Large Language Models Are State-of-the-Art Evaluators of Translation Quality

Tom Kocmi and **Christian Federmann**

Microsoft, One Microsoft Way, Redmond, WA-98052, USA

{tomkocmi,chrife}@microsoft.com

Abstract

We describe GEMBA, a GPT-based metric for assessment of translation quality, which works both with a reference translation and without. In our evaluation, we focus on zero-shot prompting, comparing four prompt variants in two modes, based on the availability of the reference. We investigate nine versions of GPT models, including ChatGPT and GPT-4. We show that our method for translation quality assessment only works with GPT 3.5 and larger models. Comparing to results from WMT22's Metrics shared task, our method achieves state-of-the-art accuracy in both modes when compared to MQM-based human labels. Our results are valid on the system level for all three WMT22 Metrics shared task language pairs, namely English into German, English into Russian, and Chinese into English. This provides a first glimpse into the usefulness of pre-trained, generative large language models for quality assessment of translations. We publicly release all our code and prompt templates used for the experiments described in this work, as well as all corresponding scoring results, to allow for external validation and reproducibility.[1]

1   Introduction

One of the interesting properties of large language models (LLMs) such as GPT (Brown et al., 2020b) is their (implicit) support for multilingual Q&A. Prompting the model in the right way allows us to translate text between languages (Vilar et al., 2022). This is surprising as GPT has not been fine-tuned for the translation task.

Hendy et al. (2023) show that GPT-enabled translation achieves high quality when applied for the translation of high-resource languages, but still lacks in terms of translation quality for under-represented languages. Building on this finding—*if the model can translate, it may be able to differen-tiate good from bad translations*—we apply GPT for the task of translation quality assessment.

In the remainder of this paper, inspired by recent progress on generative, pre-trained large language models (LLMs), we explore how these models can be applied for automated assessment of translation quality. The primary query for this study centers around the question: *Can LLMs be used for effective quality assessment of translations?*

We propose GEMBA, which stands for *GPT Estimation Metric Based Assessment*. The metric evaluates each segment translation in isolation and then averages across all obtained scores for the final, system-level score.

We define and evaluate several prompt variants for zero-shot assessment of translation quality in two modes, either with a human reference translation, as a quality metric, or without one, as a quality estimation task.

We design the main prompts based on the DA+SQM template used for human assessment of translation quality as implemented in the Appraise framework (Federmann, 2018) for WMT22 (Kocmi et al., 2022), building on previous work conducted by Freitag et al. (2021a).

To the best of our knowledge, our research represents the pioneering effort in exploring the utilization of large language models (LLMs) for the purpose of quality assessment. Subsequent to the publishing of our findings, Lu et al. (2023) independently published a related report, corroborating the high performance of LLMs.

The main contributions of this paper are:

- We demonstrate state-of-the-art capabilities of GPT-based translation quality assessment on the latest WMT22 metrics evaluation data (on the system level);

- We experiment with four prompt templates, showing that the least constrained template achieves the best performance;

---

[1] https://github.com/MicrosoftTranslator/GEMBA

- We evaluate nine different models of GPT, showing that only GPT 3.5 and larger models are capable of translation quality assessment;

- We show that GEMBA with GPT-4 model is only slightly behind on segment-level scores to the best-performing metrics.

2 The GEMBA Metric

To assess translation quality via prompting an LLM, the following parameters are needed:

- prompt variant (from a pre-defined set)
- source language name, e.g., "Chinese"
- target language name, e.g., "English"
- source segments $src_{1..N}$
- candidate translations $hyp_{1..N}$
- optionally, reference translations $ref_{1..N}$

We generate a GPT request for every segment, querying as individual zero-shot problems, and then aggregate results. For this initial proof of concept, we leave improvements such as few-shot queries or document-level context to future work.

2.1 Prompt variants

We experiment with four distinct prompt types: modeling two scoring and two classification tasks. For the scoring tasks, first, one based on **direct assessment** (*GEMBA-DA*), second, another based on recent research efforts on **scalar quality metrics** (*GEMBA-SQM*).[2] As scoring translation quality may be an unnatural task for an LLM, we also design two classification tasks. The first is based on **one-to-five stars ranking** (*GEMBA-stars*), which is a style often used when users are asked to review various services or products. The second prompt asks the LLM to label translation quality as one of five discrete **quality classes** (*GEMBA-classes*).

For each of these four prompt types, we experiment with two modes that differ with respect to the wording of the corresponding query templates which either have access to a human reference or not. As an example, we show the *GEMBA-DA* prompt in Figure 1. Based on token count, this is the least constrained prompt template that we experiment with. The complete set of prompt templates is available in Appendix A. For naming convention, we mark quality estimation metrics (without reference) with the suffix "[noref]".

