	10.3 [‡]	89.7 [‡]
YiSi	0.698 [*]	0.416 [*]	80.9	66.3	84.0 [†]	77.4 [‡]	87.0 [‡]	90.6 [‡]	72.1	81.3 [‡]	84.2 [‡]	60.4	54.4	16.7 [†]	73.6
chrF	0.679 [†]	0.418 [*]	80.9	66.3	84.1 [‡]	77.7 [‡]	86.9 [*]	90.6 [‡]	72.1	81.8 [‡]	83.9	61.4 [‡]	55.0 [‡]	16.4 [‡]	76.3 [*]
chrF++	0.686 [*]	0.439	80.9	66.2	83.9 [*]	77.6 [‡]	86.9 [*]	90.5 [‡]	72.2 [†]	81.6 [‡]	83.9 [*]	61.3 [‡]	55.2 [‡]	17.1 [*]	75.2
sentBLEU	0.772	0.490	80.4	66.0	83.2	76.7	86.8	90.3	71.9	80.8	83.8	59.9	54.2	17.9	71.5
TER	0.835	0.495	79.8 [*]	65.2 [*]	82.5 [‡]	75.9 [‡]	86.5	89.9 [*]	71.6 [*]	80.2 [‡]	83.6	59.0 [‡]	53.2 [‡]	17.2	69.5 [‡]
rankAvg:all	0.419 [‡]	0.175 [‡]	87.2 [‡]	71.3 [‡]	87.5 [‡]	79.7 [‡]	87.8 [‡]	91.4 [‡]	73.8 [‡]	83.2 [‡]	83.5 [*]	59.6 [†]	53.5 [‡]	16.0 [‡]	75.6
rankAvg:qe	0.388 [‡]	0.105 [‡]	87.8 [‡]	71.3 [‡]	88.8 [‡]	81.0 [‡]	88.1 [‡]	91.1 [‡]	72.8 [‡]	82.8 [‡]	82.2 [‡]	56.1 [‡]	49.7 [‡]	12.5 [‡]	81.3 [‡]
rankAvg:top	0.321 [‡]	0.0985 [‡]	90.8 [‡]	73.1 [‡]	88.7 [‡]	80.8 [‡]	87.7 [‡]	91.0 [‡]	72.9 [‡]	83.1 [‡]	81.8 [‡]	55.6 [‡]	49.1 [‡]	12.2 [‡]	83.3 [‡]
rankAvg:topQe	0.385 [‡]	0.0913 [‡]	88.4 [‡]	70.9 [‡]	89.1 [‡]	81.3 [‡]	87.7 [‡]	90.8 [‡]	72.7 [‡]	82.8 [‡]	81.8 [‡]	55.2 [‡]	48.8 [‡]	12.0 [‡]	83.3 [‡]
rankAvg:mxmxqe	0.229 [‡]	0.0563 [‡]	87.4 [‡]	67.9 [‡]	85.8 [‡]	77.8 [‡]	86.9	90.2	72.1	82.2 [‡]	80.8 [‡]	53.9 [‡]	47.4 [‡]	11.0 [‡]	86.3 [‡]
rankAvg:noLex	0.368 [‡]	0.136 [‡]	88.8 [‡]	72.5 [‡]	88.1 [‡]	80.2 [‡]	87.9 [‡]	91.5 [‡]	74.1 [‡]	83.7 [‡]	83.0 [‡]	58.0 [‡]	51.7 [‡]	13.9 [‡]	79.3 [‡]
rankAvg:noNC	0.441 [‡]	0.181 [‡]	85.2 [‡]	69.5 [‡]	86.1 [‡]	78.7 [‡]	87.5 [‡]	91.4 [‡]	74.0 [‡]	83.2 [‡]	83.7	60.1	54.0	16.6 [‡]	74.0
rankAvg:noNCnoLex	0.366 [‡]	0.136 [‡]	86.2 [‡]	70.1 [‡]	86.5 [‡]	79.0 [‡]	87.5 [‡]	91.5 [‡]	74.5 [‡]	83.9 [‡]	83.2 [‡]	58.5 [‡]	52.2 [‡]	14.5 [‡]	78.4 [‡]
allQE(32)allMBR	0.414 [‡]	0.184 [‡]	87.6 [‡]	71.3 [‡]	87.3 [‡]	79.6 [‡]	87.7 [‡]	91.4 [‡]	73.8 [‡]	83.0 [‡]	83.4 [‡]	59.3 [‡]	53.1 [‡]	15.4 [‡]	76.3 [*]
allQE(32)nolexMBR	0.370 [‡]	0.151 [‡]	89.2 [‡]	72.5 [‡]	87.8 [‡]	79.9 [‡]	87.8 [‡]	91.5 [‡]	74.2 [‡]	83.7 [‡]	82.9 [‡]	57.9 [‡]	51.7 [‡]	14.2 [‡]	79.3 [‡]
topQE(32)topMBR	0.308 [‡]	0.121 [‡]	91.0 [‡]	73.6 [‡]	88.1 [‡]	79.8 [‡]	87.5 [‡]	90.8 [‡]	72.8 [‡]	82.7 [‡]	81.8 [‡]	55.5 [‡]	49.2 [‡]	12.7 [‡]	82.5 [‡]
noncQE(32)noncMBR	0.399 [‡]	0.138 [‡]	85.6 [‡]	69.1 [‡]	86.3 [‡]	78.4 [‡]	87.5 [‡]	91.3 [‡]	73.6 [‡]	82.8 [‡]	83.4 [‡]	59.1 [‡]	53.0 [‡]	15.9 [‡]	75.8 [*]
noncQE(32)noncnolexMBR	0.343 [‡]	0.117 [‡]	86.8 [‡]	69.9 [‡]	86.7 [‡]	78.9 [‡]	87.5 [‡]	91.4 [‡]	74.3 [‡]	83.7 [‡]	82.9 [‡]	57.8 [‡]	51.5 [‡]	14.2 [‡]	78.9 [‡]
mxQE(32)mxMBR	0.230 [‡]	0.0857 [‡]	87.5 [‡]	68.2 [‡]	85.7 [‡]	77.8 [‡]	86.9 [*]	90.3	72.2	82.3 [‡]	81.0 [‡]	54.2 [‡]	47.7 [‡]	11.2 [‡]	85.9 [‡]
ckQE(32)xcMBR	0.432 [‡]	0.176 [‡]	93.2 [‡]	69.5 [‡]	88.3 [‡]	78.7 [‡]	87.3 [‡]	90.4 [*]	72.1	82.3 [‡]	81.2 [‡]	54.1 [‡]	47.7 [‡]	11.6 [‡]	84.1 [‡]
mxQE(32)xcMBR	0.387 [‡]	0.115 [‡]	93.2 [‡]	68.9 [‡]	86.9 [‡]	77.9 [‡]	87.1 [‡]	90.3	72.0	82.3 [‡]	81.0 [‡]	53.4 [‡]	47.1 [‡]	11.4 [‡]	84.4 [‡]
ckQE(32)mxMBR	0.254 [‡]	0.102 [‡]	88.2 [‡]	69.8 [‡]	88.0 [‡]	78.9 [‡]	87.3 [‡]	90.6 [‡]	72.5 [‡]	82.7 [‡]	81.5 [‡]	54.9 [‡]	48.4 [‡]	11.8 [‡]	84.6 [‡]

Table 16: Reference-based and QE evaluation scores for greedy and MBR/QE decoding (1st block), and ensembles (2nd block), on en-ta (FLORES200 test dataset). Higher scores are better, except MetricX, MetricX-QE, and TER, where lower is better. Green is better than greedy, red is worse. Ensembles are defined in Table 2. Significant differences from greedy (pairwise t-test) indicated by * for p<0.05, † for p<0.01, ‡ for p<0.001. The green diagonal in the 1st block shows metrics prefer outputs from MBR/QE decoding using the same utility metric.

