

What do Large Language Models Need for Machine Translation Evaluation?

Shenbin Qian¹, Archchana Sindhujan^{2*}, Minnie Kabra^{3*}

Diptesh Kanojia², Constantin Orăsan¹, Tharindu Ranasinghe⁴ and Frédéric Blain⁵

¹Centre for Translation Studies, University of Surrey, United Kingdom,

²Institute for People-Centred AI, University of Surrey, United Kingdom,

³Independent Researcher, India,

⁴Lancaster University, United Kingdom, ⁵Tilburg University, The Netherlands

{s.qian, a.sindhujan, d.kanojia, c.orasan}@surrey.ac.uk, minniekabra@gmail.com,

t.ranasinghe@lancaster.ac.uk, f.l.g.blain@tilburguniversity.edu

Abstract

Leveraging large language models (LLMs) for various natural language processing tasks has led to superlative claims about their performance. For the evaluation of machine translation (MT), existing research shows that LLMs are able to achieve results comparable to fine-tuned multilingual pre-trained language models. In this paper, we explore what translation information, such as the source, reference, translation errors and annotation guidelines, is needed for LLMs to evaluate MT quality. In addition, we investigate prompting techniques such as zero-shot, Chain of Thought (CoT) and few-shot prompting for eight language pairs covering high-, medium- and low-resource languages, leveraging varying LLM variants. Our findings indicate the importance of reference translations for an LLM-based evaluation. While larger models do not necessarily fare better, they tend to benefit more from CoT prompting, than smaller models. We also observe that LLMs do not always provide a numerical score when generating evaluations, which poses a question on their reliability for the task. Our work presents a comprehensive analysis for resource-constrained and training-less LLM-based evaluation of machine translation. We release the accrued prompt templates, code and data publicly for reproducibility¹.

1 Introduction

Recent surge in the use of large language models (LLMs) for natural language processing (NLP) tasks like question answering (Kocoń et al., 2023; Tan et al., 2023) has taken strides, and significantly improved their applications to other downstream tasks such as machine translation (MT), text summarization, information retrieval and *etc.*, due to advancements in natural language understanding capabilities, contextual awareness, and a versatile

knowledge base (Kocmi and Federmann, 2023b; Zhu et al., 2023; Zhang et al., 2024).

For automatic evaluation of MT quality, traditional approaches use metrics such as BLEU (Papineni et al., 2002), BLEURT (Sellam et al., 2020) or BERTScore (Zhang* et al., 2020) to compare MT output with a reference translation. When references are not available, quality estimation (QE) methods such as fine-tuning multilingual pre-trained language models (PTLMs) on human evaluation data like Direct Assessment (DA) scores (Graham et al., 2013) are often used to predict estimated scores to approximate human evaluation (Specia et al., 2018). Recent studies leverage prompting techniques and *instruct* LLMs to output a score for translation quality, claim to achieve promising results (Kocmi and Federmann, 2023b,a).

However, there exists no systematic exploration of what translation information LLMs need for quality evaluation, and whether different prompting techniques, such as Chain-of-Thought (CoT) (Wei et al., 2024) or few-shot prompting, can help boost the performance of LLMs. To that end, we conduct this investigation to systematically explore the ability of LLMs in quality evaluation in a training-less scenario. Our contributions can be summarized as:

- We investigate what translation information, *i.e.*, source, reference, translation errors and annotation guidelines LLMs need to evaluate translation for 8 language pairs covering high-, medium- and low-resource languages.
- We explore different ways of prompting, *i.e.*, zero-shot, CoT and few-shot prompting for LLMs to evaluate MT quality. Our code, data and prompts are released with the paper.
- We compare our prompting methods with fine-tuning of encoder-based multilingual PTLMs and find LLM performance still lags behind.

*Both authors contributed equally to this work.

¹https://github.com/surrey-nlp/LLM4MT_eval

- Our analyses of the results on various prompt templates indicate that references are important for accurate translation evaluation with LLMs, and while larger models are not always better, they tend to benefit more from CoT prompting than smaller model variants.

The rest of the paper is structured as follows: Section 2 discusses relevant work in quality evaluation, while Section 3 introduces the dataset we utilize for this work. Section 4 describes the prompting methods and the baselines with the experimental setup. Results and discussion are presented in Section 5. Section 6 concludes our study and outlines future directions.

2 Related Work

Traditional automatic MT quality evaluation metrics such as BLEU, BLEURT and BERTScore compare the MT output to one or several references, whilst metrics like Translation Error Rate (TER) (Snover et al., 2006) are based on the number of edits required for MT output to become reference, and neither takes semantic variations into account.

Training supervised machine learning systems on human-annotated data based on metrics such as DA or Multi-dimensional Quality Metrics (MQM) (Lommel et al., 2014) can help predict translation quality without any references (Deoghare et al., 2023b). Ranasinghe et al. (2020b, 2021) proposed the TransQuest framework to utilize the source text and MT output only and fine-tune XLM-RoBERTa-large (Conneau et al., 2020) to predict a DA score as an estimation of translation quality. COMET (Rei et al., 2020; Stewart et al., 2020; Rei et al., 2022b) was proposed initially to incorporate references along with the source and MT output to train multilingual PTLMs for quality evaluation, but later it also supported reference-less evaluation. Wan et al. (2022) proposed a unified translation evaluation framework that could include source or reference or both as input for quality evaluation. Various approaches achieved promising results in the QE shared task of the Conference on Machine Translation (WMT) (Specia et al., 2020, 2021; Zerva et al., 2022; Blain et al., 2023), however, most require supervision and training (Deoghare et al., 2023a; Kanojia et al., 2021).

The advent of LLMs prompted its application to translation quality evaluation. Kocmi and Federmann (2023b) proposed a zero-shot prompting

technique, called GEMBA for DA score prediction using GPT-4 (OpenAI et al., 2024), claiming LLMs can achieve performance comparable to state-of-the-art models fine-tuned on multilingual data. Based on the GEMBA prompt, Fernandes et al. (2023) proposed to use LLMs for both DA score prediction and error categorization via fine-tuning to achieve more fine-grained evaluation. Previous research focused on whether LLMs can be better translation evaluators than state-of-the-art models. To the best of our knowledge, only Huang et al. (2024) investigated how LLMs leverage the source and reference for quality evaluation. However, they only perform zero-shot prompting for three language pairs. Our work comprehensively examines factors such as translation errors and annotation guidelines across eight language pairs, eight prompt templates, and three different prompting techniques.

3 Data

We utilized the DA score prediction data released with WMT22 QE shared task (Zerva et al., 2022). This dataset includes the source (mainly from news articles), MT output (from different MT engines) and (post-edited) human references for eight language pairs, *i.e.*, English-German (EN-DE), English-Marathi (EN-MR), English-Chinese (EN-ZH), Estonian-English (ET-EN), Nepali-English (NE-EN), Romanian-English (RO-EN), Russian-English (RU-EN) and Sinhala-English (SI-EN). For each source-MT pair, the dataset contains a DA score ranging from 0 to 100, rated by human annotators for quality assessment.

To include annotated errors in the MT output into our prompts, we obtained word-level QE data from WMT22, where tokens of the MT output have sequence labels with either “OK” or “BAD” indicating translation quality at word level. This dataset also involves the above 8 language pairs, which contains the source, MT output and the tags for translation quality. For each MT output, we extracted the tokens that were tagged as “BAD” as error words.

Since source-MT segments from sentence-level QE data might differ from those of word-level, we compared each source-MT pair in the two datasets and used the overlapping as the main resource of our research. It includes source, MT output, reference translations and error words for the 8 language pairs, covering high-, medium- and low-resource

languages. We present different prompt templates in Section 4.2 to selectively include source, reference and error words to test what translation information LLMs need for quality evaluation.

We split the data into training, validation, and test sets in proportions of 80%, 10%, and 10% respectively. We inferred with LLMs on the test set to obtain evaluation results. Training and validation sets were used to sample examples for few-shot learning (see Section 4.4). The size of the test set for each language pair can be seen in Table 1.

LP	EN-DE	EN-MR	EN-ZH	ET-EN	NE-EN	RO-EN	RU-EN	SI-EN
N	891	2598	890	897	761	867	900	343

Table 1: The number of instances (N) in the **test set** for each language pair (LP)

4 Methodology

This section presents the baselines, and our zero-shot, CoT and few-shot prompting methods.

4.1 Baselines

We utilize TransQuest and COMET, two widely used reference-less and reference-based² QE frameworks as baselines. For TransQuest, we employed the fine-tuned MonoTransQuest models proposed by Ranasinghe et al. (2020a) on each language pair except EN-MR. For EN-MR, we used the English to any model released with TransQuest. For COMET, we utilized a fine-tuned multilingual model “Unbabel/wmt22-comet-da” (Rei et al., 2022a) for all language pairs as it does not have models for each language pair.