2.2 Scoring process

The expected scores are in $[0, 100]$ for *GEMBA-DA* and *GEMBA-SQM* prompts, same as for human assessment (Graham et al., 2013); for *GEMBA-stars* the output ranges from $[1, 5]$ and *GEMBA-classes* assigns one of five class labels.

We average segment-level scores to obtain system-level scores. For the *GEMBA-classes* metric variant, we assign classes a numerical value $[0 - 4]$, based on the label, before averaging.

Depending on the GPT model we query, sometimes answers are returned outside these ranges, as text. When we observe such an *invalid* answer, we add randomness and sample more responses, selecting the first answer matching the output range as the final result.

2.3 GPT models

We experiment with seven GPT models—*ranging from GPT 2 up to the latest GPT-4 model*—that are described in Table 1.[3] We use the GPT-4 model as the default model for most experiments and compare the performance of other models in Section 4.3. Specifically, we use these models with brief description:

GPT 2 We use models provided by Radford et al. (2019), assessing if GPT 2 may be useful for quality assessment—*we find that it is not*;

Ada GPT 3. Max request size of 2,048 tokens and training data up to October 2019 (Brown et al., 2020a);

Babbage GPT 3. More capable than Ada (Brown et al., 2020a);

Curie GPT 3. More capable than Babbage (Brown et al., 2020a);

Davinci-002 GPT 3.5. Max request size of 4,000 tokens and training data up to June 2021. Uses `FeedME` training;

ChatGPT Improved GPT 3.5 model, fine-tuned using Reinforcement Learning from Human Feedback (RLHF);

Davinci-003 GPT 3.5.1. Uses `PPO` training;

GPT-3.5-turbo Davinci-003 model optimized for speed;

GPT-4 there is only limited information about GPT-4, see OpenAI (2023).

```
Score the following translation from {source_lang} to {target_lang} 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".

{source_lang} source: "{source_seg}"
{target_lang} human reference: {reference_seg}
{target_lang} translation: "{target_seg}"
Score:
```

Figure 1: The best-performing prompt based on Direct Assessment expecting a score between 0–100. Template **portions in bold face** are used only when a human reference translation is available.

Model name	Abbrev.	Model used
GPT-2	—	Radford et al. (2019)
Ada	—	text-ada-001
Babbage	Bab	text-babbage-001
Curie	Curie	text-curie-001
Davinci-002	Dav2	text-davinci-002
ChatGPT	Chat	text-chat-davinci-002
Davinci-003	Dav3	text-davinci-003
GPT-3.5-turbo	Turbo	gpt-3.5-turbo
GPT-4	GPT4	gpt-4

Table 1: Details of all models used in this work. Models are sorted from oldest to newest which also reflects their respective power. GPT 2 and Ada do not work.

GPT 3 models are based on Ouyang et al. (2022). The models are sorted based on their estimated power or date of release. We acknowledge that OpenAI has not released detailed information about the architecture and training data behind given models. Most importantly, OpenAI claims that models have been trained with data up until September 2021. It is important as we use testsets prepared and released by December 2022.

3 Experiments

To measure the performance of the proposed GEMBA metric, we follow the methodology and use test data provided by the WMT22 Metrics shared task (Freitag et al., 2022b) which hosts an annual evaluation of automatic metrics, benchmarking them against human gold labels. Effectively, we compare GEMBA against the best-performing automatic metrics: COMET (Rei et al., 2020, 2022), BLEURT (Sellam et al., 2020), or the non-public winner MetricX XXL.

3.1 Test set

We use the MQM 2022 test set which contains human judgments for the following three translation directions: English into German, English into Russian, and Chinese into English. The test set contains a total of 54 machine translation system outputs or human translations. It contains a total of 106k segments. Translation systems are mainly from participants of the WMT22 General MT shared task (Kocmi et al., 2022).

The source segments and human reference translations for each language pair contain around 2,000 sentences from four different texts domains: news, social, conversational, and e-commerce. The gold standard for scoring translation quality is based on human MQM ratings, annotated by professionals who mark individual errors in each translation, as described in Freitag et al. (2021a).

3.2 Evaluation methods

To determine how well automatic metrics correlate with humans, we measure system-level, pairwise accuracy (*accuracy*, Kocmi et al., 2021). For segment-level evaluation, we use Kendall's Tau (τ, Freitag et al., 2022a).

Here, accuracy is defined as the number of system pairs ranked correctly by the metric with respect to the human ranking divided by the total number of system pair comparisons.

Formally:

$$\text{Accuracy} = \frac{|\text{sign}(\text{metric}\Delta) == \text{sign}(\text{human}\Delta)|}{|\text{all system pairs}|}$$

The variant of Kendall's Tau used for metric evaluation changed over the years. Initially, Callison-Burch et al. (2011) proposed to use Kendall's Tau-a ignoring human rankings that tied, while penalising ties in automatic metrics.

$$\tau = \frac{|\text{Concordant}| - |\text{Discordant}|}{|\text{Concordant}| + |\text{Discordant}|}$$

where Concordant is the set of all human segment comparisons for which a given metric suggests the same order of systems and Discordant is the set of all human comparisons for which a given metric disagrees.