G.7 Results for English-Gujarati (en-gu) on FLORES200 test dataset

MBR/QE Method	Evaluated Metric														
	MetricX	MetricX-QE	XCOMET-XXL	XCOMET-XL	CometKiwi23-XXL	CometKiwi23-XL	CometKiwi22	COMET22	IndicCOMET	BLEURT	YiSi	chrF	chrF++	sentBLEU	TER
Greedy	0.794	0.263	94.3	86.2	76.2	69.1	87.7	89.1	95.2	83.2	84.7	54.9	51.7	21.8	65.0
MetricX	0.196 ‡	0.0327 ‡	96.2 ‡	87.9 ‡	76.5	68.6 *	87.7	89.1	95.3	84.0 ‡	81.8 ‡	48.0 ‡	44.5 ‡	14.7 ‡	78.9 ‡
MetricX-QE	0.583 ‡	0.00737 ‡	94.5	84.1 ‡	73.5 ‡	67.0 ‡	87.2 ‡	88.0 ‡	93.6 ‡	81.8 ‡	81.3 ‡	45.6 ‡	42.4 ‡	14.0 ‡	78.3 ‡
XCOMET-XXL	0.547 ‡	0.107 ‡	98.5 ‡	87.4 ‡	77.6 ‡	68.3 ‡	87.6	88.7 ‡	94.3 ‡	82.8 ‡	81.5 ‡	46.5 ‡	43.1 ‡	14.1 ‡	77.9 ‡
XCOMET-XL	0.468 ‡	0.0935 ‡	96.5 ‡	92.9 ‡	78.3 ‡	70.5 ‡	88.2 ‡	89.5 ‡	95.2	84.4 ‡	82.7 ‡	49.3 ‡	45.9 ‡	16.0 ‡	74.7 ‡
CometKiwi23-XXL	0.609 ‡	0.133 ‡	96.2 ‡	87.9 ‡	86.2 ‡	71.7 ‡	88.2 ‡	89.1	94.8 ‡	83.3	82.6 ‡	49.2 ‡	45.8 ‡	15.7 ‡	74.9 ‡
CometKiwi23-XL	0.713 †	0.186 ‡	94.7 *	88.0 ‡	79.8 ‡	75.8 ‡	88.1 ‡	88.8 †	94.3 ‡	82.7 *	82.3 ‡	48.3 ‡	45.0 ‡	15.0 ‡	77.7 ‡
CometKiwi22	0.631 ‡	0.132 ‡	95.2 ‡	87.9 ‡	79.2 ‡	70.8 ‡	89.3 ‡	89.6 ‡	95.4 *	83.6 *	82.6 ‡	48.5 ‡	45.2 ‡	15.4 ‡	74.5 ‡
COMET22	0.578 ‡	0.134 ‡	95.6 ‡	88.7 ‡	78.5 ‡	70.5 ‡	88.5 ‡	91.0 ‡	96.5 ‡	84.2 ‡	83.8 ‡	51.9 ‡	48.6 ‡	17.9 ‡	70.3 ‡
IndicCOMET	0.745	0.233	94.4	85.4 *	75.4 *	68.5 †	87.8	89.3 *	99.4 ‡	83.2	82.8 ‡	49.5 ‡	46.1 ‡	16.1 ‡	73.4 ‡
BLEURT	0.608 ‡	0.168 ‡	94.9 †	87.3 ‡	77.0 *	69.2	87.9 ‡	89.1	95.1	88.0 ‡	82.3 ‡	48.1 ‡	44.6 ‡	14.4 ‡	77.7 ‡
YiSi	0.762	0.252	94.3	86.4	77.6 ‡	69.9 ‡	87.8	89.5 ‡	95.4 *	83.6 †	85.3 ‡	55.1	51.8	21.0 †	65.1
chrF	0.772	0.269	94.2	86.5	77.4 ‡	70.2 ‡	87.8	89.4 ‡	95.4 †	83.6 †	84.9 *	55.8 ‡	52.4 ‡	21.2 *	66.9 ‡
chrF++	0.767	0.266	94.2	86.4	77.4 ‡	70.2 ‡	87.8	89.4 ‡	95.4 *	83.6 †	85.0 *	55.8 ‡	52.5 ‡	21.5	66.7 ‡
sentBLEU	0.813	0.290	94.0	85.7	75.8	68.7 *	87.5 †	89.1	95.0 *	82.9	84.7	54.4 *	51.4	21.8	64.1 *
TER	0.827	0.298	94.0	85.2 ‡	74.6 ‡	67.6 ‡	87.4 ‡	88.9 †	94.7 ‡	82.5 ‡	84.5 *	53.6 ‡	50.5 ‡	21.4	62.4 ‡
rankAvg:all	0.449 ‡	0.0762 ‡	96.7 ‡	90.0 ‡	81.3 ‡	72.0 ‡	88.6 ‡	90.3 ‡	96.7 ‡	84.9 ‡	84.6	54.2 †	50.8 ‡	20.1 ‡	67.1 ‡
rankAvg:qe	0.466 ‡	0.0402 ‡	96.4 ‡	89.7 ‡	83.7 ‡	73.8 ‡	88.9 ‡	89.8 ‡	95.6 ‡	83.9 ‡	83.0 ‡	49.9 ‡	46.5 ‡	16.3 ‡	74.6 ‡
rankAvg:top	0.346 ‡	0.0339 ‡	97.4 ‡	91.1 ‡	83.3 ‡	73.3 ‡	88.4 ‡	89.7 ‡	95.7 ‡	84.4 ‡	83.1 ‡	50.5 ‡	47.1 ‡	16.5 ‡	74.8 ‡
rankAvg:topQe	0.487 ‡	0.0313 ‡	96.4 ‡	89.5 ‡	84.3 ‡	74.3 ‡	88.4 ‡	89.4 †	95.1	83.5	82.8 ‡	49.7 ‡	46.3 ‡	16.1 ‡	75.3 ‡
rankAvg:mxmxqe	0.201 ‡	0.0128 ‡	96.1 ‡	87.9 ‡	76.5	68.5 *	87.7	89.0	95.3	84.0 ‡	81.8 ‡	47.8 ‡	44.3 ‡	14.5 ‡	79.0 ‡
rankAvg:noLex	0.397 ‡	0.0527 ‡	97.0 ‡	90.6 ‡	82.2 ‡	72.6 ‡	88.7 ‡	90.3 ‡	97.0 ‡	85.4 ‡	84.1 ‡	52.6 ‡	49.2 ‡	18.3 ‡	70.6 ‡
rankAvg:noNC	0.467 ‡	0.0805 ‡	95.7 ‡	88.8 ‡	78.9 ‡	70.2 ‡	88.2 ‡	90.2 ‡	96.7 ‡	84.9 ‡	84.6	54.3 †	51.0 †	20.6 ‡	66.2 †
rankAvg:noNCnoLex	0.381 ‡	0.0535 ‡	96.2 ‡	89.3 ‡	79.1 ‡	70.4 ‡	88.3 ‡	90.3 ‡	97.3 ‡	85.7 ‡	84.4 ‡	53.2 ‡	49.8 ‡	18.8 ‡	69.3 ‡
allQE(32)allMBR	0.450 ‡	0.0803 ‡	96.7 ‡	90.1 ‡	81.4 ‡	71.8 ‡	88.6 ‡	90.3 ‡	96.6 ‡	84.9 ‡	84.5 †	54.0 ‡	50.7 ‡	20.0 ‡	67.6 ‡
allQE(32)nolexMBR	0.384 ‡	0.0615 ‡	97.2 ‡	90.9 ‡	81.6 ‡	71.8 ‡	88.6 ‡	90.3 ‡	97.0 ‡	85.5 ‡	84.0 ‡	52.7 ‡	49.3 ‡	18.6 ‡	70.5 ‡
topQE(32)topMBR	0.312 ‡	0.0438 ‡	97.7 ‡	91.5 ‡	81.8 ‡	71.8 ‡	88.4 ‡	89.7 ‡	95.7 ‡	84.6 ‡	82.8 ‡	50.0 ‡	46.5 ‡	16.3 ‡	75.0 ‡
noncQE(32)noncMBR	0.417 ‡	0.0465 ‡	95.9 ‡	88.6 ‡	78.7 ‡	69.9 ‡	88.2 ‡	90.0 ‡	96.3 ‡	84.7 ‡	84.2 ‡	53.2 ‡	49.9 ‡	19.5 ‡	68.2 ‡
noncQE(32)noncnolexMBR	0.361 ‡	0.0395 ‡	96.2 ‡	89.0 ‡	79.1 ‡	70.2 ‡	88.3 ‡	90.1 ‡	96.8 ‡	85.4 ‡	83.7 ‡	51.6 ‡	48.2 ‡	17.7 ‡	71.9 ‡
mxQE(32)mxMBR	0.217 ‡	0.0204 ‡	96.3 ‡	88.1 ‡	76.7	68.6 *	87.7	89.1	95.1	83.9 ‡	82.1 ‡	48.2 ‡	44.8 ‡	15.1 ‡	77.4 ‡
ckQE(32)xcMBR	0.550 ‡	0.0946 ‡	98.2 ‡	88.1 ‡	82.1 ‡	70.0 ‡	88.0 ‡	89.1	94.7 ‡	83.4	82.3 ‡	48.2 ‡	44.8 ‡	15.0 ‡	76.0 ‡
mxQE(32)xcMBR	0.455 ‡	0.0397 ‡	98.3 ‡	87.8 ‡	77.7 ‡	68.3 ‡	87.7	88.8 *	94.5 ‡	83.2	81.7 ‡	46.7 ‡	43.3 ‡	14.0 ‡	77.5 ‡
ckQE(32)mxMBR	0.244 ‡	0.0348 ‡	96.7 ‡	89.0 ‡	82.1 ‡	70.8 ‡	88.2 ‡	89.5 ‡	95.5 *	84.4 ‡	82.5 ‡	49.5 ‡	46.0 ‡	15.7 ‡	76.6 ‡

Table 17: Reference-based and QE evaluation scores for greedy and MBR/QE decoding (1st block), and ensembles (2nd block), on en-gu (FLORES200 test dataset). Higher scores are better, except MetricX, MetricX-QE, and TER, where lower is better. Green is better than greedy, red is worse. Ensembles are defined in Table 2. Significant differences from greedy (pairwise t-test) indicated by * for p<0.05, † for p<0.01, ‡ for p<0.001. The green diagonal in the 1st block shows metrics prefer outputs from MBR/QE decoding using the same utility metric.