4.2 Zero-shot Prompting

For LLMs to predict translation quality, our prompt includes 1) instructions to perform the task such as “Score the following translation”, and 2) translation information such as source or reference.

Since Kocmi and Federmann (2023b) have shown that their prompt template can achieve state-of-the-art performance using GPT-4, we mainly followed their template to create our prompt instruction as shown in Figure 1. We used it as our base template and changed slightly different translation information to test what is needed for LLMs to evaluate MT quality. We constructed prompt **Template 1** containing “source + MT output” as

²COMET also supports reference-less evaluation.

translation information, **Template 2** “MT output + reference”, **Template 3** “source + MT output + reference” (exact GEMBA prompt), **Template 4** “source + MT output + error words”, **Template 5** “source + MT output + reference + error words”.

We augmented the base prompt with summarized guidelines used during human evaluation, as **Template 6** to test if this could help LLMs evaluate MT quality. These guidelines instruct evaluators to give a DA score by considering multiple factors including accuracy, contextual understanding, grammar, syntax and overall readability.

4.3 CoT Prompting

Apart from the translation information and guidelines added in the prompt, we also tested whether CoT prompting could improve LLMs’ performance by utilizing reasoning-based steps for quality evaluation. We devised **Template 7** which includes two-step prompts to score MT quality, as shown in Figure 2. In the first prompt, we give translation information (including source, MT output and reference) to the LLM and ask it to analyze step by step where the machine translation is different from the reference. In the second prompt, we instruct the LLM to score the machine translation based on its previous output, *i.e.*, the analysis of machine translation based on reference. Instruction to output a score in JSON format is given to ensure it produces the score first, like other templates.

4.4 Few-shot Learning

In addition to zero-shot and CoT prompting, we also added 5 examples based on Template 3, to show how human annotators score machine translations from 0 – 100. We split the training and validation sets into 5 buckets for each language pair according to the score ranges of 0 – 20, 21 – 40, 41 – 60, 61 – 80, 81 – 100. We randomly sampled 1 example from each range. The selected 5 examples for each language pair were given before the instruction for scoring as a prefix (see Figure 3) of the base prompt in Figure 1. We call this prompt **Template 8**.

4.5 Model Selection

We chose 6 models from a variety of open-source LLMs according to their size, popularity and type such as mixture of expert (MoE) (Shazeer et al., 2017) and dense models, and based on our compute capability. For 7-billion-parameter models, we selected Llama-2-7B from Meta (Touvron et al.,

Score the following translation from {source_lang} to {target_lang} with respect to the {source/human_reference/error_words} on a continuous scale from 0 to 100, where score of zero means "no meaning preserved" and score of one hundred means "perfect meaning and grammar".

{translation_information}

Score:

Figure 1: Base Prompt Template

Prompt 1:

You are going to evaluate the quality for {language_pair} translation. You need to think step by step. First read the following source, machine translation and reference translation. Analyze where the machine translation is different from the reference translation.

Source: {source_sentence}

Machine translation: {target_sentence}

Reference translation: {reference_translation}

Prompt 2:

A large language model did an evaluation of machine translation quality for the {source_language} sentence, which is given as below: {output_from_Prompt1}

Based on your analysis, score the machine translation quality on a continuous scale from 0 to 100, where score of zero means "no meaning preserved" and score of one hundred means "perfect meaning and grammar". Provide the score strictly in JSON format.

Figure 2: Prompt Template 7

2023), Gemma-7B from Google (Gemma Team et al., 2024) and OpenChat3.5 which was trained with mixed-quality data using Conditional Reinforcement Learning from Human Feedback (Wang et al., 2024). For 13-billion-parameter models, we opted for Llama-2-13B and Qwen1.5-14B which was specifically tested on a diverse set of 12 languages and showed impressive multilingual capabilities (Bai et al., 2023). We also included the Mixtral-8x7B model (Jiang et al., 2024) as our MoE model, but due to the limit of our compute capability, we used the activation-aware weight quantized (AWQ) version (Lin et al., 2024). For all 6 selected models, we used the instruction-tuned version, *i.e.*, the chat model, for zero-shot, CoT and few-shot inference. Additionally, we experimented with TowerLLM (Alves et al., 2024) for EN-ZH via HuggingFace³, but results are not discussed in the paper because the model output is mostly identical to the input for most instances (see Appendix A).

4.6 Experimental Setup

All our experiments were run using $1 \times$ NVIDIA A100 40G, $1 \times$ A40, and $2 \times$ RTX A5000 GPUs,

³<https://huggingface.co/Unbabel/TowerInstruct-7B-v0.1>

for different LLM variants. We used vLLM (Kwon et al., 2023) to save inference time. Detailed settings for hyperparameters, formatting and evaluation metrics are provided below.

Hyperparameters We chose the default hyperparameter settings in vLLM for all our experiments, *i.e.*, 0.8 as temperature⁴, 0.95 for top_p. The input sequence length was chosen as 1024 for zero-shot and CoT inference and 3000 for few-shot inference.

Formatting As chat models were fine-tuned on certain formats to interact with humans, it is suggested to use the specific format that was used to train the model while inferencing. As vLLM does not support formatting natively, we formatted all our prompt templates before they were fed into the models based on the format of each model in Section 4.5.

Evaluation Since LLMs usually output a score and some explanations about their evaluation, we used regular expression to extract the score from the LLM output. Spearman correlation⁵ (Spearman,

⁴We also experimented with 0 temperature, but we did not observe huge differences in terms of score distribution.

⁵Pearson's r and Kendall's τ are in Appendix A.

You are going to evaluate the quality of machine translation given the source, machine translation and reference translation. The followings are examples of scoring translation quality.
{5_examples}

Figure 3: Prefix for the base prompt to create Template 8

LP	TransQuest	COMET
EN-DE	0.3811	0.3579
EN-MR	0.2489	0.5135
EN-ZH	0.6360	0.5410
ET-EN	0.8148	0.7018
NE-EN	0.8034	0.6393
RO-EN	0.8739	0.7699
RU-EN	0.8252	0.6482
SI-EN	0.7233	0.5874

Table 2: Spearman ρ achieved by models using TransQuest and COMET on each language pair (LP).

1904) was used to evaluate how the predicted scores are correlated with the (mean of) human annotated scores. However, not all LLMs would output a score for all the instances. Sometimes, LLMs failed to score the input translation. In such cases, we dropped these instances (denoted as D in Table 3) during the process of correlation calculation, but they were noted as a metric for robustness.

5 Results and Discussion

This section displays our baseline and prompting results using existing QE frameworks and LLMs.

5.1 Baselines

Table 2 shows our baseline results using TransQuest and COMET. Since TransQuest models were fine-tuned on data from each language pair with the exception of EN-MR, the Spearman correlation scores of these reference-less models, are higher than those of reference-based COMET models. Except EN-DE and EN-MR, the correlation scores for most language pairs are relatively high.

5.2 Zero-shot Inference

Table 3 shows our zero-shot results of Templates 1 to 6 for open-source LLMs including OpenChat3.5, Llama-2-7B, Gemma7B, Llama-2-13B, Qwen1.5-14B and Mixtral-8x7B-AWQ.

Comparing results among LLMs We observe that OpenChat3.5 achieved the highest Spearman

correlation scores for most language pairs, despite having only 7 billion parameters – roughly half the size of Llama-2-13B and Qwen1.5-14B. It excelled not only in Spearman scores but also in consistently providing valid quality evaluation scores, with very few dropped rows. Among the 6 models, Llama-2 (both the 7 and 13 billion variants) performed poorly in generating evaluations with valid scores. Many rows were dropped, and the Spearman scores were low, indicating a weak correlation with the true scores. The MoE model, Mixtral-8x7B-AWQ, did not outperform OpenChat3.5 on most language pairs and prompt templates for our task.

Comparing with the baselines We find that models fine-tuned for each language pair by TransQuest, performed much better than the zero-shot prompting results for all language pairs. For some language pairs like ET-EN, NE-EN and RO-EN, our best zero-shot prompting results (Template 6 of OpenChat3.5 in Table 3) were comparable to the reference-based models fine-tuned on multilingual data using COMET. For some other pairs like SI-EN, our best zero-shot prompting results were even slightly better than the COMET models.

Comparison among Templates When we fix the model variable as OpenChat3.5, we can compare the performance of different prompt templates. Looking at the OpenChat3.5 results in Table 3, we observe that LLM performance is generally better when the source and reference are included in the prompt, as in Templates 3, 5, and 6, compared to prompts without them, such as Templates 1 and 2. This pattern holds true for other LLMs such as the Llama-2 models and Gemma, as shown in the table. Notably, the Spearman scores are obviously higher when the source is incorporated into the prompt, as seen by comparing Templates 2 and 3. This suggests that the source is an essential component for evaluating MT quality using LLMs, contrary to the results in Huang et al. (2024), who indicate that the source provides a negative impact.