		Metric		
		$s_1 < s_2$	$s_1 = s_2$	$s_1 > s_2$
Human	$s_1 < s_2$	Conc	Disc	Disc
	$s_1 = s_2$	–	–	–
	$s_1 > s_2$	Disc	Disc	Conc

This definition was later updated by Macháček and Bojar (2014), who handle ties as a separate group in contrast to Concordant and Discordant. Metrics shared tasks Mathur et al. (2020) and Freitag et al. (2021b) changed this back to the 2011 version. Last year, Freitag et al. (2022a) changed it to Kendall's Tau-b, which makes adjustments for ties, we use the latest definition in our experiments. Overall, ties in automatic metrics are rare for non-identical translations but are an issue when a method outputs only a discrete set of scores (as in our case). Additionally, Kendall's Tau is susceptible to noise in gold pairwise rankings (Freitag et al., 2022a).

We reproduced all scores reported in the WMT22 Metrics shared task findings paper with the official WMT22 script.[4] Reported scores match Table 11 of the WMT22 metrics findings paper (Freitag et al., 2022b).

4 Results

We investigate GEMBA's performance for two modes: with a reference translation and without reference translation (in a quality estimation setting). Table 2 reports pairwise accuracy on the system level, comparing *GEMBA-DA* against the best-performing metrics from the WMT22 Metrics shared task (Freitag et al., 2022b). We use GPT-4 as the main model and *GEMBA-DA* as the main style for some experiments.

4.1 Reference-based

The results in Table 2 show that our reference-based **GEMBA-GPT4-DA** metric sets a new state of the art. It outperforms all of the other reference-based metrics from the WMT22 Metrics shared task. The observed level of metric performance is unexpected, especially considering that human labels used as a gold standard in itself are noisy and therefore an accuracy of 100% is impossible to obtain for an automatic metric.

4.2 Quality estimation

Table 2 shows that our reference-less metric **GEMBA-GPT4-DA[noref]** achieves the highest

Metric	Accuracy
GEMBA-GPT4-DA	89.8%
GEMBA-GPT4-DA[noref]	87.6%
MetricX XXL	85.0%
BLEURT-20	84.7%
COMET-22	83.9%
COMET-20	83.6%
UniTE	82.8%
MS-COMET-22	82.8%
MATESE	81.0%
YiSi-1	79.2%
COMETKiwi[noref]	78.8%
COMET-QE[noref]	78.1%
BERTScore	77.4%
UniTE-src[noref]	75.9%
MS-COMET-QE-22[noref]	75.5%
MATESE-QE[noref]	74.8%
f200spBLEU	74.1%
chrF	73.4%
BLEU	70.8%

Table 2: Results for the system-level pairwise accuracy compared to the current automatic metric. Metrics marked as "[noref]" do not use a reference translation.

performance for the quality estimation mode, and strongly outperforms all of the other reference-less metrics. Moreover, it also outperforms all of the other reference-based metrics, performing only slightly worse than **GEMBA-GPT4-DA**. Again, the observed level of assessment quality is unexpectedly high, highlighting the potential of using LLMs for translation quality assessment tasks.

4.3 Comparison of GPT models

We compare the performance of various GPT versions as an automatic metric. Table 3 shows results for all models we have experimented with and all prompt variants tested.

We do not show results for GPT-2 or Ada models. Neither of those have produced replies in the specific scoring range and neither seemed to be producing any meaningful replies. We list a couple of their answers in Appendix C. Based on our experiments, we conclude that they are not powerful enough to understand the zero-shot prompts.

By contrast, Babbage and Curie models appear to understand what type of answer they should produce, but the quality of their scores seems to be close to random guessing. Thus, both Babbage and Curie are useless for translation quality assessment.

The main performance jump occurs for GPT 3.5 and larger models, i.e., Davinci-002, ChatGPT, Davinci-003, Turbo, and GPT-4. Each of them achieves highly competitive results for all of the prompt variants we have tested. Interestingly, Chat-

	Bab	Curie	Dav2	Chat	Dav3	Turbo	GPT4
DA	39.1%	54.4%	**85.8%**	81.0%	**88.0%**	**86.5%**	**89.8%**
DA[noref]	55.8%	51.8%	83.9%	82.1%	**86.1%**	**86.9%**	**87.6%**
SQM	51.8%	40.5%	**85.8%**	**85.0%**	85.4%	**87.2%**	**88.7%**
SQM[noref]	51.1%	41.6%	82.8%	81.0%	82.5%	**87.6%**	**89.1%**
Stars	48.2%	37.2%	**88.3%**	**85.0%**	85.8%	**89.4%**	**91.2%**
Stars[noref]	58.4%	54.7%	79.6%	83.6%	83.2%	84.3%	**89.1%**
Classes	47.4%	43.4%	79.6%	**87.2%**	85.4%	82.5%	**89.1%**
Classes[noref]	35.0%	61.7%	78.1%	83.6%	78.8%	62.0%	**91.2%**

Table 3: Accuracy of the system-level pairwise accuracy for quality estimation methods for most combinations of prompts and different GPT models. The evaluation is based on three language pairs and MQM human labels. All results higher than the WMT22 winner of Metrics shared task MetricX XXL are bolded.