G.8 Results for English-Malayalam (en-ml) on FLORES200 test dataset

MBR/QE Method	Evaluated Metric														
	MetricX	MetricX-QE	XCOMET-XXL	XCOMET-XL	CometKiwi23-XXL	CometKiwi23-XL	CometKiwi22	COMET22	IndicCOMET	BLEURT	YiSi	chrF	chrF++	sentBLEU	TER
Greedy	0.849	0.393	91.7	83.6	82.6	77.1	86.9	88.5	94.3	80.9	82.4	56.5	50.9	15.9	74.9
MetricX	0.245‡	0.0652‡	95.8‡	86.3‡	84.6‡	77.8‡	86.9	89.0†	94.5	81.8‡	79.4‡	50.4‡	44.3‡	10.3‡	87.7‡
MetricX-QE	0.572‡	0.0197‡	93.3‡	82.5†	82.5	76.4‡	86.4‡	87.7‡	92.8‡	79.6‡	78.8‡	48.2‡	42.2‡	9.63‡	89.1‡
XCOMET-XXL	0.526‡	0.150‡	98.5‡	86.3‡	85.1‡	77.8‡	87.1*	88.9	93.7†	81.2	79.7‡	50.6‡	44.5‡	10.1‡	86.9‡
XCOMET-XL	0.464‡	0.129‡	96.0‡	91.9‡	85.8‡	79.6‡	87.7‡	89.8‡	95.0‡	83.0‡	80.7‡	52.6‡	46.7‡	12.2‡	82.4‡
CometKiwi23-XXL	0.553‡	0.134‡	96.5‡	87.4‡	89.3‡	79.5‡	87.9‡	89.5‡	94.6	82.0‡	80.6‡	51.8‡	45.8‡	11.1‡	82.9‡
CometKiwi23-XL	0.577‡	0.165‡	95.0‡	88.9‡	86.4‡	81.6‡	87.9‡	89.5‡	94.5	82.1‡	80.5‡	52.1‡	46.1‡	11.4‡	84.2‡
CometKiwi22	0.623‡	0.208‡	94.0‡	86.7‡	85.3‡	78.9‡	88.9‡	89.6‡	94.7†	81.7‡	80.7‡	52.0‡	46.1‡	11.6‡	82.5‡
COMET22	0.519‡	0.181‡	94.7‡	87.4‡	85.5‡	78.9‡	87.9‡	91.2‡	96.2‡	83.0‡	81.7‡	55.1‡	49.0‡	13.0‡	79.4‡
IndicCOMET	0.738†	0.295†	92.3	83.7	82.9	77.1	87.2‡	89.5‡	99.1‡	81.6‡	80.5‡	52.0‡	45.9‡	11.1‡	83.3‡
BLEURT	0.621‡	0.237‡	93.3‡	84.8‡	84.0‡	77.8‡	87.4‡	89.2‡	94.3	87.0‡	79.9‡	50.9‡	44.7‡	9.60‡	89.5‡
YiSi	0.754*	0.332*	92.1	84.1	83.5‡	77.7‡	87.3‡	89.5‡	95.1‡	81.9‡	83.3‡	57.9‡	52.0‡	15.4	74.3
chrF	0.790	0.348	91.2	83.4	83.3‡	77.5†	87.2‡	89.5‡	95.1‡	81.9‡	83.0‡	59.1‡	52.9‡	14.9‡	76.9‡
chrF++	0.789	0.384	91.3	83.6	83.1†	77.5†	87.2‡	89.4‡	94.8‡	81.9‡	82.9‡	58.7‡	52.8‡	15.2*	76.5†
sentBLEU	0.877	0.415	91.1*	83.0	82.1*	76.7*	86.8	89.0‡	94.4	81.1	82.7*	56.8	51.5	16.9‡	72.4‡
TER	0.914	0.454	91.1*	82.7*	80.9‡	75.9‡	86.4‡	88.5	94.1	80.5	82.4	55.5†	50.2*	15.9	70.0‡
rankAvg:all	0.424‡	0.114‡	96.1‡	88.9‡	86.7‡	79.7‡	88.2‡	90.5‡	96.5‡	83.6‡	82.6	57.1	51.2	15.1*	75.2
rankAvg:qe	0.435‡	0.0638‡	96.1‡	89.1‡	88.1‡	80.6‡	88.5‡	90.1‡	95.5‡	83.0‡	81.1‡	53.3‡	47.2‡	12.0‡	82.1‡
rankAvg:top	0.345‡	0.0578‡	97.6‡	90.1‡	87.9‡	80.4‡	88.0‡	90.1‡	95.4‡	83.3‡	81.0‡	53.3‡	47.3‡	12.5‡	82.1‡
rankAvg:topQe	0.422‡	0.0521‡	96.2‡	89.0‡	88.4‡	80.9‡	88.0‡	89.9‡	95.1‡	82.7‡	80.9‡	52.8‡	46.7‡	11.6‡	82.9‡
rankAvg:mxmxqe	0.264‡	0.0280‡	95.9‡	86.2‡	84.8‡	77.9‡	86.9	89.0*	94.5	81.6†	79.4‡	50.3‡	44.2‡	10.3‡	87.6‡
rankAvg:noLex	0.375‡	0.0845‡	96.8‡	89.8‡	87.3‡	80.1‡	88.3‡	90.6‡	96.9‡	84.2‡	82.2	55.8	49.7‡	13.6‡	79.2‡
rankAvg:noNC	0.464‡	0.127‡	94.7‡	87.4‡	85.5‡	78.8‡	87.9‡	90.5‡	96.5‡	83.6‡	82.7*	57.4†	51.5	15.3	74.4
rankAvg:noNCnoLex	0.385‡	0.0801‡	95.6‡	88.2‡	86.1‡	79.2‡	88.0‡	90.6‡	97.1‡	84.5‡	82.2	55.9	49.7‡	13.5‡	79.1‡
allQE(32)allMBR	0.413‡	0.112‡	96.2‡	89.2‡	86.6‡	79.6‡	88.1‡	90.5‡	96.4‡	83.7‡	82.5	56.8	50.9	14.8‡	76.1*
allQE(32)nolexMBR	0.367‡	0.0948‡	96.8‡	90.0‡	86.8‡	79.8‡	88.1‡	90.6‡	96.9‡	84.3‡	82.1*	55.6*	49.5‡	13.7‡	79.1‡
topQE(32)topMBR	0.325‡	0.0696‡	97.7‡	90.5‡	86.9‡	79.7‡	87.8‡	89.9‡	95.3‡	83.1‡	80.8‡	53.0‡	46.9‡	12.2‡	82.7‡
noncQE(32)noncMBR	0.427‡	0.0816‡	95.0‡	87.5‡	85.5‡	78.7‡	87.8‡	90.1‡	96.2‡	83.0‡	82.2	56.2	50.3	14.5‡	76.8†
noncQE(32)noncnolexMBR	0.376‡	0.0669‡	95.7‡	88.0‡	85.8‡	79.0‡	87.9‡	90.4‡	96.7‡	84.0‡	81.8‡	54.8‡	48.8‡	13.1‡	80.1‡
mxQE(32)mxMBR	0.259‡	0.0452‡	95.9‡	86.6‡	84.9‡	78.0‡	87.0	89.0*	94.5	81.8‡	79.7‡	50.8‡	44.7‡	10.5‡	87.1‡
ckQE(32)xcMBR	0.487‡	0.131‡	98.4‡	87.2‡	86.9‡	78.5‡	87.5‡	89.4‡	94.5	81.9‡	80.3‡	51.6‡	45.6‡	11.2‡	83.7‡
mxQE(32)xcMBR	0.453‡	0.0756‡	98.4‡	87.0‡	85.3‡	77.9‡	87.2*	89.1†	94.1	81.3	79.9‡	50.6‡	44.6‡	10.7‡	85.8‡
ckQE(32)mxMBR	0.266‡	0.0570‡	96.4‡	87.6‡	87.0‡	78.9‡	87.5‡	89.5‡	95.1‡	82.5‡	80.1‡	51.6‡	45.4‡	11.0‡	85.1‡

Table 18: Reference-based and QE evaluation scores for greedy and MBR/QE decoding (1st block), and ensembles (2nd block), on en-ml (FLORES200 test dataset). Higher scores are better, except MetricX, MetricX-QE, and TER, where lower is better. Green is better than greedy, red is worse. Ensembles are defined in Table 2. Significant differences from greedy (pairwise t-test) indicated by * for p<0.05, † for p<0.01, ‡ for p<0.001. The green diagonal in the 1st block shows metrics prefer outputs from MBR/QE decoding using the same utility metric.