Our results (on Templates 4, 5, 6) suggest that including error words and annotation guidelines does not consistently help LLMs evaluate MT qual-

LP	T1		T2		T3		T4		T5		T6	
	ρ	D	ρ	D	ρ	D	ρ	D	ρ	D	ρ	D
OpenChat3.5												
EN-DE	0.2258	1	0.2209	2	0.2849	0	0.2599	0	0.2812	0	0.2960	0
EN-MR	0.2295	3	0.3110	9	<u>0.3546</u>	0	0.3347	0	0.3565	0	0.3446	0
EN-ZH	0.2722	0	0.2603	4	<u>0.3995</u>	0	0.3002	0	0.3333	0	0.3635	0
ET-EN	0.5402	0	0.5798	2	0.6980	0	0.5879	0	0.6700	0	0.6925	0
NE-EN	0.3784	9	0.4855	25	<u>0.5937</u>	0	0.5008	0	0.5832	0	0.6073	0
RO-EN	0.4712	2	0.5669	25	0.7294	0	0.6900	0	0.7096	0	0.7385	0
RU-EN	0.5714	0	0.5320	13	0.6066	0	0.5494	0	0.5322	0	0.5938	0
SI-EN	0.4120	4	0.4201	7	0.6034	0	0.4364	0	0.5990	0	0.5963	0
Llama-2-7B												
EN-DE	0.0663	5	0.0397	18	0.0876	73	-0.0166	3	0.0887	2	<u>0.0957</u>	27
EN-MR	0.0417	26	0.0024	85	<u>0.1255</u>	377	0.0154	2	0.0861	1	<u>0.0943</u>	140
EN-ZH	<u>0.0956</u>	4	0.0553	15	0.0946	86	0.0273	2	0.0607	2	0.0791	47
ET-EN	0.0439	3	0.1643	3	<u>0.3715</u>	54	-0.0431	1	0.2527	0	0.3319	31
NE-EN	0.1825	47	0.1018	7	<u>0.2207</u>	85	0.0461	1	0.2026	2	<u>0.2629</u>	26
RO-EN	0.3068	0	0.1322	4	<u>0.4514</u>	50	-0.0059	1	0.2444	0	0.3619	20
RU-EN	0.1718	13	0.1389	44	<u>0.4253</u>	64	-0.0081	15	0.2170	12	0.2404	24
SI-EN	0.0801	7	-0.0238	11	0.2212	36	0.0639	2	0.2288	2	<u>0.2530</u>	18
Gemma-7B												
EN-DE	0.1516	1	0.1241	0	0.1624	0	0.1074	0	<u>0.1856</u>	0	0.1820	0
EN-MR	<u>0.3070</u>	0	0.2332	0	0.1479	0	0.1529	0	0.2177	0	0.1815	0
EN-ZH	0.2046	0	0.1362	0	0.1805	0	<u>0.2734</u>	0	0.2444	0	0.2342	0
ET-EN	0.3490	0	0.4074	0	0.3772	0	0.4125	0	<u>0.5552</u>	0	0.5169	0
NE-EN	0.3329	0	0.2732	2	0.2921	0	0.3439	0	<u>0.4098</u>	0	0.4098	0
RO-EN	<u>0.6238</u>	0	0.4393	0	0.4429	0	0.5858	0	0.5816	0	0.5911	0
RU-EN	0.3265	2	0.3697	13	0.4399	0	0.3709	0	0.4450	0	<u>0.5012</u>	0
SI-EN	0.2740	0	0.2610	0	0.3519	0	0.2816	0	<u>0.3980</u>	0	<u>0.3741</u>	0
Llama-2-13B												
EN-DE	-0.0062	535	-0.0092	83	0.0316	118	0.0716	10	<u>0.1161</u>	8	0.1061	123
EN-MR	0.0229	201	-0.0692	282	0.0685	224	0.0193	2	<u>0.1051</u>	2	0.1044	483
EN-ZH	0.0002	104	0.0032	78	<u>0.1412</u>	118	0.0821	5	<u>0.0967</u>	2	0.0974	206
ET-EN	0.2159	268	0.0973	84	<u>0.4042</u>	54	0.2196	3	0.3755	5	<u>0.4392</u>	123
NE-EN	0.0890	78	0.2337	42	<u>0.3178</u>	76	0.1175	4	0.1259	3	0.2895	138
RO-EN	0.2787	417	0.2484	67	0.4616	50	0.2661	0	0.3224	0	<u>0.5102</u>	133
RU-EN	0.3931	216	0.1298	99	0.4074	64	0.3328	25	0.3076	26	<u>0.4422</u>	105
SI-EN	0.0152	79	0.3020	28	0.2669	32	0.0498	2	0.0928	3	<u>0.3659</u>	58
Qwen1.5-14B												
EN-DE	0.1363	16	0.2286	27	0.2182	12	0.1579	0	0.2245	0	<u>0.2359</u>	20
EN-MR	0.3011	16	0.3647	48	0.3131	12	0.2151	0	0.2838	0	0.3033	17
EN-ZH	0.3758	68	<u>0.2500</u>	30	0.4131	11	0.3166	0	0.3504	1	0.4367	18
ET-EN	0.4836	86	0.5240	132	0.6467	26	0.4741	0	0.5483	0	<u>0.6516</u>	34
NE-EN	0.3485	213	0.4777	268	<u>0.5114</u>	33	0.3349	0	0.4466	2	<u>0.4651</u>	29
RO-EN	0.2201	124	0.5161	124	<u>0.7200</u>	17	0.5569	0	0.5790	0	0.6992	18
RU-EN	0.5157	27	0.5196	96	<u>0.5597</u>	12	0.4743	0	0.5397	1	0.5547	18
SI-EN	0.3828	71	0.4691	94	<u>0.5936</u>	7	0.2769	0	0.4091	0	0.5427	16
Mixtral-8x7B-AWQ												
EN-DE	0.0870	4	0.0607	4	<u>0.2631</u>	1	0.1572	2	0.1930	0	0.2309	0
EN-MR	0.1067	19	0.0799	22	<u>0.1825</u>	2	0.0872	8	<u>0.2078</u>	8	0.1936	1
EN-ZH	0.3390	1	0.1253	1	<u>0.3720</u>	0	0.2104	0	0.2746	5	0.3434	0
ET-EN	0.3128	9	0.2081	3	<u>0.6229</u>	1	0.3499	3	0.4338	3	0.5903	0
NE-EN	0.4019	7	0.4025	8	<u>0.4891</u>	0	0.1212	3	0.2279	1	0.4684	0
RO-EN	0.3204	13	0.2557	16	<u>0.6526</u>	0	0.4053	2	0.4404	2	0.6164	1
RU-EN	0.4750	4	0.3742	13	<u>0.5831</u>	0	0.4160	1	0.5022	3	0.5528	0
SI-EN	0.2652	6	0.2139	9	<u>0.4563</u>	0	0.0966	2	0.2220	1	0.4124	1

Table 3: Spearman ρ correlation scores achieved by zero-shot inference using **Templates 1-6 (T1-6)** on various open-source LLMs for each language pair (LP). D -> rows dropped as LLM generated output without a score. The underlined scores represent the best result among templates, while the bold represent the best among LLMs.




Figure 4: Density plots of the predicted (red) and true (mean) DA scores (blue) for high-resource language pairs *i.e.*, EN-DE (left) and EN-ZH (right).

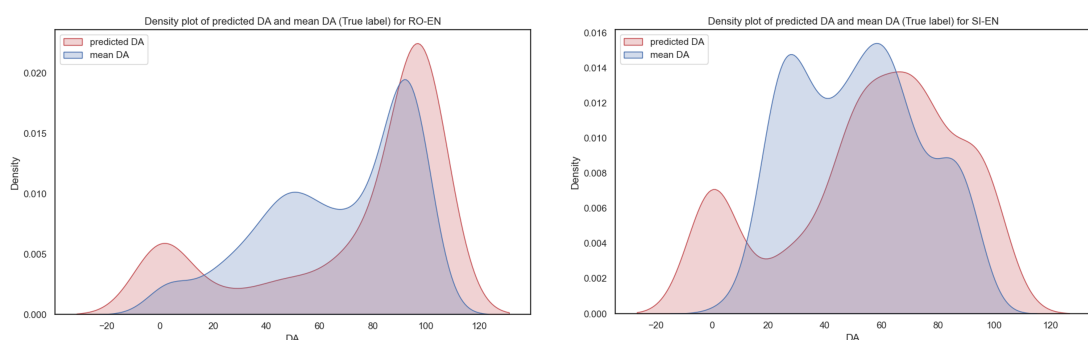


Figure 5: Density plots of the predicted (red) and true (mean) DA scores (blue) for medium- and low-resource language pairs *i.e.*, RO-EN (left) and SI-EN (right).

ity across different language pairs when compared to using just the plain GEMBA prompt (Template 3). For most language pairs like EN-ZH, ET-EN, RU-EN, and SI-EN, Template 3 had the highest correlation with human judgments. However, removing reference translations (Template 4) clearly lowered correlation scores, highlighting their importance for accurate MT evaluation.