GPT in DA style appears to have the lowest quality among those models. In addition, ChatGPT and Turbo frequently reply with a score followed by an explanation of why it has assigned that score. One possible reason may be in the form of the prompt, which wasn't modified to instruct ChatGPT not to generate an explanation.

Unsurprisingly, the best performance is obtained by the most powerful LLM, GPT-4. Moreover, we can see that over time, each generation of models is slightly better. This confirms the findings of Hendy et al. (2023) who demonstrated superior translation capabilities with Davinci-003 over all other previous GPT variants.

4.4 Segment-level performance

All previous results are reported on the system level. We also investigate how well the GEMBA metric performs on the segment level, with respect to the human gold annotations. We present Kendall's Tau results for each language pair separately in Table 4 for GPT-4 and Davinci-003 (results for all metrics are in Appendix B).

GPT-4 models are slightly behind the top-performing metrics but continue to have a high correlation with human judgment. On the other hand, quality estimation **GEMBA-Dav3-DA [noref]** has significantly lower segment-level performance in contrast to other top-performing metrics.

The lower performance of a segment-level correlation could be attributed to Kendall's Tau, which penalizes ties. Our metric in contrast to other automatic metrics returns a discrete value between 0–100. There is a high probability that two translations will obtain an equal score.

In order to investigate this further, we collect all answers across all systems and all three language pairs and then calculate the frequency of each distinct answer value.

Metric	Acc	en-de	en-ru	zh-en
GEMBA-GPT4-DA	89.8%	0.36	0.36	0.38
GEMBA-Dav3-DA	88.0%	0.31	0.33	0.37
GEMBA-GPT4-DA[noref]	87.6%	0.31	0.40	0.41
GEMBA-Dav3-DA[noref]	86.1%	0.18	0.26	0.29
MetricX XXL	85.0%	0.36	**0.42**	**0.43**
BLEURT-20	84.7%	0.34	0.36	0.36
COMET-22	83.9%	**0.37**	0.40	**0.43**
UniTE	82.8%	**0.37**	0.38	0.36
COMETKiwi[noref]	78.8%	0.29	0.36	0.36
COMET-QE[noref]	78.1%	0.28	0.34	0.36
chrF	73.4%	0.21	0.17	0.15
BLEU	70.8%	0.17	0.14	0.14

Table 4: Kendall's Tau (τ) segment-level evaluation. Full results are in Appendix B.

We can notice several interesting observations in Table 5. The DA reference-based prompt generates mostly multiples of five. Over three-quarters of all scores are either score 80, 95, or 100. This could reflect the actual quality of the system translations as the underlying systems are provably high-quality. This is also a finding of Freitag et al. (2022b) that many metrics fall into the same significance cluster.

When we investigate the "DA[noref]", we notice that 60.5% of all scores are of value "95". Despite this fact, the metric still manages to differentiate the systems from each other and outperform all other quality estimation metrics on the system level. This is contributed to the fact that better-performing systems obtain more segments with a score 95 than worse-performing systems, therefore getting a lower average score. We should note, that there are no system-level ties.

We conjecture that frequent segment-level ties and the discrete scale thus may contribute to the lower Kendall's Tau segment-level performance.

4.5 Failure rate

As we described earlier, LLMs may answer with an invalid answer, for example with a textual answer instead of a score, mostly explaining its decision. When such a situation happens, we iteratively in-

Answers	DA	DA[noref]	SQM	SQM[noref]
0	0.1%	0.1%	0.1%	0.1%
5	0.0%	0.0%	0.0%	0.0%
10	0.0%	0.0%	0.0%	0.1%
15	—	—	0.0%	0.0%
20	0.2%	0.3%	0.2%	0.3%
25	—	—	0.0%	—
30	0.1%	0.2%	0.1%	0.1%
35	—	—	0.0%	—
40	0.5%	0.6%	0.5%	0.6%
45	0.0%	0.0%	0.0%	0.0%
50	0.0%	0.1%	0.1%	0.0%
55	—	—	0.0%	—
60	2.1%	2.3%	2.0%	2.1%
65	—	0.0%	0.0%	0.0%
70	1.3%	0.4%	1.9%	0.6%
75	0.5%	1.0%	0.7%	0.7%
80	6.3%	4.5%	7.0%	5.7%
85	4.4%	2.7%	6.0%	2.9%
87	—	—	0.0%	—
88	—	—	0.0%	—
90	21.3%	13.0%	27.6%	14.5%
92	—	—	0.0%	—
93	—	—	0.0%	—
94	—	—	0.0%	—
95	53.3%	60.6%	44.6%	49.4%
98	0.8%	0.0%	0.4%	0.0%
99	0.4%	—	0.2%	—
100	8.6%	14.1%	8.5%	22.8%

Table 5: Distribution of all distinct segment-level score values for MQM 2022 for model GPT-4.

crease the temperature—*adding randomness to the model*—and take the first answer matching the expected score output range.