G.9 Results for English-Vietnamese (en-vi) on FLORES200 test dataset

MBR/QE Method	Evaluated Metric													
	MetricX	MetricX-QE	XCOMET-XXL	XCOMET-XL	CometKiwi23-XXL	CometKiwi23-XL	CometKiwi22	COMET22	BLEURT	YiSi	chrF	chrF++	sentBLEU	TER
Greedy	1.16	0.555	93.7	92.7	91.6	79.6	85.9	90.4	76.8	89.2	62.6	62.6	42.7	41.3
MetricX	0.486‡	0.211‡	96.6‡	94.7‡	93.8‡	81.0‡	86.4‡	90.2†	75.9‡	86.8‡	56.4‡	56.2‡	32.7‡	53.4‡
MetricX-QE	0.673‡	0.117‡	95.7‡	93.8‡	93.3‡	80.5‡	86.2‡	89.8‡	74.7‡	86.4‡	54.5‡	54.3‡	31.6‡	53.5‡
XCOMET-XXL	0.755‡	0.298‡	98.4‡	94.8‡	94.6‡	81.1‡	86.4‡	90.1†	75.4‡	86.9‡	56.2‡	56.0‡	33.1‡	52.5‡
XCOMET-XL	0.725‡	0.303‡	96.9‡	96.6‡	94.1‡	82.1‡	86.6‡	90.5	76.6	87.6‡	57.9‡	57.8‡	35.3‡	49.4‡
CometKiwi23-XXL	0.809‡	0.310‡	97.0‡	94.7‡	96.0‡	81.7‡	86.7‡	90.4	76.0‡	87.5‡	57.3‡	57.1‡	34.7‡	50.2‡
CometKiwi23-XL	0.832‡	0.328‡	96.1‡	95.3‡	94.4‡	83.6‡	86.7‡	90.5	76.1‡	87.4‡	57.3‡	57.2‡	34.5‡	50.5‡
CometKiwi22	0.912‡	0.366‡	95.4‡	94.3‡	93.8‡	81.4‡	87.6‡	90.4	75.9‡	87.3‡	57.1‡	56.9‡	34.2‡	50.7‡
COMET22	0.907‡	0.410‡	95.4‡	94.3‡	93.3‡	81.0‡	86.6‡	91.5‡	77.2‡	88.6‡	60.9‡	60.9‡	39.2‡	44.5‡
BLEURT	0.923‡	0.420‡	95.0‡	94.2‡	93.1‡	80.7‡	86.5‡	90.7‡	78.9‡	88.6‡	60.4‡	60.4‡	39.2‡	44.8‡
YiSi	1.08*	0.519†	94.0	93.1*	92.1†	80.0‡	86.1‡	90.5†	76.8	89.3*	62.5	62.5	42.1*	42.1†
chrF	1.09	0.531	94.1*	92.9	92.1†	80.0‡	86.0‡	90.5	76.7	89.1	63.1†	63.0*	41.4‡	43.4‡
chrF++	1.09*	0.527*	94.1†	92.9	92.1‡	79.9‡	86.0†	90.5	76.7	89.1	63.1†	63.1*	41.5‡	43.4‡
sentBLEU	1.11	0.546	93.9	92.9	91.7	79.8	86.0	90.4	76.8	89.2	62.6	62.6	42.5	41.3
TER	1.21	0.592*	93.6	92.5	91.3*	79.4*	85.8*	90.4	76.5	89.2	62.1*	62.2*	42.5	39.7‡
rankAvg:all	0.759‡	0.302‡	96.4‡	95.3‡	94.4‡	81.7‡	86.8‡	91.1‡	77.8‡	89.0*	62.3	62.2	41.2‡	43.1‡
rankAvg:qe	0.688‡	0.198‡	97.0‡	95.4‡	95.4‡	82.7‡	87.3‡	90.6†	76.5	87.6‡	57.8‡	57.7‡	35.2‡	49.7‡
rankAvg:top	0.608‡	0.203‡	97.8‡	95.9‡	95.3‡	82.5‡	86.8‡	90.6†	76.7	87.5‡	58.1‡	57.9‡	35.3‡	50.4‡
rankAvg:topQe	0.670‡	0.185‡	97.2‡	95.4‡	95.5‡	82.8‡	86.9‡	90.5	76.3*	87.4‡	57.6‡	57.4‡	34.8‡	50.7‡
rankAvg:mxmxqe	0.518‡	0.144‡	96.7‡	94.7‡	94.0‡	81.1‡	86.5‡	90.1†	75.6‡	86.7‡	56.0‡	55.8‡	32.3‡	53.8‡
rankAvg:noLex	0.680‡	0.254‡	97.2‡	95.8‡	94.9‡	82.2‡	87.0‡	91.1‡	77.9‡	88.6‡	60.9‡	60.8‡	39.1‡	45.4‡
rankAvg:noNC	0.795‡	0.324‡	95.7‡	94.5‡	93.4‡	81.0‡	86.5‡	91.0‡	77.9‡	89.2	62.5	62.5	41.6‡	42.3‡
rankAvg:noNCnoLex	0.715‡	0.264‡	96.1‡	94.9‡	93.8‡	81.3‡	86.7‡	91.2‡	78.1‡	89.0†	61.8‡	61.8‡	40.5‡	43.7‡
allQE(32)allMBR	0.735‡	0.287‡	96.7‡	95.4‡	94.6‡	81.8‡	86.9‡	91.0‡	77.8‡	88.8‡	61.5‡	61.5‡	40.4‡	44.1‡
allQE(32)nolexMBR	0.681‡	0.261‡	97.2‡	95.8‡	94.8‡	82.0‡	86.9‡	91.1‡	77.9‡	88.6‡	60.7‡	60.7‡	39.1‡	45.4‡
topQE(32)topMBR	0.589‡	0.220‡	97.8‡	96.0‡	94.9‡	82.2‡	86.8‡	90.5	76.5	87.4‡	57.8‡	57.7‡	35.1‡	50.8‡
noncQE(32)noncMBR	0.717‡	0.242‡	96.2‡	94.8‡	93.8‡	81.2‡	86.6‡	91.0‡	77.7‡	88.8‡	61.3‡	61.2‡	40.1‡	44.1‡
noncQE(32)noncnolexMBR	0.650‡	0.226‡	96.3‡	95.0‡	94.0‡	81.4‡	86.7‡	91.1‡	77.8‡	88.6‡	60.9‡	60.8‡	39.5‡	44.8‡
mxQE(32)mxMBR	0.520‡	0.177‡	96.5‡	94.6‡	93.9‡	81.1‡	86.5‡	90.2†	75.7‡	86.9‡	56.4‡	56.2‡	33.1‡	53.2‡
ckQE(32)xcMBR	0.739‡	0.280‡	98.3‡	94.9‡	95.2‡	81.6‡	86.5‡	90.3	75.7‡	87.1‡	56.7‡	56.6‡	34.0‡	51.4‡
mxQE(32)xcMBR	0.684‡	0.213‡	98.1‡	94.8‡	94.7‡	81.3‡	86.4‡	90.1†	75.6‡	87.0‡	56.1‡	55.9‡	33.1‡	52.1‡
ckQE(32)mxMBR	0.529‡	0.213‡	97.1‡	95.0‡	94.9‡	81.6‡	86.6‡	90.4	76.3†	87.2‡	57.3‡	57.2‡	34.1‡	51.9‡

Table 19: Reference-based and QE evaluation scores for greedy and MBR/QE decoding (1st block), and ensembles (2nd block), on en-vi (FLORES200 test dataset). Higher scores are better, except MetricX, MetricX-QE, and TER, where lower is better. Green is better than greedy, red is worse. Ensembles are defined in Table 2. Significant differences from greedy (pairwise t-test) indicated by * for p<0.05, † for p<0.01, ‡ for p<0.001. The green diagonal in the 1st block shows metrics prefer outputs from MBR/QE decoding using the same utility metric.