Although incorporating error words does not seem to improve performance, they are surprisingly useful in helping LLMs provide scores in their outputs. As shown in Table 3, there are fewer dropped rows when using Templates 4 and 5, which include error words. Outputs from Templates 4 and 5 are the most stable across models, unlike other templates that are more model-dependent.

Results among different language pairs For high-resource language pairs like EN-DE and EN-ZH, correlation scores tend to be lower than those of medium- and low-resource pairs such as NE-EN, RO-EN, and RU-EN. This pattern holds true across most models, including the fine-tuned ones from TransQuest and COMET.

To further investigate the reasons, we selected EN-DE and EN-ZH as high-resource language pairs, and RO-EN and SI-EN as medium- and low-resource language pairs. We plotted the distributions of the predicted (from OpenChat3.5) *vs* true scores (mean of all annotators) as shown in Figures 4 and 5. For high-resource language pairs, the predictions are skewed towards higher DA scores. Well-trained MT systems, due to abundant resources, tend to produce high-quality translations, leading to higher DA scores. However, LLM-based evaluation systems may amplify these imbalanced distributions and are more likely to predict scores within the high range.

In contrast, for medium- and low-resource language pairs, there are fewer resources for training MT systems. As a result, low-quality translations (with low DA scores) are better represented than in high-resource pairs. Quality evaluation systems can better recognize low-quality translations and produce a more balanced score distribution. This imbalance in the score representation could be the reason why predicted DA scores for high-resource languages are less correlated with true scores than

LP	OpenChat3.5		Llama-2-7B		Gemma-7B		Llama-2-13B		Qwen1.5-14B		Mixtral-8x7B-AWQ	
	T7	T3	T7	T3	T7	T3	T7	T3	T7	T3	T7	T3
EN-DE	0.2433	0.2849	-0.0353	<u>0.0876</u>	0.0048	<u>0.1624</u>	<u>0.0345</u>	0.0316	<u>0.2388</u>	0.2182	0.2213	<u>0.2631</u>
EN-MR	0.2937	0.3546	-0.0021	<u>0.1255</u>	0.0859	<u>0.1479</u>	<u>0.0804</u>	0.0685	<u>0.3455</u>	0.3131	<u>0.1906</u>	0.1825
EN-ZH	0.3324	0.3995	0.0354	<u>0.0946</u>	0.1609	<u>0.1805</u>	<u>0.0703</u>	<u>0.1412</u>	0.3429	<u>0.4131</u>	0.2479	<u>0.3720</u>
ET-EN	0.6110	0.6980	0.1459	<u>0.3715</u>	0.3191	<u>0.3772</u>	0.2558	<u>0.4042</u>	0.5845	<u>0.6467</u>	0.4628	<u>0.6229</u>
NE-EN	0.5160	0.5937	0.1363	<u>0.2207</u>	<u>0.3221</u>	0.2921	<u>0.3315</u>	0.3178	0.4791	<u>0.5114</u>	0.4373	<u>0.4891</u>
RO-EN	0.7175	0.7294	0.1859	<u>0.4514</u>	<u>0.4550</u>	0.4429	0.3403	<u>0.4616</u>	0.7019	<u>0.7200</u>	0.6360	<u>0.6526</u>
RU-EN	0.5317	0.6066	0.1618	<u>0.4253</u>	0.2979	<u>0.4399</u>	0.2519	<u>0.4074</u>	0.5203	<u>0.5597</u>	0.5191	<u>0.5831</u>
SI-EN	0.5124	0.6034	0.1818	<u>0.2212</u>	0.2808	<u>0.3519</u>	<u>0.2854</u>	0.2669	0.4680	<u>0.5936</u>	<u>0.4691</u>	0.4563

Table 4: Spearman ρ correlation scores achieved using **Template 7 (T7)**, the CoT prompt template, on various LLMs for each language pair (LP). Results of **Template 3 (T3)** from Table 3 are listed here for reference. The underlined scores represent better result between the templates, while the bold represent the best among LLMs.

LP	OpenChat3.5		Llama-2-7B		Gemma-7B		Llama-2-13B		Qwen1.5-14B		Mixtral-8x7B-AWQ	
	T8	T3	T8	T3	T8	T3	T8	T3	T8	T3	T8	T3
EN-DE	0.1756	0.2849	0.0327	<u>0.0876</u>	0.0343	<u>0.1624</u>	<u>0.0655</u>	0.0316	0.1804	<u>0.2182</u>	0.2155	<u>0.2631</u>
EN-MR	0.2543	0.3546	0.0078	<u>0.1255</u>	<u>0.1651</u>	0.1479	0.0159	<u>0.068</u>	0.2706	<u>0.3131</u>	<u>0.2410</u>	0.1825
EN-ZH	0.2801	<u>0.3995</u>	0.0283	<u>0.0946</u>	<u>0.1831</u>	0.1805	0.0875	<u>0.1412</u>	0.2946	0.4131	0.2970	<u>0.3720</u>
ET-EN	0.5779	0.6980	-0.0026	<u>0.3715</u>	<u>0.4134</u>	0.3772	0.2328	<u>0.4042</u>	0.4320	<u>0.6467</u>	0.5566	<u>0.6229</u>
NE-EN	0.4621	0.5937	0.1428	<u>0.2207</u>	<u>0.3117</u>	0.2921	0.1907	<u>0.3178</u>	0.3349	<u>0.5114</u>	<u>0.5143</u>	0.4891
RO-EN	0.6881	0.7294	0.0405	<u>0.4514</u>	<u>0.4693</u>	0.4429	0.2574	<u>0.4616</u>	0.4498	<u>0.7200</u>	<u>0.6712</u>	0.6526
RU-EN	0.5774	0.6066	0.1680	<u>0.4253</u>	0.2531	<u>0.4399</u>	0.1951	<u>0.4074</u>	0.4798	<u>0.5597</u>	0.5239	<u>0.5831</u>
SI-EN	0.4277	0.6034	0.0352	<u>0.2212</u>	0.3048	<u>0.3519</u>	0.1368	<u>0.2669</u>	0.4207	<u>0.5936</u>	<u>0.4642</u>	0.4563

Table 5: Spearman ρ correlation scores achieved using **Template 8 (T8)**, the few-shot prompt template, on various LLMs for each language pair (LP). Results of **Template 3 (T3)** from Table 3 are listed here for reference. The underlined scores represent better result between the templates, while the bold represent the best among LLMs.

for medium- and low-resource pairs.

5.3 CoT and Few-shot Inference

Tables 4 and 5 show results of CoT (Template 7) and 5-shot inference (Template 8) together with the results of Template 3 for the 6 selected LLMs. Dropped rows for the two templates are presented in Table 6. Both Templates 7 and 8 were built upon Template 3, *i.e.*, including the source, MT output and reference. We expect the model performance to be improved when more reasoning steps or evaluation examples were given. However, for 7 billion parameter variants, CoT prompting resulted in worse performance, as Spearman correlation scores of Template 7 were obviously lower than those of Template 3. For the larger 13 billion parameter variants, results were mixed for different language pairs. For language pairs such as EN-DE and EN-MR, CoT prompting improved the performance in the prediction of DA scores. This indicates that CoT may work better on larger models than smaller models. While CoT prompting did not consistently improve model performance as measured by the Spearman correlation scores, it shows relatively more consistent output than other prompt templates. Table 6 suggests that fewer rows were dropped when using Template 7, especially

for Llama-2 models.

Interestingly, 5-shot inference results are not better than zero-shot results, posing a question on context utilization by LLMs. Performance varies on the LLMs and the specific language pairs. This could relate to the language data available for training these LLMs, as well as the quality of the evaluation examples chosen for different languages pairs.

5.4 Discussion

Based on our results, *Template 3*, which includes the source, MT output and reference, but excludes error words and detailed guidelines, performed the best in terms of Spearman correlation scores. Prompting with CoT and few-shot learning may yield better results for larger models, but more experiments are needed to confirm this.

While larger language models often perform better, our results show that a 7-billion parameter model outperformed other models for most language pairs. Surprisingly, even much smaller COMET models fine-tuned on multilingual data, rather than data for specific language pairs, usually outperformed our LLM prompting results. However, due to the high computational cost, we could not test models with 70 billion or more parameters.