This adds non-determinism to our evaluation, therefore we investigate how frequently this phenomenon happens. Table 6 shows the number of invalid answers. For almost all combinations of models and prompts, except of SQM-style, LLMs understand the prompt and provide answers in a valid range with less than 1% of the answers being invalid.[5] This has a minimal effect on the final system-level score and therefore, we conclude that the metric is mostly deterministic.

Additionally, we confirm that a temperature equal to zero always returns the same answer, which we evaluated by re-running GEMBA-Dav2-DA[noref].

Processing answers is straightforward as it is usually a stand-alone number. In some occasions, LLMs give a numerical score and continue with a textual explanation, for such cases, we parse only the first number. A more complex approach needs to be taken for *GEMBA-stars* prompts where the model provides different answers which we parse separately. Here are some examples of two-star answers: "2", "two", "**", "★★", "two stars", or "2 stars". For non-English target languages the answer may be produced in the target language, e.g.,

---
[5]Roughly 1,000 answers equal to 1% of the total volume.

"一星", or "五". We have not observed attempts to translate output for other prompts.

5 Conclusion

We have presented our work on GEMBA, a GPT-based estimation metric-based assessment method. Comparing our metrics to other automated metrics from the WMT22 Metrics shared task we report state-of-the-art performance on the MQM 2022 test set across three language pairs: English to German, English to Russian, and Chinese to English.

We intend to continue research on the application of GPT models for quality assessment. Further research will focus on the switch to few-shot (as opposed to our current zero-shot methodology) as well as model fine-tuning. Both of which promise to increase GEMBA accuracy. Furthermore, modifying prompts to support MQM error-based evaluation or post-editing efforts may lead to further improvements.

GPT-enhanced evaluation metrics may allow us to make progress with respect to document-level evaluation (due to their ability to use much larger context windows). This could be beneficial as there is little research into document-level metrics (Vernikos et al., 2022).

Limitations

While preliminary results indicate that the GEMBA metric performs very well when compared to other automated metrics evaluated as part of the WMT22 Metrics shared task, it is important to note that these results are based on human labels for *only three language pairs*. We expect that the metrics performance may suffer for other language pairs, mainly under-resourced languages similar to Hendy et al. (2023) who show low translation quality for such languages. In addition, GEMBA's state-of-the-art performance only holds for the system level, while segment-level scores still have room for improvement. Reported results are indicative of the potential performance LLMs could achieve for the translation quality assessment task in the long run. However, more analysis is needed before using it as the main tool for deciding translation quality.

An additional limitation to consider in this study is the inability to definitively ascertain that the evaluation data have not been included in OpenAI's training dataset. Nevertheless, the available evidence strongly indicates that this is unlikely. OpenAI claims that their data compilation only extends

	Bab	Curie	Dav2	Chat	Dav3	GPT4
DA	750	8,048	7	565	0	0
DA[noref]	146	862	0	935	53	0
SQM	89,599	129	4,827	45	1,279	—
SQM[noref]	15,577	95,131	1,763	59	1	0
Stars	18,074	—	135	1,064	58	—
Stars[noref]	—	86,593	135	1,924	1	0
Classes	74	12	0	10	0	—
Classes[noref]	115	15	0	12	0	—

Table 6: Number of invalid answers (full set size 106,758) that needed to be re-prompted with added randomness. The evaluation of ChatGPT and parts of GPT-4 were excluded due to their late integration and changes in our codebase.

up to September 2021, while the test set employed in this research was generated during the second half of 2022 and made publicly available in December 2022. Our initial positive results using the Davinci-002 model were obtained in early February, which presents a narrow timeframe for OpenAI to incorporate and process the evaluation data. Furthermore, the test set is not readily accessible in plaintext format, necessitating pre-processing prior to utilization in training.

Acknowledgments

References

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020a. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.

Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020b. Language models are few-shot learners.

Chris Callison-Burch, Philipp Koehn, Christof Monz, and Omar Zaidan. 2011. Findings of the 2011 workshop on statistical machine translation. In *Proceedings of the Sixth Workshop on Statistical Machine Translation*, pages 22–64, Edinburgh, Scotland. Association for Computational Linguistics.