G.10 Results for English-Hungarian (en-hu) on FLORES200 test dataset

MBR/QE Method	Evaluated Metric													
	MetricX	MetricX-QE	XCOMET-XXL	XCOMET-XL	CometKiwi23-XXL	CometKiwi23-XL	CometKiwi22	COMET22	BLEURT	YiSi	chrF	chrF++	sentBLEU	TER
Greedy	0.589	0.338	96.2	94.1	93.7	82.7	87.9	90.9	90.5	84.6	60.6	57.4	26.7	58.7
MetricX	0.117 ‡	0.0585 ‡	97.1 ‡	95.8 ‡	95.4 ‡	83.7 ‡	88.0	90.6 †	92.5 ‡	81.6 ‡	54.1 ‡	50.3 ‡	18.5 ‡	71.5 ‡
MetricX-QE	0.344 ‡	0.0166 ‡	96.0	94.2	93.9	82.6	87.6 ‡	89.9 ‡	90.6	81.1 ‡	52.5 ‡	48.6 ‡	17.2 ‡	71.9 ‡
XCOMET-XXL	0.367 ‡	0.162 ‡	99.2 ‡	95.6 ‡	96.0 ‡	83.5 ‡	87.9	90.4 ‡	91.4 ‡	81.8 ‡	54.1 ‡	50.4 ‡	18.7 ‡	69.8 ‡
XCOMET-XL	0.336 ‡	0.140 ‡	97.9 ‡	98.0 ‡	96.0 ‡	85.6 ‡	88.5 ‡	91.2 ‡	92.9 ‡	82.9 ‡	56.3 ‡	52.8 ‡	21.2 ‡	67.0 ‡
CometKiwi23-XXL	0.369 ‡	0.149 ‡	98.1 ‡	96.1 ‡	97.7 ‡	85.2 ‡	88.5 ‡	91.1 *	92.1 ‡	83.0 ‡	56.3 ‡	52.8 ‡	21.3 ‡	66.6 ‡
CometKiwi23-XL	0.375 ‡	0.155 ‡	97.5 ‡	96.9 ‡	96.3 ‡	87.5 ‡	88.6 ‡	91.1	92.2 ‡	82.8 ‡	56.1 ‡	52.5 ‡	20.8 ‡	66.6 ‡
CometKiwi22	0.449 ‡	0.179 ‡	96.7 *	95.7 ‡	95.5 ‡	84.5 ‡	89.5 ‡	91.3 ‡	91.8 ‡	83.1 ‡	56.6 ‡	53.0 ‡	20.9 ‡	66.1 ‡
COMET22	0.414 ‡	0.212 ‡	97.1 ‡	95.7 ‡	95.3 ‡	84.4 ‡	88.7 ‡	92.4 ‡	92.2 ‡	84.4 *	59.6 ‡	56.2 ‡	24.6 ‡	61.7 ‡
BLEURT	0.402 ‡	0.205 ‡	96.5	95.5 ‡	95.0 ‡	83.7 ‡	88.1 †	90.5 ‡	96.4 ‡	81.3 ‡	53.6 ‡	49.8 ‡	17.4 ‡	75.1 ‡
YiSi	0.561	0.330	96.0	94.4	94.0	83.1 †	88.1 †	91.1 ‡	90.9 †	85.3 ‡	61.2 †	58.0 †	27.4 *	57.8 *
chrF	0.506 ‡	0.317	96.0	94.4	94.1 *	83.2 *	88.0 *	91.2 ‡	91.0 ‡	85.1 ‡	62.4 ‡	59.0 ‡	27.3	59.0
chrF++	0.516 ‡	0.317	96.0	94.4	94.2 *	83.1 *	88.1 †	91.2 ‡	91.1 ‡	85.1 ‡	62.1 ‡	58.8 ‡	27.6 †	59.1
sentBLEU	0.589	0.338	95.8 *	94.1	93.6	82.7	87.8	90.8	90.4	84.8 *	60.6	57.6	28.0 ‡	57.2 ‡
TER	0.621	0.351	95.7 †	93.8	93.0 ‡	82.2 ‡	87.6 ‡	90.6 ‡	89.7 ‡	84.6	59.7 ‡	56.7 *	27.4	55.4 ‡
rankAvg:all	0.282 ‡	0.116 ‡	98.1 ‡	96.8 ‡	96.4 ‡	85.5 ‡	88.8 ‡	91.9 ‡	93.0 ‡	84.9 †	61.2 *	57.8	26.8	58.9
rankAvg:qe	0.274 ‡	0.0509 ‡	97.9 ‡	97.0 ‡	97.1 ‡	86.6 ‡	89.2 ‡	91.5 ‡	92.8 ‡	83.3 ‡	57.3 ‡	53.8 ‡	22.2 ‡	65.2 ‡
rankAvg:top	0.196 ‡	0.0476 ‡	98.7 ‡	97.4 ‡	97.1 ‡	86.4 ‡	88.7 ‡	91.4 ‡	93.0 ‡	83.1 ‡	57.1 ‡	53.5 ‡	21.6 ‡	66.4 ‡
rankAvg:topQe	0.280 ‡	0.0397 ‡	98.1 ‡	97.0 ‡	97.3 ‡	86.8 ‡	88.7 ‡	91.3 ‡	92.6 ‡	83.0 ‡	56.6 ‡	53.1 ‡	21.4 ‡	65.9 ‡
rankAvg:mxmxqe	0.127 ‡	0.0239 ‡	97.3 ‡	95.9 ‡	95.5 ‡	83.9 ‡	88.0	90.6 †	92.5 ‡	81.6 ‡	54.1 ‡	50.3 ‡	18.5 ‡	71.2 ‡
rankAvg:noLex	0.231 ‡	0.0761 ‡	98.5 ‡	97.2 ‡	96.8 ‡	86.0 ‡	89.0 ‡	92.0 ‡	93.7 ‡	84.5	60.0 *	56.5 †	25.0 ‡	61.9 ‡
rankAvg:noNC	0.296 ‡	0.117 ‡	97.3 ‡	95.9 ‡	95.6 ‡	84.5 ‡	88.5 ‡	91.8 ‡	92.8 ‡	85.1 ‡	61.5 ‡	58.3 ‡	27.5 *	57.8 *
rankAvg:noNCnoLex	0.226 ‡	0.0690 ‡	97.5 ‡	96.3 ‡	95.9 ‡	84.8 ‡	88.6 ‡	92.0 ‡	93.7 ‡	84.6	60.1 *	56.7 *	25.4 ‡	61.2 ‡
allQE(32)allMBR	0.288 ‡	0.119 ‡	98.1 ‡	96.8 ‡	96.4 ‡	85.6 ‡	88.8 ‡	91.9 ‡	93.0 ‡	84.8	60.8	57.5	26.7	59.2
allQE(32)nolexMBR	0.228 ‡	0.0929 ‡	98.5 ‡	97.2 ‡	96.7 ‡	85.8 ‡	88.8 ‡	91.9 ‡	93.7 ‡	84.3 *	59.8 †	56.4 ‡	24.6 ‡	61.6 ‡
topQE(32)topMBR	0.182 ‡	0.0686 ‡	98.7 ‡	97.4 ‡	96.7 ‡	85.7 ‡	88.6 ‡	91.3 ‡	93.0 ‡	82.9 ‡	56.5 ‡	52.9 ‡	21.0 ‡	66.9 ‡
noncQE(32)noncMBR	0.265 ‡	0.0709 ‡	97.4 ‡	96.1 ‡	95.6 ‡	84.5 ‡	88.4 ‡	91.6 ‡	92.7 ‡	84.4	59.8 †	56.4 ‡	25.3 ‡	60.3 †
noncQE(32)noncnolexMBR	0.206 ‡	0.0591 ‡	97.5 ‡	96.4 ‡	95.8 ‡	84.8 ‡	88.5 ‡	91.8 ‡	93.6 ‡	84.0 ‡	58.7 ‡	55.2 ‡	23.7 ‡	63.2 ‡
mxQE(32)mxMBR	0.126 ‡	0.0373 ‡	97.2 ‡	95.8 ‡	95.4 ‡	83.9 ‡	88.0	90.6 †	92.4 ‡	81.6 ‡	53.9 ‡	50.1 ‡	18.3 ‡	71.2 ‡
ckQE(32)xcMBR	0.346 ‡	0.139 ‡	99.1 ‡	95.9 ‡	96.9 ‡	84.5 ‡	88.3 ‡	90.8	92.0 ‡	82.4 ‡	55.2 ‡	51.5 ‡	19.9 ‡	68.1 ‡
mxQE(32)xcMBR	0.298 ‡	0.0600 ‡	98.9 ‡	95.8 ‡	96.0 ‡	83.8 ‡	88.0	90.5 ‡	91.4 ‡	81.9 ‡	54.1 ‡	50.4 ‡	18.7 ‡	69.2 ‡
ckQE(32)mxMBR	0.132 ‡	0.0554 ‡	98.0 ‡	96.3 ‡	96.8 ‡	84.8 ‡	88.3 ‡	91.0	92.8 ‡	82.3 ‡	55.5 ‡	51.8 ‡	19.8 ‡	68.8 ‡

Table 20: Reference-based and QE evaluation scores for greedy and MBR/QE decoding (1st block), and ensembles (2nd block), on en-hu (FLORES200 test dataset). Higher scores are better, except MetricX, MetricX-QE, and TER, where lower is better. Green is better than greedy, red is worse. Ensembles are defined in Table 2. Significant differences from greedy (pairwise t-test) indicated by * for p<0.05, † for p<0.01, ‡ for p<0.001. The green diagonal in the 1st block shows metrics prefer outputs from MBR/QE decoding using the same utility metric.