Different models excel at various language pairs

LP	OpenChat3.5		Llama-2-7B		Gemma-7B		Llama-2-13B		Qwen1.5-14B		Mixtral-8x7B-AWQ	
	T7	T8	T7	T8	T7	T8	T7	T8	T7	T8	T7	T8
EN-DE	0	0	0	1	0	0	0	0	3	9	0	0
EN-MR	0	0	0	7	0	0	0	0	1	8	0	2
EN-ZH	0	0	0	21	0	1	0	0	1	24	0	0
ET-EN	0	0	0	0	0	0	0	0	38	22	0	0
NE-EN	0	0	1	1	0	0	1	0	36	23	0	2
RO-EN	0	0	0	0	0	1	0	0	27	69	1	0
RU-EN	0	0	4	7	4	0	3	1	17	29	1	0
SI-EN	0	0	1	2	0	0	0	2	4	5	1	0

Table 6: **Dropped rows for [Template 7 \(T7\)](#) and [Template 8 \(T8\)](#), i.e., the CoT and few-shot prompt templates, using various LLMs for each language pair (LP).**

while struggling with others. Even for a single model, performance fluctuates across different language pairs. This variability could stem from whether a language is considered high-resource, but further research is necessary to understand the underlying causes.

Our experiments with prompting LLMs for translation evaluation reveal that these models are often inconsistent in *generating numerical scores*. In most cases, LLMs tend to generate scores accompanied by lengthy and unstructured explanations. While using regular expressions for extraction can be helpful, it is not always reliable. For models like Llama-2, we observed numerous instances where LLMs failed to produce a valid score. Our empirical findings demonstrate that employing CoT prompting or incorporating error words into the prompt can enhance the consistency of the model outputs.

6 Conclusion and Future Work

In this paper, we explored what translation information is needed for LLMs to evaluate MT quality. We conducted a comprehensive investigation into different prompting techniques such as zero-shot, CoT and few-shot prompting using different translation information for 8 language pairs and 6 LLMs of different sizes and types. Our findings suggest that the source, MT output and reference are essential compared to other information such as translation errors for quality evaluation. Larger models may not necessarily perform better than smaller models, but CoT prompting works better on larger than smaller model variants. We also observe that LLMs do not always provide a numerical score when generating evaluations, which makes their assessments less reliable. For future research, we plan to explore whether fine-tuning LLMs could improve their performance in quality

evaluation. We also plan to thoroughly investigate error explainability of LLMs using MQM and other fine-grained error identification techniques. These future studies can inform downstream error correction through automatic post-editing, contributing to a more comprehensive evaluation and correction framework.

Limitations and Ethical Considerations

Our results were achieved on a limited number of LLMs which are mostly smaller than 14 billion parameters due to the constraints of our computational capabilities. Larger models may perform differently in this translation evaluation task. The examples used in the few-shot scenario were randomly sampled since we do not have the knowledge to prepare good-quality examples for all language pairs. Results might be different if these examples were carefully chosen by native speakers.

Our experiments in the paper were conducted solely on publicly available datasets as described in Section 3, requiring no ethical approval.

Acknowledgements

We would like to thank the European Association for Machine Translation (EAMT) for funding QE data curation of Indic languages used in this paper (UoS/RN0580).

References

- Duarte Miguel Alves, José Pombal, Nuno M Guerreiro, Pedro Henrique Martins, João Alves, Amin Farajian, Ben Peters, Ricardo Rei, Patrick Fernandes, Sweta Agrawal, Pierre Colombo, José G. C. de Souza, and Andre Martins. 2024. [Tower: An Open Multilingual Large Language Model for Translation-Related Tasks](#). In *First Conference on Language Modeling*.
- Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei

- Huang, Binyuan Hui, Luo Ji, Mei Li, Junyang Lin, Runji Lin, Dayiheng Liu, Gao Liu, Chengqiang Lu, Keming Lu, Jianxin Ma, Rui Men, Xingzhang Ren, Xuancheng Ren, Chuanqi Tan, Sinan Tan, Jianhong Tu, Peng Wang, Shijie Wang, Wei Wang, Shengguang Wu, Benfeng Xu, Jin Xu, An Yang, Hao Yang, Jian Yang, Shusheng Yang, Yang Yao, Bowen Yu, Hongyi Yuan, Zheng Yuan, Jianwei Zhang, Xingxuan Zhang, Yichang Zhang, Zhenru Zhang, Chang Zhou, Jingren Zhou, Xiaohuan Zhou, and Tianhang Zhu. 2023. [Qwen Technical Report](#).
- Frederic Blain, Chrysoula Zerva, Ricardo Rei, Nuno M. Guerreiro, Diptesh Kanojia, José G. C. de Souza, Beatriz Silva, Tânia Vaz, Yan Jingxuan, Fatemeh Azadi, Constantin Orasan, and André Martins. 2023. [Findings of the WMT 2023 shared task on quality estimation](#). In *Proceedings of the Eighth Conference on Machine Translation*, pages 629–653, Singapore. Association for Computational Linguistics.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Sourabh Deoghare, Paramveer Choudhary, Diptesh Kanojia, Tharindu Ranasinghe, Pushpak Bhattacharyya, and Constantin Orăsan. 2023a. [A multi-task learning framework for quality estimation](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 9191–9205, Toronto, Canada. Association for Computational Linguistics.
- Sourabh Deoghare, Diptesh Kanojia, Fred Blain, Tharindu Ranasinghe, and Pushpak Bhattacharyya. 2023b. [Quality estimation-assisted automatic post-editing](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 1686–1698, Singapore. Association for Computational Linguistics.
- Patrick Fernandes, Daniel Deutsch, Mara Finkelstein, Parker Riley, André Martins, Graham Neubig, Ankush Garg, Jonathan Clark, Markus Freitag, and Orhan Firat. 2023. [The devil is in the errors: Leveraging large language models for fine-grained machine translation evaluation](#). In *Proceedings of the Eighth Conference on Machine Translation*, pages 1066–1083, Singapore. Association for Computational Linguistics.
- Gemma Team, Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak, Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, Pouya Tafti, Léonard Hussenot, Pier Giuseppe Sessa, Aakanksha Chowdhery, Adam Roberts, Aditya Barua, Alex Botev, Alex Castro-Ros, Ambrose Slone, Amélie Héliou, Andrea Tacchetti, Anna Bulanova, Antonia Paterson, Beth Tsai, Bobak Shahriari, Charline Le Lan, Christopher A. Choquette-Choo, Clément Crepy, Daniel Cer, Daphne Ippolito, David Reid, Elena Buchatskaya, Eric Ni, Eric Noland, Geng Yan, George Tucker, George-Christian Muraru, Grigory Rozhdestvenskiy, Henryk Michalewski, Ian Tenney, Ivan Grishchenko, Jacob Austin, James Keeling, Jane Labanowski, Jean-Baptiste Lespiau, Jeff Stanway, Jenny Brennan, Jeremy Chen, Johan Ferret, Justin Chiu, Justin Mao-Jones, Katherine Lee, Kathy Yu, Katie Millican, Lars Lowe Sjoesund, Lisa Lee, Lucas Dixon, Machel Reid, Maciej Mikuła, Mateo Wirth, Michael Sharman, Nikolai Chinaev, Nithum Thain, Olivier Bachem, Oscar Chang, Oscar Wahltinez, Paige Bailey, Paul Michel, Petko Yotov, Rahma Chaabouni, Ramona Comanescu, Reena Jana, Rohan Anil, Ross McIlroy, Ruibo Liu, Ryan Mullins, Samuel L Smith, Sebastian Borgeaud, Sertan Girgin, Sholto Douglas, Shree Pandya, Siamak Shakeri, Soham De, Ted Klimentko, Tom Hennigan, Vlad Feinberg, Wojciech Stokowiec, Yu hui Chen, Zafarali Ahmed, Zhitao Gong, Tris Warkentin, Ludovic Peran, Minh Giang, Clément Farabet, Oriol Vinyals, Jeff Dean, Koray Kavukcuoglu, Demis Hassabis, Zoubin Ghahramani, Douglas Eck, Joelle Barral, Fernando Pereira, Eli Collins, Armand Joulin, Noah Fiedel, Evan Senter, Alek Andreev, and Kathleen Kenealy. 2024. [Gemma: Open Models Based on Gemini Research and Technology](#). *Preprint*, arXiv:2403.08295.
- Yvette Graham, Timothy Baldwin, Alistair Moffat, and Justin Zobel. 2013. [Continuous measurement scales in human evaluation of machine translation](#). In *Proceedings of the 7th Linguistic Annotation Workshop and Interoperability with Discourse*, pages 33–41, Sofia, Bulgaria. Association for Computational Linguistics.
- Xu Huang, Zhirui Zhang, Xiang Geng, Yichao Du, Jiajun Chen, and Shujian Huang. 2024. [Lost in the Source Language: How Large Language Models Evaluate the Quality of Machine Translation](#). In *Findings of the Association for Computational Linguistics ACL 2024*, pages 3546–3562, Bangkok, Thailand and virtual meeting. Association for Computational Linguistics.
- Albert Q. Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, Gianna Lengyel, Guillaume Bour, Guillaume Lample, Léo Renard Lavaud, Lucile Saulnier, Marie-Anne Lachaux, Pierre Stock, Sandeep Subramanian, Sophia Yang, Szymon Antoniak, Teven Le Scao, Théophile Gervet, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. 2024. [Mixture of Experts](#). *Preprint*, arXiv:2401.04088.
- Diptesh Kanojia, Marina Fomicheva, Tharindu Ranasinghe, Frédéric Blain, Constantin Orăsan, and Lucia Specia. 2021. [Pushing the right buttons: Adversarial evaluation of quality estimation](#). In *Proceedings of the Sixth Conference on Machine Translation*, pages