Christian Federmann. 2018. Appraise evaluation framework for machine translation. In *Proceedings of the 27th International Conference on Computational Linguistics: System Demonstrations*, pages 86–88, Santa Fe, New Mexico. Association for Computational Linguistics.

Markus Freitag, George Foster, David Grangier, Viresh Ratnakar, Qijun Tan, and Wolfgang Macherey. 2021a. Experts, errors, and context: A large-scale study of human evaluation for machine translation. *Transactions of the Association for Computational Linguistics*, 9:1460–1474.

Markus Freitag, Ricardo Rei, Nitika Mathur, Chi-kiu Lo, Craig Stewart, Eleftherios Avramidis, Tom Kocmi, George Foster, Alon Lavie, and André F. T. Martins. 2022a. Results of WMT22 metrics shared task: Stop using BLEU – neural metrics are better and more robust. In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 46–68, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.

Markus Freitag, Ricardo Rei, Nitika Mathur, Chi-kiu Lo, Craig Stewart, Eleftherios Avramidis, Tom Kocmi, George Foster, Alon Lavie, and Andre F. T. Martins. 2022b. Results of wmt22 metrics shared task: Stop using bleu neural metrics are better and more robust. In *Proceedings of the Seventh Conference on Machine Translation*, pages 46–68, Abu Dhabi. Association for Computational Linguistics.

Markus Freitag, Ricardo Rei, Nitika Mathur, Chi-kiu Lo, Craig Stewart, George Foster, Alon Lavie, and Ondřej Bojar. 2021b. Results of the WMT21 metrics shared task: Evaluating metrics with expert-based human evaluations on TED and news domain. In *Proceedings of the Sixth Conference on Machine Translation*, pages 733–774, Online. Association for Computational Linguistics.

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.

Amr Hendy, Mohamed Abdelrehim, Amr Sharaf, Vikas Raunak, Mohamed Gabr, Hitokazu Matsushita, Young Jin Kim, Mohamed Afify, and Hany Hassan Awadalla. 2023. How good are gpt models at machine translation? a comprehensive evaluation. *arXiv preprint arXiv:2302.09210*.

Tom Kocmi, Rachel Bawden, Ondrej 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 Novak, Martin Popel, Maja Popovic, and Mariya Shmatova. 2022. Findings of the 2022 conference on machine translation (wmt22). In *Proceedings of the Seventh Conference on Machine Translation*, pages 1–45, Abu Dhabi. Association for Computational Linguistics.

Tom Kocmi, Christian Federmann, Roman Grundkiewicz, Marcin Junczys-Dowmunt, Hitokazu Matsushita, and Arul Menezes. 2021. To ship or not to ship: An extensive evaluation of automatic metrics for machine translation. In *Proceedings of the Sixth Conference on Machine Translation*, pages 478–494, Online. Association for Computational Linguistics.

Qingyu Lu, Baopu Qiu, Liang Ding, Liping Xie, and Dacheng Tao. 2023. Error analysis prompting enables human-like translation evaluation in large language models: A case study on chatgpt. *arXiv preprint*.

Matouš Macháček and Ondřej Bojar. 2014. Results of the WMT14 metrics shared task. In *Proceedings of the Ninth Workshop on Statistical Machine Translation*, pages 293–301, Baltimore, Maryland, USA. Association for Computational Linguistics.

Nitika Mathur, Johnny Wei, Markus Freitag, Qingsong Ma, and Ondřej Bojar. 2020. Results of the WMT20 metrics shared task. In *Proceedings of the Fifth Conference on Machine Translation*, pages 688–725, Online. Association for Computational Linguistics.

OpenAI. 2023. Gpt-4 technical report.

Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. 2022. Training language models to follow instructions with human feedback.

Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.

Ricardo Rei, Jose G. C. de Souza, Duarte Alves, Chrysoula Zerva, Ana C Farinha, Taisiya Glushkova, Alon Lavie, Luisa Coheur, and Andre F. T. Martins. 2022. Comet-22: Unbabel-ist 2022 submission for the metrics shared task. In *Proceedings of the Seventh Conference on Machine Translation*, pages 578–585, Abu Dhabi. 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.

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.

Giorgos Vernikos, Brian Thompson, Prashant Mathur, and Marcello Federico. 2022. Embarrassingly easy document-level mt metrics: How to convert any pretrained metric into a document-level metric. In *Proceedings of the Seventh Conference on Machine Translation*, pages 118–128, Abu Dhabi. Association for Computational Linguistics.

David Vilar, Markus Freitag, Colin Cherry, Jiaming Luo, Viresh Ratnakar, and George Foster. 2022. Prompting palm for translation: Assessing strategies and performance. *arXiv preprint arXiv:2211.09102*.

A   Appendix: Prompt Templates

Below we provide our prompt templates which we use for the experiments described in this paper. Template **portions in bold face** are used only when a human reference translation is available.