G.11 Results for English-German (en-de) on WMT2023 dataset

MBR/QE Method	Evaluated Metric													
	MetricX	MetricX-QE	XCOMET-XXL	XCOMET-XL	CometKiwi23-XXL	CometKiwi23-XL	CometKiwi22	COMET22	BLEURT	YiSi	chrF	chrF++	sentBLEU	TER
Greedy	1.24	1.42	90.1	87.2	79.6	70.7	81.3	85.6	73.5	87.9	70.1	68.2	45.4	42.1
MetricX	0.571 ‡	0.794 ‡	92.0 ‡	87.8 *	79.7	70.4	80.1 ‡	83.9 ‡	73.2	82.2 ‡	58.9 ‡	55.9 ‡	27.4 ‡	63.4 ‡
MetricX-QE	0.630 ‡	0.494 ‡	91.8 ‡	87.5	79.9	70.3	80.3 ‡	83.7 ‡	72.8 *	82.2 ‡	58.2 ‡	55.1 ‡	26.3 ‡	64.8 ‡
XCOMET-XXL	0.915 ‡	1.03 ‡	94.5 ‡	88.3 ‡	81.6 ‡	71.6 ‡	80.6 ‡	84.0 ‡	72.9	83.1 ‡	59.7 ‡	57.0 ‡	29.8 ‡	60.3 ‡
XCOMET-XL	0.907 ‡	1.04 ‡	92.1 ‡	90.8 ‡	81.1 ‡	72.5 ‡	81.2	84.5 ‡	73.3	83.8 ‡	61.1 ‡	58.4 ‡	31.3 ‡	59.0 ‡
CometKiwi23-XXL	1.06 †	1.09 ‡	92.0 ‡	88.4 ‡	85.5 ‡	72.6 ‡	81.5 *	85.1 *	73.2	84.9 ‡	63.4 ‡	60.9 ‡	34.0 ‡	54.7 ‡
CometKiwi23-XL	1.05 ‡	1.15 ‡	91.5 ‡	89.1 ‡	82.6 ‡	75.7 ‡	81.9 ‡	85.1 †	73.5	84.9 ‡	63.8 ‡	61.3 ‡	34.5 ‡	55.3 ‡
CometKiwi22	1.11 ‡	1.20 ‡	91.0 ‡	88.0 ‡	81.5 ‡	72.3 ‡	83.4 ‡	85.5 ‡	73.3	85.5 ‡	64.7 ‡	62.2 ‡	35.8 ‡	52.7 ‡
COMET22	1.01 ‡	1.23 ‡	91.4 ‡	88.1 ‡	80.6 ‡	71.3 ‡	81.6 *	87.0 ‡	74.7 ‡	86.4 ‡	67.3 ‡	65.0 ‡	39.7 ‡	47.3 ‡
BLEURT	0.874 ‡	0.999 ‡	91.5 ‡	88.2 ‡	80.8 ‡	71.4 ‡	81.3	85.4	77.4 ‡	84.8 ‡	63.9 ‡	61.3 ‡	34.0 ‡	54.6 ‡
YiSi	1.27	1.43	90.1	87.2	79.2	70.5	81.1	85.7	73.8	88.1	70.0	67.9	44.3 *	42.4
chrF	1.22	1.42	90.1	87.2	79.7	71.0	81.4	85.8	73.8	87.7	70.4	68.2	43.3 ‡	44.7 ‡
chrF++	1.23	1.42	90.2	87.3	79.7	70.9	81.3	85.7	73.8	87.7	70.3	68.3	44.0 †	43.7 †
sentBLEU	1.29	1.48	90.0	87.2	79.3	70.3 *	81.0 †	85.5	73.6	87.8	69.7	67.8	44.9	42.2
TER	1.40 *	1.55 *	90.0	87.2	78.6 ‡	69.8 ‡	80.7 ‡	85.0 †	73.3	87.2 †	68.5 ‡	66.5 ‡	44.0 †	41.5
rankAvg:all	0.948 ‡	1.09 ‡	92.2 ‡	89.1 ‡	82.0 ‡	72.6 ‡	81.9 ‡	86.3 ‡	75.4 ‡	87.1 †	68.7 ‡	66.5 ‡	42.0 ‡	44.7 ‡
rankAvg:qe	0.868 ‡	0.800 ‡	92.3 ‡	88.9 ‡	83.9 ‡	74.1 ‡	82.6 ‡	85.6	74.4 †	85.2 ‡	64.3 ‡	61.8 ‡	35.3 ‡	53.7 ‡
rankAvg:top	0.762 ‡	0.822 ‡	93.3 ‡	89.7 ‡	83.3 ‡	73.7 ‡	81.6 *	85.2	74.4 †	84.3 ‡	63.1 ‡	60.5 ‡	33.4 ‡	56.0 ‡
rankAvg:topQe	0.798 ‡	0.738 ‡	92.5 ‡	88.9 ‡	84.1 ‡	74.3 ‡	81.8 ‡	85.2 *	74.0	84.6 ‡	62.8 ‡	60.2 ‡	33.3 ‡	56.3 ‡
rankAvg:mxmxqe	0.616 ‡	0.633 ‡	92.1 ‡	87.7 *	79.9	70.6	80.2 ‡	83.8 ‡	73.2	82.2 ‡	58.9 ‡	55.8 ‡	27.1 ‡	63.6 ‡
rankAvg:noLex	0.873 ‡	0.964 ‡	92.7 ‡	89.4 ‡	82.8 ‡	73.2 ‡	82.1 ‡	86.1 †	75.6 ‡	86.2 ‡	66.8 ‡	64.4 ‡	38.5 ‡	48.9 ‡
rankAvg:noNC	0.964 ‡	1.07 ‡	91.4 ‡	88.1 ‡	80.6 ‡	71.5 ‡	81.5 *	86.2 ‡	75.2 ‡	87.1 †	68.8 ‡	66.7 ‡	42.2 ‡	44.4 ‡
rankAvg:noNCnoLex	0.856 ‡	0.931 ‡	91.9 ‡	88.2 ‡	81.0 ‡	71.7 ‡	81.6 †	86.3 ‡	75.7 ‡	86.5 ‡	67.2 ‡	64.9 ‡	39.7 ‡	47.6 ‡
allQE(32)allMBR	0.945 ‡	1.09 ‡	92.2 ‡	89.1 ‡	82.1 ‡	72.7 ‡	82.0 ‡	86.2 †	75.1 ‡	86.8 ‡	68.0 ‡	65.7 ‡	40.9 ‡	46.1 ‡
allQE(32)nolexMBR	0.861 ‡	0.986 ‡	92.8 ‡	89.6 ‡	82.3 ‡	72.8 ‡	81.9 ‡	86.1 †	75.6 ‡	86.2 ‡	66.5 ‡	64.2 ‡	38.7 ‡	48.9 ‡
topQE(32)topMBR	0.739 ‡	0.828 ‡	93.6 ‡	89.9 ‡	82.4 ‡	72.8 ‡	81.4	85.0 *	74.3 *	84.2 ‡	62.7 ‡	60.0 ‡	33.3 ‡	55.6 ‡
noncQE(32)noncMBR	0.825 ‡	0.862 ‡	91.7 ‡	88.3 ‡	81.0 ‡	71.4 ‡	81.6 †	85.9	74.9 ‡	86.3 ‡	66.6 ‡	64.2 ‡	38.6 ‡	48.1 ‡
noncQE(32)noncnolexMBR	0.774 ‡	0.834 ‡	92.1 ‡	88.4 ‡	81.1 ‡	71.6 ‡	81.4	85.8	75.4 ‡	85.6 ‡	65.1 ‡	62.5 ‡	36.3 ‡	51.3 ‡
mxQE(32)mxMBR	0.552 ‡	0.666 ‡	92.0 ‡	87.9 †	80.0	70.4	80.2 ‡	83.9 ‡	73.3	82.5 ‡	59.2 ‡	56.1 ‡	27.5 ‡	63.5 ‡
ckQE(32)xcMBR	0.951 ‡	1.04 ‡	94.2 ‡	88.4 ‡	83.4 ‡	72.2 ‡	81.1	84.7 ‡	73.5	83.9 ‡	61.7 ‡	59.1 ‡	32.1 ‡	56.8 ‡
mxQE(32)xcMBR	0.810 ‡	0.826 ‡	94.3 ‡	88.4 ‡	81.6 ‡	71.4 ‡	80.9 †	84.4 ‡	73.6	83.3 ‡	60.2 ‡	57.4 ‡	30.1 ‡	59.6 ‡
ckQE(32)mxMBR	0.627 ‡	0.776 ‡	92.7 ‡	88.4 ‡	83.0 ‡	71.9 ‡	81.2	84.7 ‡	74.0	83.6 ‡	61.3 ‡	58.5 ‡	30.8 ‡	58.6 ‡

Table 21: Reference-based and QE evaluation scores for greedy and MBR/QE decoding (1st block), and ensembles (2nd block), on en-de (WMT2023 dataset). Higher scores are better, except MetricX, MetricX-QE, and TER, where lower is better. Green is better than greedy, red is worse. Ensembles are defined in Table 2. Significant differences from greedy (pairwise t-test) indicated by * for p<0.05, † for p<0.01, ‡ for p<0.001. The green diagonal in the 1st block shows metrics prefer outputs from MBR/QE decoding using the same utility metric.

G.12 Results for German-English (de-en) on WMT2023 dataset

MBR/QE Method	Evaluated Metric													
	MetricX	MetricX-QE	XCOMET-XXL	XCOMET-XL	CometKiwi23-XXL	CometKiwi23-XL	CometKiwi22	COMET22	BLEURT	YiSi	chrF	chrF++	sentBLEU	TER
Greedy	2.00	1.82	87.3	89.0	76.5	68.4	79.3	85.4	74.7	88.8	68.1	66.5	46.0	39.5
MetricX	1.31‡	1.59	89.0‡	89.2	77.3	68.4	78.7†	83.9‡	72.0‡	83.8‡	58.9‡	56.7‡	31.2‡	57.5‡
MetricX-QE	1.33‡	0.839‡	89.4‡	89.0	78.5‡	69.3*	79.3	84.3‡	72.2‡	85.0‡	59.4‡	57.1‡	31.7‡	56.2‡
XCOMET-XXL	1.71*	1.70	93.6‡	89.8‡	79.4‡	69.6‡	79.4	84.2‡	72.3‡	85.0‡	60.3‡	58.1‡	33.7‡	54.8‡
XCOMET-XL	1.65‡	1.64	90.4‡	92.4‡	79.3‡	71.0‡	80.0‡	85.1	74.0	86.4‡	62.6‡	60.7‡	36.9‡	50.2‡
CometKiwi23-XXL	1.72‡	1.53‡	90.2‡	90.3‡	83.2‡	70.8‡	80.3‡	84.9*	73.3‡	86.4‡	62.6‡	60.6‡	36.8‡	50.1‡
CometKiwi23-XL	1.83	1.66*	89.4‡	90.5‡	80.0‡	73.6‡	80.0‡	84.5‡	72.8‡	86.0‡	62.0‡	60.0‡	35.4‡	51.9‡
CometKiwi22	1.88	1.65*	88.8‡	89.9‡	78.8‡	70.3‡	81.8‡	85.0	73.7†	86.3‡	63.3‡	61.2‡	37.1‡	49.2‡
COMET22	1.84	1.75	89.3‡	89.8‡	77.8‡	69.1‡	79.7†	86.2†	75.1	87.2‡	65.8‡	64.0‡	42.1‡	42.9‡
BLEURT	1.66‡	1.56‡	89.2‡	89.9‡	78.0‡	69.2‡	79.8‡	85.8†	76.6‡	88.0‡	66.2‡	64.3‡	42.2‡	42.5‡
YiSi	1.98	1.77	88.0*	89.4	77.5†	68.9†	79.5	85.7*	75.4†	89.3†	68.3	66.7	46.1	38.8
chrF	1.91	1.80	88.1*	89.3	77.7‡	69.0†	79.5*	85.6	75.0	89.0	69.0	67.2	45.2	41.1
chrF++	1.91	1.80	88.0*	89.2	77.7‡	69.0†	79.5	85.7*	75.2	89.1	69.1*	67.5*	46.3	39.7
sentBLEU	2.00	1.79	87.4	89.2	76.9	68.7	79.3	85.4	75.1	88.9	68.0	66.5	46.5	38.5
TER	2.46†	2.18†	86.1‡	88.2‡	75.6*	67.7†	78.6‡	83.5‡	73.5‡	86.2‡	63.8‡	62.0‡	41.7‡	39.5
rankAvg:all	1.73†	1.53†	90.2‡	90.6‡	79.3‡	70.3‡	80.1‡	85.9*	75.9‡	88.4	67.8	66.1	44.9	40.0
rankAvg:qe	1.54‡	1.17‡	90.9‡	90.8‡	81.5‡	72.1‡	81.0‡	85.6	74.6	87.2‡	64.4‡	62.5‡	38.6‡	47.2‡
rankAvg:top	1.47‡	1.26‡	91.8‡	91.2‡	81.0‡	71.4‡	80.1‡	85.3	74.5	86.5‡	63.8‡	61.9‡	38.4‡	48.4‡
rankAvg:topQe	1.46‡	1.09‡	90.8‡	90.7‡	81.9‡	72.3‡	80.3‡	85.4	74.1	86.9‡	63.4‡	61.4‡	37.1‡	49.2‡
rankAvg:mxmxqe	1.36‡	1.14‡	89.1‡	89.0	78.0‡	69.1*	79.1	84.2‡	72.4‡	84.4‡	60.1‡	58.0‡	33.0‡	54.9‡
rankAvg:noLex	1.60‡	1.41‡	91.2‡	90.9‡	80.2‡	70.9‡	80.4‡	86.0*	76.0‡	88.0*	66.7*	64.9†	42.6‡	42.2‡
rankAvg:noNC	1.78*	1.56*	89.2‡	89.8†	78.1‡	69.3‡	79.7*	85.7	75.7‡	88.1	67.6	66.0	45.2	39.9
rankAvg:noNCnoLex	1.63‡	1.43‡	89.7‡	89.9‡	78.4‡	69.5‡	79.7†	85.8	75.8‡	87.9*	66.6*	64.9†	43.5‡	41.6†
allQE(32)allMBR	1.69†	1.50‡	90.4‡	90.8‡	79.8‡	70.5‡	80.3‡	86.0*	75.9‡	88.5	67.9	66.3	44.7	40.5
allQE(32)nolexMBR	1.56‡	1.42‡	91.1‡	91.0‡	80.0‡	70.7‡	80.4‡	86.1†	76.0‡	88.2	67.1	65.3*	43.4‡	41.6†
topQE(32)topMBR	1.45‡	1.30‡	92.2‡	91.5‡	80.4‡	70.9‡	80.1‡	85.4	74.5	86.6‡	64.0‡	62.1‡	38.7‡	47.9‡
noncQE(32)noncMBR	1.51‡	1.24‡	89.8‡	90.1‡	78.6‡	69.6‡	79.8‡	85.8†	75.4*	88.2*	66.4‡	64.6‡	42.8‡	42.0‡
noncQE(32)noncnolexMBR	1.40‡	1.20‡	90.2‡	90.1‡	78.5‡	69.6‡	79.9‡	86.0‡	75.6†	88.1†	65.8‡	64.0‡	41.9‡	43.2‡
mxQE(32)mxMBR	1.11‡	1.08‡	89.7‡	89.5*	78.3‡	69.2*	79.2	84.7†	72.8‡	85.2‡	60.8‡	58.6‡	33.4‡	55.0‡
ckQE(32)xcMBR	1.67†	1.61	93.2‡	90.5‡	81.2‡	70.3‡	80.0‡	84.7*	73.3‡	85.8‡	62.2‡	60.1‡	36.0‡	50.9‡
mxQE(32)xcMBR	1.43‡	1.17‡	93.3‡	90.3‡	79.7‡	69.9‡	79.7†	84.9	73.3‡	85.9‡	61.8‡	59.6‡	35.1‡	52.3‡
ckQE(32)mxMBR	1.25‡	1.31‡	90.7‡	90.1‡	80.6‡	70.0‡	79.8†	85.0	73.8†	86.0‡	62.5‡	60.5‡	36.0‡	51.8‡