- 625–638, Online. Association for Computational Linguistics.
- Tom Kocmi and Christian Federmann. 2023a. [GEMBA-MQM: Detecting translation quality error spans with GPT-4](#). In *Proceedings of the Eighth Conference on Machine Translation*, pages 768–775, Singapore. Association for Computational Linguistics.
- Tom Kocmi and Christian Federmann. 2023b. [Large language models are state-of-the-art evaluators of translation quality](#). In *Proceedings of the 24th Annual Conference of the European Association for Machine Translation*, pages 193–203, Tampere, Finland. European Association for Machine Translation.
- Jan Kocoń, Igor Cichecki, Oliwier Kaszyca, Mateusz Kochanek, Dominika Szydło, Joanna Baran, Julita Bielaniec, Marcin Gruza, Arkadiusz Janz, Kamil Kanclerz, Anna Kocoń, Bartłomiej Koptyra, Wiktoria Mieszczenko-Kowszewicz, Piotr Miłkowski, Marcin Oleksy, Maciej Piasecki, Łukasz Radliński, Konrad Wojtasik, Stanisław Woźniak, and Przemysław Kazienko. 2023. [ChatGPT: Jack of all trades, master of none](#). *Information Fusion*, 99:101861.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph Gonzalez, Hao Zhang, and Ion Stoica. 2023. [Efficient Memory Management for Large Language Model Serving with PagedAttention](#). In *Proceedings of the 29th Symposium on Operating Systems Principles, SOSP '23*, page 611–626, New York, NY, USA. Association for Computing Machinery.
- Ji Lin, Jiaming Tang, Haotian Tang, Shang Yang, Weiming Chen, Wei-Chen Wang, Guangxuan Xiao, Xingyu Dang, Chuang Gan, and Song Han. 2024. [AWQ: Activation-aware Weight Quantization for On-Device LLM Compression and Acceleration](#). In *Proceedings of Machine Learning and Systems*, volume 6, pages 87–100.
- Arle Richard Lommel, Aljoscha Burchardt, and Hans Uszkoreit. 2014. [Multidimensional Quality Metrics: A Flexible System for Assessing Translation Quality](#). *Tradumàtica: tecnologies de la traducció*, 0:455–463.
- OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altmenschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mohammad Bavarian, Jeff Belgum, Irwan Bello, Jake Berdine, Gabriel Bernadett-Shapiro, Christopher Berner, Lenny Bogdonoff, Oleg Boiko, Madelaine Boyd, Anna-Luisa Brakman, Greg Brockman, Tim Brooks, Miles Brundage, Kevin Button, Trevor Cai, Rosie Campbell, Andrew Cann, Brittany Carey, Chelsea Carlson, Rory Carmichael, Brooke Chan, Che Chang, Fotis Chantzis, Derek Chen, Sully Chen, Ruby Chen, Jason Chen, Mark Chen, Ben Chess, Chester Cho, Casey Chu, Hyung Won Chung, Dave Cummings, Jeremiah Currier, Yunxing Dai, Cory Decareaux, Thomas Degry, Noah Deutsch, Damien Deville, Arka Dhar, David Dohan, Steve Dowling, Sheila Dunning, Adrien Ecoffet, Atty Eleti, Tyna Eloundou, David Farhi, Liam Fedus, Niko Felix, Simón Posada Fishman, Juston Forte, Isabella Fulford, Leo Gao, Elie Georges, Christian Gibson, Vik Goel, Tarun Gogineni, Gabriel Goh, Rapha Gontijo-Lopes, Jonathan Gordon, Morgan Grafstein, Scott Gray, Ryan Greene, Joshua Gross, Shixiang Shane Gu, Yufei Guo, Chris Hallacy, Jesse Han, Jeff Harris, Yuchen He, Mike Heaton, Johannes Heidecke, Chris Hesse, Alan Hickey, Wade Hickey, Peter Hoeschele, Brandon Houghton, Kenny Hsu, Shengli Hu, Xin Hu, Joost Huizinga, Shantanu Jain, Shawn Jain, Joanne Jang, Angela Jiang, Roger Jiang, Haozhun Jin, Denny Jin, Shino Jomoto, Billie Jonn, Heewoo Jun, Tomer Kaftan, Łukasz Kaiser, Ali Kamali, Ingmar Kanitscheider, Nitish Shirish Keskar, Tabarak Khan, Logan Kilpatrick, Jong Wook Kim, Christina Kim, Yongjik Kim, Jan Hendrik Kirchner, Jamie Kiros, Matt Knight, Daniel Kokotajlo, Łukasz Kondraciuk, Andrew Kondrich, Aris Konstantinidis, Kyle Kosic, Gretchen Krueger, Vishal Kuo, Michael Lampe, Ikai Lan, Teddy Lee, Jan Leike, Jade Leung, Daniel Levy, Chak Ming Li, Rachel Lim, Molly Lin, Stephanie Lin, Mateusz Litwin, Theresa Lopez, Ryan Lowe, Patricia Lue, Anna Makanju, Kim Malfacini, Sam Manning, Todor Markov, Yaniv Markovski, Bianca Martin, Katie Mayer, Andrew Mayne, Bob McGrew, Scott Mayer McKinney, Christine McLeavey, Paul McMillan, Jake McNeil, David Medina, Aalok Mehta, Jacob Menick, Luke Metz, Andrey Mishchenko, Pamela Mishkin, Vinnie Monaco, Evan Morikawa, Daniel Mossing, Tong Mu, Mira Murati, Oleg Murk, David Mély, Ashvin Nair, Reiichiro Nakano, Rameez Nayak, Arvind Neelakantan, Richard Ngo, Hyeonwoo Noh, Long Ouyang, Cullen O’Keefe, Jakub Pachocki, Alex Paino, Joe Palermo, Ashley Pantuliano, Giambattista Parascandolo, Joel Parish, Emy Parparita, Alex Passos, Mikhail Pavlov, Andrew Peng, Adam Perelman, Filipe de Avila Belbute Peres, Michael Petrov, Henrique Ponde de Oliveira Pinto, Michael, Pokorny, Michelle Pokrass, Vitchyr H. Pong, Tolly Powell, Alethea Power, Boris Power, Elizabeth Proehl, Raul Puri, Alec Radford, Jack Rae, Aditya Ramesh, Cameron Raymond, Francis Real, Kendra Rimbach, Carl Ross, Bob Rotsted, Henri Roussez, Nick Ryder, Mario Saltarelli, Ted Sanders, Shibani Santurkar, Girish Sastry, Heather Schmidt, David Schnurr, John Schulman, Daniel Selsam, Kyla Sheppard, Toki Sherbakov, Jessica Shieh, Sarah Shoker, Pranav Shyam, Szymon Sidor, Eric Sigler, Maddie Simens, Jordan Sitkin, Katarina Slama, Ian Sohl, Benjamin Sokolowsky, Yang Song, Natalie Staudacher, Felipe Petroski Such, Natalie Summers, Ilya Sutskever, Jie Tang, Nikolas Tezak, Madeleine B. Thompson, Phil Tillet, Amin Tootoonchian, Elizabeth Tseng, Preston Tuggle, Nick Turley, Jerry Tworek, Juan Felipe Cerón Uribe, Andrea Vallone, Arun Vijayarvigiya, Chelsea Voss, Carroll Wainwright, Justin Jay Wang, Alvin Wang, Ben Wang, Jonathan Ward, Jason Wei, CJ Weinmann, Akila Welihinda, Peter Welinder, Jiayi Weng, Lilian Weng, Matt Wiethoff, Dave Willner,