A.1   DA: Direct Assessment
Output scores range from $0 - 100$.

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

{source_lang} source: "{source_seg}"
{target_lang} human reference: {reference_seg}
{target_lang} translation: "{target_seg}"
Score:
```

A.2   SQM: Scalar Quality Metrics
Output scores range from $0 - 100$.

```
Score the following translation from {source_lang} to
{target_lang} with respect to the human reference on a continuous
scale from 0 to 100 that starts with "No meaning preserved", goes
through "Some meaning preserved", then "Most meaning preserved and
few grammar mistakes", up to "Perfect meaning and grammar".

{source_lang} source: "{source_seg}"
{target_lang} human reference: "{reference_seg}"
{target_lang} translation: "{target_seg}"
Score (0-100):
```

A.3   Stars: One to Five Stars Ranking
Output scores range from $1 - 5$. Special care is taken for answers containing non-numerical answers, such as "Three stars", "****", or "1 star".

```
Score the following translation from {source_lang} to {target_lang}
with respect to the human reference with one to five stars.
Where one star means "Nonsense/No meaning preserved",
two stars mean "Some meaning preserved, but not understandable",
three stars mean "Some meaning preserved and understandable",
four stars mean "Most meaning preserved with possibly few grammar mistakes",
and five stars mean "Perfect meaning and grammar".

{source_lang} source: "{source_seg}"
{target_lang} human reference: "{reference_seg}"
{target_lang} translation: "{target_seg}"
Stars:
```

A.4   Classes: Quality Class Labels
Output label one of `"No meaning preserved"`, `"Some meaning preserved, but not understandable"`, `"Some meaning preserved and understandable"`, `"Most meaning preserved, minor issues"`, `"Perfect translation"`.

```
Classify the quality of translation from {source_lang} to {target_lang}
with respect to the human reference into one of following classes: "No meaning
preserved", "Some meaning preserved, but not understandable", "Some meaning
preserved and understandable", "Most meaning preserved, minor issues", "Perfect
translation".