Table 22: Reference-based and QE evaluation scores for greedy and MBR/QE decoding (1st block), and ensembles (2nd block), on de-en (WMT2023 dataset). Higher scores are better, except MetricX, MetricX-QE, and TER, where lower is better. Green is better than greedy, red is worse. Ensembles are defined in Table 2. Significant differences from greedy (pairwise t-test) indicated by * for $p < 0.05$, † for $p < 0.01$, ‡ for $p < 0.001$. The green diagonal in the 1st block shows metrics prefer outputs from MBR/QE decoding using the same utility metric.

G.13 Results for English-Chinese (en-zh) on WMT2023 dataset

MBR/QE Method	Evaluated Metric													
	MetricX	MetricX-QE	XCOMET-XXL	XCOMET-XL	CometKiwi23-XXL	CometKiwi23-XL	CometKiwi22	COMET22	BLEURT	YiSi	chrF	chrF++	sentBLEU	TER
Greedy	1.36	1.24	89.4	85.9	75.5	70.1	80.2	87.0	73.3	88.0	46.9	41.0	11.6	97.1
MetricX	0.682‡	0.711‡	92.8‡	88.1‡	82.4‡	74.1‡	81.8‡	86.3‡	71.1‡	83.9‡	33.2‡	28.8‡	6.39‡	102.‡
MetricX-QE	0.827‡	0.553‡	92.2‡	87.5‡	82.3‡	74.0‡	81.7‡	85.6‡	70.0‡	83.5‡	32.2‡	27.8‡	6.04‡	103.‡
XCOMET-XXL	0.925‡	0.837‡	96.2‡	88.5‡	84.1‡	74.4‡	81.7‡	86.5‡	70.9‡	84.4‡	34.7‡	29.9‡	6.55‡	101.‡
XCOMET-XL	0.927‡	0.864‡	93.5‡	92.3‡	82.3‡	75.7‡	82.0‡	87.1‡	72.4‡	85.2‡	37.1‡	32.1‡	6.63‡	101.†
CometKiwi23-XXL	0.996‡	0.842‡	93.5‡	87.8‡	88.4‡	75.3‡	82.1‡	86.2‡	70.6‡	84.5‡	34.3‡	29.8‡	6.50‡	102.‡
CometKiwi23-XL	1.02‡	0.876‡	92.4‡	89.1‡	83.6‡	78.9‡	82.3‡	86.4‡	70.9‡	84.6‡	34.4‡	29.8‡	6.31‡	101.‡
CometKiwi22	0.995‡	0.869‡	92.2‡	88.2‡	82.5‡	75.0‡	84.2‡	87.1‡	71.8‡	85.2‡	35.9‡	31.1‡	7.23‡	103.‡
COMET22	1.04‡	0.999‡	91.6‡	87.9‡	80.2‡	73.3‡	81.9‡	89.3‡	74.0‡	87.4‡	43.4‡	37.9‡	10.4‡	97.6
BLEURT	1.08‡	1.05‡	91.1‡	87.6‡	79.3‡	72.5‡	81.4‡	87.5‡	76.7‡	87.2‡	42.6‡	37.1‡	8.42‡	98.0
YiSi	1.29‡	1.21*	89.9‡	86.3†	77.0‡	71.0‡	80.7‡	87.7‡	74.2‡	89.0‡	48.4‡	42.3‡	11.5	96.6
chrF	1.28‡	1.20†	90.0‡	86.3†	77.1‡	71.0‡	80.7‡	87.8‡	74.2‡	88.8‡	49.6‡	43.4‡	12.4	97.6
chrF++	1.28‡	1.19†	89.9‡	86.3†	77.1‡	71.0‡	80.7‡	87.8‡	74.2‡	88.8‡	49.4‡	43.6‡	12.7†	97.5
sentBLEU	1.40	1.26	88.7‡	84.8‡	75.4	69.7*	79.9†	86.5‡	71.9‡	86.9‡	43.3‡	38.2‡	15.1‡	106.‡
TER	1.43†	1.26	88.5‡	84.3‡	75.6	69.6†	79.7‡	86.0‡	71.7‡	86.5‡	41.8‡	36.5‡	8.65‡	94.3*
rankAvg:all	0.931‡	0.869‡	93.4‡	89.7‡	83.0‡	74.9‡	82.4‡	88.6‡	75.3‡	88.2*	46.7	40.9	11.6	94.5‡
rankAvg:qe	0.853‡	0.694‡	93.9‡	89.8‡	86.2‡	77.2‡	83.4‡	87.2	72.2‡	85.1‡	36.1‡	31.4‡	7.69‡	100.†
rankAvg:top	0.792‡	0.698‡	94.9‡	90.9‡	85.8‡	76.9‡	82.7‡	87.3*	72.6‡	85.2‡	37.0‡	32.2‡	7.70‡	99.4*
rankAvg:topQe	0.854‡	0.673‡	93.8‡	89.6‡	86.7‡	77.5‡	82.6‡	86.7	71.6‡	84.7‡	35.2‡	30.6‡	7.23‡	100.†
rankAvg:mxmxqe	0.712‡	0.608‡	93.1‡	88.3‡	82.9‡	74.4‡	81.9‡	86.3‡	71.1‡	83.9‡	33.4‡	29.0‡	6.49‡	101.‡
rankAvg:noLex	0.860‡	0.802‡	94.0‡	90.3‡	84.2‡	75.9‡	82.9‡	88.5‡	75.2‡	87.5‡	43.7‡	38.1‡	10.1†	96.2
rankAvg:noNC	0.997‡	0.920‡	91.9‡	88.3‡	80.3‡	73.1‡	81.8‡	88.5‡	75.4‡	88.5‡	47.8†	41.9†	11.9	94.0†
rankAvg:noNCnoLex	0.911‡	0.835‡	92.4‡	88.8‡	81.4‡	73.8‡	82.1‡	88.6‡	75.4‡	88.0	45.0‡	39.3‡	10.5†	96.5
allQE(32)allMBR	0.901‡	0.841‡	93.6‡	89.8‡	83.4‡	75.2‡	82.6‡	88.5‡	75.0‡	87.7†	44.9‡	39.3‡	10.8	95.2
allQE(32)nolexMBR	0.868‡	0.813‡	94.0‡	90.2‡	83.9‡	75.6‡	82.7‡	88.5‡	74.9‡	87.4‡	43.4‡	38.0‡	10.2†	96.4
topQE(32)topMBR	0.778‡	0.730‡	95.1‡	91.0‡	84.8‡	76.2‡	82.4‡	87.3*	72.6†	85.3‡	37.2‡	32.4‡	8.18‡	99.0
noncQE(32)noncMBR	0.915‡	0.781‡	92.6‡	88.7‡	81.7‡	73.9‡	82.1‡	88.1‡	74.5‡	87.5‡	44.3‡	38.7‡	9.85‡	97.2
noncQE(32)noncnolexMBR	0.862‡	0.755‡	92.9‡	89.0‡	82.3‡	74.3‡	82.3‡	88.2‡	74.6‡	87.2‡	42.8‡	37.3‡	9.37‡	97.8
mxQE(32)mxMBR	0.695‡	0.664‡	93.1‡	88.3‡	82.7‡	74.2‡	81.8‡	86.3‡	71.2‡	84.0‡	33.6‡	29.1‡	6.46‡	102.‡
ckQE(32)xcMBR	0.912‡	0.823‡	96.0‡	88.7‡	85.8‡	75.0‡	82.0‡	86.6‡	71.2‡	84.5‡	35.1‡	30.4‡	7.20‡	100.*
mxQE(32)xcMBR	0.860‡	0.724‡	95.9‡	88.8‡	84.4‡	74.8‡	82.1‡	86.5‡	71.2‡	84.4‡	34.8‡	30.1‡	6.93‡	101.†
ckQE(32)mxMBR	0.723‡	0.720‡	93.5‡	88.6‡	85.2‡	74.9‡	82.1‡	86.7*	71.7‡	84.5‡	34.8‡	30.2‡	7.37‡	100.†