- Clemens Winter, Samuel Wolrich, Hannah Wong, Lauren Workman, Sherwin Wu, Jeff Wu, Michael Wu, Kai Xiao, Tao Xu, Sarah Yoo, Kevin Yu, Qiming Yuan, Wojciech Zaremba, Rowan Zellers, Chong Zhang, Marvin Zhang, Shengjia Zhao, Tianhao Zheng, Juntang Zhuang, William Zhuk, and Barret Zoph. 2024. [GPT-4 Technical Report](#).
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Tharindu Ranasinghe, Constantin Orasan, and Ruslan Mitkov. 2020a. [TransQuest at WMT2020: Sentence-level direct assessment](#). In *Proceedings of the Fifth Conference on Machine Translation*, pages 1049–1055, Online. Association for Computational Linguistics.
- Tharindu Ranasinghe, Constantin Orasan, and Ruslan Mitkov. 2020b. [TransQuest: Translation quality estimation with cross-lingual transformers](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 5070–5081, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Tharindu Ranasinghe, Constantin Orasan, and Ruslan Mitkov. 2021. [An exploratory analysis of multilingual word-level quality estimation with cross-lingual transformers](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 434–440, Online. Association for Computational Linguistics.
- Ricardo Rei, José G. C. de Souza, Duarte Alves, Chrysoula Zerva, Ana C Farinha, Taisiya Glushkova, Alon Lavie, Luisa Coheur, and André F. T. Martins. 2022a. [COMET-22: Unbabel-IST 2022 submission for the metrics shared task](#). In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 578–585, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Ricardo Rei, Craig Stewart, Ana C Farinha, and Alon Lavie. 2020. [COMET: A neural framework for MT evaluation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2685–2702, Online. Association for Computational Linguistics.
- Ricardo Rei, Marcos Treviso, Nuno M. Guerreiro, Chrysoula Zerva, Ana C Farinha, Christine Maroti, José G. C. de Souza, Taisiya Glushkova, Duarte Alves, Luisa Coheur, Alon Lavie, and André F. T. Martins. 2022b. [CometKiwi: IST-unbabel 2022 submission for the quality estimation shared task](#). In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 634–645, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Thibault Sellam, Dipanjan Das, and Ankur Parikh. 2020. [BLEURT: Learning robust metrics for text generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7881–7892, Online. Association for Computational Linguistics.
- Noam Shazeer, Azalia Mirhoseini*, Krzysztof Maziarz*, Andy Davis, Quoc Le, Geoffrey Hinton, and Jeff Dean. 2017. [Outrageously Large Neural Networks: The Sparsely-Gated Mixture-of-Experts Layer](#). In *International Conference on Learning Representations*.
- Matthew Snover, Bonnie Dorr, Rich Schwartz, Linnea Micciulla, and John Makhoul. 2006. [A study of translation edit rate with targeted human annotation](#). In *Proceedings of the 7th Conference of the Association for Machine Translation in the Americas: Technical Papers*, pages 223–231, Cambridge, Massachusetts, USA. Association for Machine Translation in the Americas.
- Charles Spearman. 1904. [The Proof and Measurement of Association between Two Things](#). *The American Journal of Psychology*, 15:72–101.
- Lucia Specia, Frédéric Blain, Marina Fomicheva, Erick Fonseca, Vishrav Chaudhary, Francisco Guzmán, and André F. T. Martins. 2020. [Findings of the WMT 2020 shared task on quality estimation](#). In *Proceedings of the Fifth Conference on Machine Translation*, pages 743–764, Online. Association for Computational Linguistics.
- Lucia Specia, Frédéric Blain, Marina Fomicheva, Chrysoula Zerva, Zhenhao Li, Vishrav Chaudhary, and André F. T. Martins. 2021. [Findings of the WMT 2021 shared task on quality estimation](#). In *Proceedings of the Sixth Conference on Machine Translation*, pages 684–725, Online. Association for Computational Linguistics.
- Lucia Specia, Carolina Scarton, and Gustavo Henrique Paetzold. 2018. [Quality Estimation for Machine Translation](#). Springer, Cham, Germany.
- Craig Stewart, Ricardo Rei, Catarina Farinha, and Alon Lavie. 2020. [COMET - deploying a new state-of-the-art MT evaluation metric in production](#). In *Proceedings of the 14th Conference of the Association for Machine Translation in the Americas (Volume 2: User Track)*, pages 78–109, Virtual. Association for Machine Translation in the Americas.
- Yiming Tan, Dehai Min, Yu Li, Wenbo Li, Nan Hu, Yongrui Chen, and Guilin Qi. 2023. [Can ChatGPT Replace Traditional KBQA Models? An In-Depth Analysis of the Question Answering Performance of the GPT LLM Family](#). In *The Semantic Web – ISWC 2023*, pages 348–367, Cham. Springer Nature Switzerland.

- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. [Llama 2: Open Foundation and Fine-Tuned Chat Models](#). *Preprint*, arXiv:2307.09288.
- Yu Wan, Dayiheng Liu, Baosong Yang, Haibo Zhang, Boxing Chen, Derek Wong, and Lidia Chao. 2022. [UniTE: Unified translation evaluation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8117–8127, Dublin, Ireland. Association for Computational Linguistics.
- Guan Wang, Sijie Cheng, Xianyuan Zhan, Xiangang Li, Sen Song, and Yang Liu. 2024. [OpenChat: Advancing Open-source Language Models with Mixed-Quality Data](#). In *The Twelfth International Conference on Learning Representations*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. 2024. Chain-of-thought prompting elicits reasoning in large language models. In *Proceedings of the 36th International Conference on Neural Information Processing Systems, NIPS '22*, Red Hook, NY, USA. Curran Associates Inc.
- Chrysoula Zerva, Frédéric Blain, Ricardo Rei, Piyawat Lertvittayakumjorn, José G. C. de Souza, Steffen Eger, Diptesh Kanojia, Duarte Alves, Constantin Orăsan, Marina Fomicheva, André F. T. Martins, and Lucia Specia. 2022. [Findings of the WMT 2022 shared task on quality estimation](#). In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 69–99, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Tianyi Zhang*, Varsha Kishore*, Felix Wu*, Kilian Q. Weinberger, and Yoav Artzi. 2020. [BERTScore: Evaluating Text Generation with BERT](#). In *International Conference on Learning Representations*.
- Tianyi Zhang, Faisal Ladhak, Esin Durmus, Percy Liang, Kathleen McKeown, and Tatsunori B. Hashimoto. 2024. [Benchmarking Large Language Models for News Summarization](#). *Transactions of the Association for Computational Linguistics*, 12:39–57.
- Yutao Zhu, Huaying Yuan, Shuting Wang, Jiongnan Liu, Wenhan Liu, Chenlong Deng, Haonan Chen, Zhicheng Dou, and Ji-Rong Wen. 2023. [Large Language Models for Information Retrieval: A Survey](#). *Preprint*, arXiv:2308.07107.