{source_lang} source: "{source_seg}"
{target_lang} human reference: "{reference_seg}"
{target_lang} translation: "{target_seg}"
Class:
```

Templates available from: https://github.com/MicrosoftTranslator/GEMBA/gemba/prompt.py

B Appendix: Full Results

Below table lists all GEMBA results we have obtained for this work. Any missing segment-level scores are due to a subset of segments for which we could not obtain a score even after adding randomness.

Metric	Accuracy	en-de	en-ru	zh-en
GEMBA-GPT4-Classes[noref]	91.2%	0.304	0.390	0.313
GEMBA-GPT4-Stars	91.2%	0.326	0.351	0.382
GEMBA-GPT4-DA	89.8%	0.357	0.358	0.382
GEMBA-Turbo-Stars	89.4%	0.259	0.223	0.265
GEMBA-GPT4-Classes	89.1%	0.222	0.267	0.273
GEMBA-GPT4-Stars[noref]	89.1%	0.308	0.366	0.404
GEMBA-GPT4-SQM[noref]	89.1%	0.359	**0.432**	0.416
GEMBA-GPT4-SQM	88.7%	**0.380**	0.388	0.398
GEMBA-Dav2-Stars	88.3%	0.225	0.282	0.183
GEMBA-Dav3-DA	88.0%	0.306	0.332	0.371
GEMBA-GPT4-DA[noref]	87.6%	0.311	0.405	0.407
GEMBA-Turbo-SQM[noref]	87.6%	0.259	0.309	0.291
GEMBA-Chat-Classes	87.2%	0.220	0.270	0.259
GEMBA-Turbo-SQM	87.2%	0.298	0.277	0.313
GEMBA-Turbo-DA[noref]	86.9%	0.255	0.294	0.264
GEMBA-Turbo-DA	86.5%	0.250	0.234	0.255
GEMBA-Dav3-DA[noref]	86.1%	0.180	0.258	0.289
GEMBA-Dav3-Stars	85.8%	0.294	0.294	0.297
GEMBA-Dav2-SQM	85.8%	0.279	0.325	0.344
GEMBA-Dav2-DA	85.8%	0.231	0.302	0.303
GEMBA-Dav3-Classes	85.4%	0.235	0.289	0.251
GEMBA-Dav3-SQM	85.4%	0.283	0.308	0.346
MetricX XXL	85.0%	0.360	0.420	0.427
GEMBA-Chat-Stars	85.0%	0.292	0.248	0.343
GEMBA-Chat-SQM	85.0%	0.250	0.293	0.310
BLEURT-20	84.7%	0.344	0.359	0.361
GEMBA-Turbo-Stars[noref]	84.3%	0.255	0.279	0.261
COMET-22	83.9%	0.368	0.400	**0.428**
GEMBA-Dav2-DA[noref]	83.9%	0.209	0.285	0.280
COMET-20	83.6%	0.319	0.330	0.332
GEMBA-Chat-Classes[noref]	83.6%	0.193	0.306	0.256
GEMBA-Chat-Stars[noref]	83.6%	0.209	0.323	0.356
GEMBA-Dav3-Stars[noref]	83.2%	0.198	0.310	0.235
UniTE	82.8%	0.369	0.378	0.357
MS-COMET-22	82.8%	0.283	0.351	0.341
GEMBA-Dav2-SQM[noref]	82.8%	0.216	0.306	0.310
GEMBA-Dav3-SQM[noref]	82.5%	0.218	0.328	0.268
GEMBA-Turbo-Classes	82.5%	0.170	0.167	0.178
GEMBA-Chat-DA[noref]	82.1%	0.231	0.332	0.359
MATESE	81.0%	0.323	0.279	0.389
GEMBA-Chat-SQM[noref]	81.0%	0.224	0.320	0.284
GEMBA-Chat-DA	81.0%	0.307	0.328	0.361
GEMBA-Dav2-Classes	79.6%	0.173	0.260	0.184
GEMBA-Dav2-Stars[noref]	79.6%	0.142	0.203	0.193
YiSi-1	79.2%	0.235	0.227	0.296
COMETKiwi[noref]	78.8%	0.290	0.359	0.364
GEMBA-Dav3-Classes[noref]	78.8%	0.176	0.271	0.172
COMET-QE[noref]	78.1%	0.281	0.341	0.365
GEMBA-Dav2-Classes[noref]	78.1%	0.105	0.172	0.128
BERTScore	77.4%	0.232	0.192	0.316
UniTE-src[noref]	75.9%	0.287	0.342	0.343
MS-COMET-QE-22[noref]	75.5%	0.233	0.305	0.287
MATESE-QE[noref]	74.8%	0.244	0.229	0.337
f200spBLEU	74.1%	0.180	0.153	0.140
chrF	73.4%	0.214	0.168	0.147
BLEU	70.8%	0.169	0.140	0.145
GEMBA-Turbo-Classes[noref]	62.0%	-0.010	0.027	0.029
GEMBA-Curie-Classes[noref]	61.7%	0.001	-0.007	-0.053
GEMBA-Bab-Stars[noref]	58.4%	—	—	—
GEMBA-Bab-DA[noref]	55.8%	-0.119	-0.011	—
GEMBA-Curie-Stars[noref]	54.7%	—	—	—
GEMBA-Curie-DA	54.4%	—	—	—
GEMBA-Curie-DA[noref]	51.8%	—	0.054	—
GEMBA-Bab-SQM	51.8%	—	—	—
GEMBA-Bab-SQM[noref]	51.1%	-0.010	0.006	—
GEMBA-Bab-Stars	48.2%	—	—	—
GEMBA-Bab-Classes	47.4%	-0.086	-0.089	-0.066
GEMBA-Curie-Classes	43.4%	-0.011	-0.035	-0.035
GEMBA-Curie-SQM[noref]	41.6%	—	—	—
GEMBA-Curie-SQM	40.5%	0.026	0.013	0.007
GEMBA-Bab-DA	39.1%	-0.110	—	—
GEMBA-Curie-Stars	37.2%	—	—	—
GEMBA-Bab-Classes[noref]	35.0%	-0.010	-0.033	0.067

Table 7: System-level pairwise accuracy and segment-level Kendall Tau scores.

C Appendix: Ada Answers

We present some of the replies from the Ada model on various prompts in Table 8 to illustrate the frequent, invalid answers of this model. Similar behaviour was observed when querying GPT 2. Both models failed with respect to our translation quality assessment experiment.

DA	0 = no meaning preserved 1 = perfect meaning and grammar 2 = excellent meaning 3 = very good meaning 4 = excellent 5 = very good 6 = excellent 7 = excellent 8 = excellent 9 = excellent 10 = perfect
DA	No meaning preserved: 0 Perfect meaning and grammar: 100
DA	0 – It does not exist in my language 1 – It does not exist in my text 2 – It does not exist in my text 3 – It exists in my text 4 – It exists in my text 5 – It exists in my text 6 – It exists in my text 7 – It exists in my text 8 – It exists in my text 9 – It exists in my text 10 – It exists in my text
DA	This sentence does not have a score.
DA[noref]	The source for this translation is English. The translation is 1 out of 100.
DA[noref]	0 = no meaning preserved 1 = perfect meaning and grammar 2 = perfect meaning
SQM	Some meaning preserved Most meaning preserved and few grammar mistakes
SQM[noref]	No meaning preserved Some meaning preserved Most meaning preserved and few grammar mistakes

Table 8: Answers by the Ada model for various prompts. We observe that SQM prompts are closer to expected outputs than answers to the corresponding DA prompts. Similar behaviour was observed when querying GPT 2.