Table 23: Reference-based and QE evaluation scores for greedy and MBR/QE decoding (1st block), and ensembles (2nd block), on en-zh (WMT2023 dataset). Higher scores are better, except MetricX, MetricX-QE, and TER, where lower is better. Green is better than greedy, red is worse. Ensembles are defined in Table 2. Significant differences from greedy (pairwise t-test) indicated by * for p<0.05, † for p<0.01, ‡ for p<0.001. The green diagonal in the 1st block shows metrics prefer outputs from MBR/QE decoding using the same utility metric.

G.14 Results for Chinese-English (zh-en) on WMT2023 dataset

MBR/QE Method	Evaluated Metric													
	MetricX	MetricX-QE	XCOMET-XXL	XCOMET-XL	CometKiwi23-XXL	CometKiwi23-XL	CometKiwi22	COMET22	BLEURT	YiSi	chrF	chrF++	sentBLEU	TER
Greedy	2.19	2.00	90.9	88.1	78.4	70.3	79.6	82.7	71.4	82.5	55.0	52.4	26.2	65.2
MetricX	1.01‡	1.20‡	93.9‡	89.6‡	78.7	70.3	79.3*	82.1‡	70.8†	80.4‡	49.5‡	46.6‡	19.1‡	78.9‡
MetricX-QE	1.25‡	0.853‡	93.4‡	89.3‡	79.3‡	71.0‡	80.0‡	81.9‡	70.7‡	80.5‡	49.0‡	46.2‡	19.0‡	78.2‡
XCOMET-XXL	1.41‡	1.38‡	96.2‡	89.8‡	80.2‡	70.9‡	79.7	81.8‡	70.6‡	80.1‡	48.7‡	45.8‡	18.2‡	78.6‡
XCOMET-XL	1.45‡	1.41‡	93.9‡	92.4‡	80.3‡	72.5‡	80.7‡	83.1‡	72.5‡	81.9‡	52.4‡	49.7‡	22.3‡	72.8‡
CometKiwi23-XXL	1.69‡	1.52‡	93.6‡	89.9‡	85.1‡	72.7‡	80.8‡	82.8	71.6	81.6‡	51.8‡	49.0‡	21.4‡	72.7‡
CometKiwi23-XL	1.74‡	1.59‡	92.9‡	90.3‡	81.3‡	75.9‡	81.0‡	82.7	71.5	81.5‡	51.8‡	49.0‡	21.5‡	74.1‡
CometKiwi22	1.74‡	1.51‡	92.8‡	89.7‡	80.7‡	72.4‡	82.8‡	83.1†	72.2‡	81.9†	52.1‡	49.4‡	22.2‡	72.6‡
COMET22	1.73‡	1.63‡	92.5‡	89.6‡	79.8‡	71.6‡	80.6‡	84.6‡	72.6‡	82.8	54.8	52.0	25.1‡	65.7
BLEURT	1.68‡	1.60‡	92.8‡	89.9‡	79.9‡	71.5‡	80.5‡	83.4‡	74.4‡	82.5	54.1‡	51.3‡	24.5‡	67.8‡
YiSi	2.12*	1.95*	91.2†	88.4*	78.9‡	70.5	79.8†	83.0‡	71.8†	83.7‡	55.9†	53.3†	26.6	64.8
chrF	2.16	2.01	91.1	88.1	78.8†	70.7‡	79.8†	82.9*	71.8†	83.2‡	56.9‡	54.1‡	26.5	67.4‡
chrF++	2.14	1.99	91.1*	88.2	78.9†	70.8‡	79.9‡	83.0‡	72.0‡	83.3‡	56.9‡	54.3‡	26.9*	66.6†
sentBLEU	2.17	2.01	91.0	88.0	78.5	70.2	79.6	82.9*	71.5	83.3‡	56.0‡	53.5‡	27.7‡	62.8‡
TER	2.28†	2.11‡	90.5*	87.3‡	77.2‡	69.2‡	78.8‡	82.1‡	70.5‡	82.4	53.1‡	50.5‡	25.3*	60.4‡
rankAvg:all	1.57‡	1.44‡	93.6‡	90.5‡	81.3‡	72.6‡	81.0‡	84.0‡	73.4‡	83.5‡	56.4‡	53.7‡	27.0†	63.7‡
rankAvg:qe	1.44‡	1.16‡	94.0‡	90.9‡	83.3‡	74.4‡	81.9‡	83.6‡	73.0‡	82.2	53.2‡	50.5‡	23.2‡	71.1‡
rankAvg:top	1.26‡	1.17‡	95.0‡	91.6‡	82.7‡	73.9‡	81.1‡	83.5‡	72.9‡	82.1*	53.1‡	50.4‡	22.7‡	71.8‡
rankAvg:topQe	1.39‡	1.11‡	94.0‡	90.8‡	83.4‡	74.5‡	81.1‡	83.3‡	72.6‡	82.0†	52.7‡	49.9‡	22.4‡	71.9‡
rankAvg:mxmxqe	1.07‡	0.972‡	93.9‡	89.9‡	79.6‡	71.1‡	80.0‡	82.3*	71.4	80.8‡	50.5‡	47.6‡	20.0‡	76.9‡
rankAvg:noLex	1.43‡	1.30‡	94.2‡	91.1‡	82.1‡	73.3‡	81.4‡	84.1‡	73.7‡	83.2‡	55.4	52.7	25.7	66.8†
rankAvg:noNC	1.64‡	1.47‡	92.8‡	89.7‡	80.0‡	71.4‡	80.4‡	83.7‡	73.0‡	83.4‡	56.2‡	53.5‡	26.9*	63.2‡
rankAvg:noNCnoLex	1.45‡	1.30‡	93.4‡	90.2‡	80.4‡	71.8‡	80.7‡	84.0‡	73.4‡	83.2‡	55.4	52.7	25.8	65.6
allQE(32)allMBR	1.52‡	1.41‡	93.8‡	90.8‡	81.6‡	72.8‡	81.2‡	84.1‡	73.6‡	83.5‡	56.3‡	53.7‡	27.1†	64.4
allQE(32)nolexMBR	1.40‡	1.33‡	94.3‡	91.2‡	81.7‡	73.0‡	81.3‡	84.2‡	73.8‡	83.3‡	55.7†	53.1*	26.3	65.8
topQE(32)topMBR	1.22‡	1.21‡	95.1‡	91.8‡	81.8‡	73.1‡	80.9‡	83.4‡	72.9‡	82.1*	53.1‡	50.4‡	23.1‡	71.5‡
noncQE(32)noncMBR	1.47‡	1.23‡	93.3‡	90.1‡	80.4‡	71.8‡	80.7‡	83.7‡	73.1‡	82.9*	54.9	52.2	25.3†	66.0
noncQE(32)noncnolexMBR	1.36‡	1.18‡	93.6‡	90.4‡	80.6‡	71.9‡	80.8‡	83.9‡	73.4‡	82.7	54.2†	51.5‡	24.4‡	67.6‡
mxQE(32)mxMBR	1.04‡	1.08‡	93.9‡	89.7‡	79.3‡	70.8†	79.7	82.3†	71.1*	80.5‡	49.8‡	46.9‡	19.6‡	77.8‡
ckQE(32)xcMBR	1.42‡	1.36‡	95.8‡	90.3‡	82.8‡	72.0‡	80.5‡	82.5	71.6	81.0‡	50.6‡	47.7‡	20.4‡	75.4‡
mxQE(32)xcMBR	1.32‡	1.16‡	95.8‡	90.0‡	80.5‡	71.4‡	80.0‡	82.2†	71.3	80.5‡	49.7‡	46.7‡	19.4‡	77.3‡
ckQE(32)mxMBR	1.12‡	1.21‡	94.4‡	90.5‡	82.4‡	72.1‡	80.5‡	83.0†	72.2‡	81.5‡	51.8‡	49.0‡	21.5‡	73.7‡

Table 24: Reference-based and QE evaluation scores for greedy and MBR/QE decoding (1st block), and ensembles (2nd block), on zh-en (WMT2023 dataset). Higher scores are better, except MetricX, MetricX-QE, and TER, where lower is better. Green is better than greedy, red is worse. Ensembles are defined in Table 2. Significant differences from greedy (pairwise t-test) indicated by * for p<0.05, † for p<0.01, ‡ for p<0.001. The green diagonal in the 1st block shows metrics prefer outputs from MBR/QE decoding using the same utility metric.