A Pearson's r and Kendall's τ Correlation Scores

LP	T1		T2		T3		T4		T5		T6		T7		T8	
	r	τ	r	τ	r	τ	r	τ	r	τ	r	τ	r	τ	r	τ
OpenChat3.5																
EN-DE	0.2048	0.1613	0.2157	0.1556	0.2932	0.2153	0.2305	0.1956	0.3094	0.2135	0.3246	0.2251	0.2180	0.1746	0.2179	0.1279
EN-MR	0.2551	0.2031	0.2774	0.2023	0.5192	0.2757	0.3351	0.2560	0.4463	0.2711	0.4919	0.2654	0.3669	0.2173	0.3529	0.1791
EN-ZH	0.2921	0.2267	0.2655	0.1894	0.4001	0.2905	0.3063	0.2191	0.3702	0.2406	0.3865	0.2617	0.3487	0.2364	0.2971	0.1938
ET-EN	0.5474	0.4107	0.5966	0.4418	0.6776	0.5249	0.5753	0.4341	0.6617	0.5043	0.6679	0.5213	0.5937	0.4452	0.5610	0.4146
NE-EN	0.3606	0.2705	0.4968	0.3617	0.6185	0.4364	0.5447	0.3649	0.6114	0.4314	0.6246	0.4478	0.5547	0.3849	0.4809	0.3287
RO-EN	0.4312	0.3664	0.5307	0.3971	0.7891	0.5642	0.7353	0.5359	0.2487	0.5516	0.7983	0.5716	0.7634	0.5433	0.7491	0.5097
RU-EN	0.5893	0.4714	0.4760	0.3630	0.6655	0.4761	0.4957	0.4290	0.5926	0.4183	0.6643	0.4630	0.6016	0.4009	0.6413	0.4199
SI-EN	0.4060	0.3060	0.4351	0.3135	0.6001	0.4492	0.4434	0.3147	0.5957	0.4449	0.5920	0.4395	0.5139	0.3811	0.4165	0.2987
Llama-2-7B																
EN-DE	0.0436	0.0344	0.0861	0.0642	0.0216	0.0662	0.0494	-0.0110	0.0132	0.0693	-0.0185	0.0686	-0.0225	-0.0259	-0.0914	0.0235
EN-MR	0.0763	0.0597	0.0481	0.0361	-0.0484	0.0931	0.0271	0.0122	0.0127	0.0647	0.0447	0.0686	-0.0096	-0.0017	0.0222	0.0056
EN-ZH	0.0818	0.0623	0.0478	0.0361	0.0938	0.0712	-0.0069	0.0208	0.0121	0.0459	0.0470	0.0579	0.0393	0.0261	0.0368	0.0194
ET-EN	0.1231	0.0978	0.2589	0.1976	0.0452	0.2805	0.0382	-0.0301	0.0593	0.1894	0.0574	0.2418	0.1224	0.1065	-0.0529	-0.0023
NE-EN	0.2156	0.1680	0.1729	0.1308	0.1288	0.1670	0.0617	0.0344	0.0164	0.1503	0.0182	0.1939	0.1348	0.0995	0.0311	0.0994
RO-EN	0.2662	0.2072	0.2364	0.1807	0.0886	0.3404	-0.0514	-0.0068	0.0571	0.1772	0.0030	0.2622	0.0718	0.1358	-0.0357	0.0285
RU-EN	0.2342	0.1891	0.2084	0.1564	0.3123	0.3234	0.0273	-0.0040	0.0270	0.1632	0.0030	0.1759	0.1879	0.1182	0.0531	0.1167
SI-EN	0.1295	0.1019	0.0382	0.0300	0.1345	0.1638	0.0852	0.0428	0.1594	0.1713	0.1287	0.1867	0.1719	0.1339	0.0725	0.0250
Gemma-7B																
EN-DE	0.1217	0.0984	0.1578	0.1239	0.1622	0.1280	0.0994	0.0844	0.1696	0.1436	0.1693	0.1424	0.0193	0.0040	-0.0066	0.0244
EN-MR	0.1745	0.1428	0.2018	0.1599	0.2114	0.1189	0.2576	0.1256	0.2634	0.1747	0.2024	0.1468	0.1379	0.0662	0.0456	0.1182
EN-ZH	0.2724	0.2121	0.2037	0.1581	0.1100	0.1406	0.2622	0.2156	0.2091	0.1876	0.2171	0.1826	0.1516	0.1228	0.0372	0.1274
ET-EN	0.3837	0.3003	0.4749	0.3721	0.3452	0.2954	0.3922	0.3237	0.5009	0.4294	0.4522	0.3979	0.3483	0.2382	0.1742	0.2927
NE-EN	0.3794	0.3002	0.3242	0.2595	0.2635	0.2244	0.3204	0.2658	0.3677	0.3191	0.3399	0.3111	0.3051	0.2425	0.0448	0.2217
RO-EN	0.5852	0.4552	0.4672	0.3662	0.4127	0.3473	0.6256	0.4585	0.5523	0.4558	0.5630	0.4600	0.5137	0.3468	0.1365	0.3430
RU-EN	0.4205	0.3294	0.4479	0.3448	0.3331	0.3452	0.3223	0.2900	0.4826	0.3497	0.4614	0.3965	0.3384	0.2273	0.0591	0.1829
SI-EN	0.2902	0.2298	0.2876	0.2245	0.3705	0.2737	0.2831	0.2162	0.3879	0.3104	0.3740	0.2862	0.2883	0.2086	0.0476	0.2169
Llama-2-13B																
EN-DE	0.0612	0.0436	0.0111	0.0082	0.0527	0.0229	0.0467	0.0534	0.0246	0.0874	0.0386	0.0735	0.0337	0.0266	0.0018	0.0487
EN-MR	-0.0025	-0.0021	-0.0324	-0.0248	0.026	0.0501	-0.0062	0.0143	0.0341	0.0824	0.0506	0.0757	0.0957	0.0607	0.0349	0.0120
EN-ZH	0.0085	0.0067	0.0086	0.0067	0.0889	0.1033	-0.0201	0.0593	0.0496	0.0744	0.0486	0.0699	0.0455	0.0521	0.0657	0.0628
ET-EN	0.2339	0.1821	0.1992	0.1440	0.3263	0.3051	0.0239	0.1603	0.0542	0.2841	0.088	0.3112	0.2104	0.1875	-0.031	0.1707
NE-EN	0.0619	0.0492	0.3039	0.2229	0.2466	0.2359	0.0538	0.0842	0.0865	0.0933	0.038	0.2045	0.3043	0.2448	0.0385	0.1406
RO-EN	0.3048	0.2332	0.2916	0.2141	0.4189	0.3422	0.0302	0.1975	0.0698	0.2414	0.0463	0.3666	0.3420	0.2538	-0.0267	0.1936
RU-EN	0.404	0.3118	0.1973	0.1442	0.3916	0.3021	0.0843	0.2463	0.0562	0.2331	0.3697	0.3148	0.2766	0.1899	0.0645	0.1493
SI-EN	0.0295	0.0212	0.3414	0.2501	0.2160	0.1992	0.0726	0.0342	-0.063	0.0731	0.0705	0.2622	0.2478	0.2116	0.0827	0.1034
Qwen1.5-14B																
EN-DE	0.1555	0.1219	0.2543	0.1975	0.1504	0.1625	0.0717	0.1173	0.0143	0.1652	0.2114	0.1746	0.2089	0.1777	0.0254	0.1365
EN-MR	0.2572	0.2053	0.3300	0.2634	0.4550	0.2395	0.0407	0.1551	0.0161	0.2120	0.4507	0.2332	0.4174	0.2648	0.1059	0.2114
EN-ZH	0.3839	0.2914	0.2615	0.1976	0.3655	0.3028	0.0557	0.2245	0.0297	0.2497	0.4011	0.3212	0.3580	0.2479	0.0038	0.2154
ET-EN	0.5134	0.3867	0.5896	0.4510	0.5923	0.4796	0.0769	0.3477	-0.0069	0.4063	0.1954	0.4871	0.5215	0.4336	0.0487	0.3170
NE-EN	0.3265	0.2376	0.5208	0.3857	0.5149	0.3741	0.1401	0.2391	0.1960	0.3282	0.4868	0.3338	0.4712	0.3578	0.0484	0.2395
RO-EN	0.5609	0.4387	0.5621	0.4348	0.7127	0.5527	0.0985	0.4176	0.0466	0.4385	0.7293	0.5345	0.7369	0.5479	0.1085	0.3300
RU-EN	0.5047	0.3982	0.4997	0.3979	0.6002	0.4332	0.1075	0.3620	0.1665	0.4161	0.6100	0.4282	0.5546	0.3957	0.3920	0.3728
SI-EN	0.3693	0.2669	0.5096	0.3773	0.5820	0.4366	0.0609	0.2002	0.1472	0.3012	0.5249	0.3997	0.4340	0.3537	0.1646	0.3033
Mixtral-8x7B-Instruct																
EN-DE	0.1980	0.1444	0.1658	0.1195	0.0189	0.1959	-0.0161	0.1137	-0.0713	0.1388	0.0404	0.1709	-0.0827	0.1625	0.2556	0.1631
EN-MR	0.1819	0.1405	0.2115	0.1539	0.0186	0.1394	0.0140	0.0634	0.0230	0.1493	0.0707	0.1469	0.0025	0.1428	0.0299	0.1917
EN-ZH	0.3324	0.2424	0.2927	0.2109	0.0890	0.2687	-0.0062	0.1496	0.0654	0.1926	-0.0061	0.2474	0.2443	0.1813	0.3079	0.2200
ET-EN	0.4748	0.3531	0.5552	0.4095	0.2111	0.4614	0.0380	0.2517	0.0738	0.3171	0.1534	0.4336	0.2656	0.3399	0.0567	0.4052
NE-EN	0.3909	0.2872	0.4215	0.3070	0.1190	0.3572	0.0543	0.0842	-0.0219	0.1629	0.1416	0.3409	0.4209	0.3223	0.4284	0.3824
RO-EN	0.5576	0.4344	0.5416	0.4107	0.0427	0.5000	0.0748	0.3024	0.0991	0.3196	0.0947	0.4727	0.6549	0.4879	0.2054	0.5154
RU-EN	0.4823	0.3630	0.4514	0.3422	0.1760	0.4510	0.0854	0.3055	0.1199	0.3767	0.1844	0.4257	0.5816	0.3913	0.0212	0.3984
SI-EN	0.3352	0.2415	0.4459	0.3223	0.0961	0.3343	0.0072	0.0691	0.0934	0.1555	0.1099	0.2997	0.4255	0.3461	0.4334	0.3441

Table 7: Pearson's r and Kendall's τ correlation scores achieved using Templates 1-8 (T1-8) on various open-source LLMs for each language pair (LP).

Model input (before formatting):

Score the following translation from English to Chinese with respect to the human reference on a continuous scale from 0 to 100, where score of zero means "no meaning preserved" and score of one hundred means "perfect meaning and grammar". \nEnglish source: The last conquistador then rides on with his sword drawn.\nChinese human reference: 最后的征服者随后举着剑前进。 \nChinese translation: 最后的征服者骑着他的剑继续前进.\nScore:

Model output:

```
<lim_start>user\nScore the following translation from English to Chinese with respect to the human reference on a continuous scale from 0 to 100, where score of zero means "no meaning preserved" and score of one hundred means "perfect meaning and grammar". \nEnglish source: The last conquistador then rides on with his sword drawn.\nChinese human reference: 最后的征服者随后举着剑前进。 \nChinese translation: 最后的征服者骑着他的剑继续前进.\nScore:<lim_endl>\n<lim_start>assistant\n最后的征服者骑着他的剑继续前进.
```

Figure 6: An example of the TowerLLM output for scoring English-to-Chinese translation using Template 3 via HuggingFace. The output was generated in March, 2024.