Recovered in Translation: Efficient Pipeline for Automated Translation of Benchmarks and Datasets

Anonymous ACL submission

Abstract

The reliability of multilingual Large Language Model (LLM) evaluation is currently compromised by the inconsistent quality of translated benchmarks. Existing resources often suffer from semantic drift and context loss, which can lead to misleading performance metrics. In this work, we present a fully automated framework designed to address these challenges by enabling scalable, high-quality translation of datasets and benchmarks. We demonstrate that adapting test-time compute scaling strategies, specifically Universal Self-Improvement (USI) and our proposed multi-round ranking method, T-RANK, allows for significantly higher quality outputs compared to traditional pipelines. By effectively applying these methods, our framework ensures that benchmarks preserve their original task structure and linguistic nuances during localization. We apply this approach to translate popular benchmarks and datasets into eight Eastern and Southern European languages. Evaluations using both reference-based metrics and LLM-as-a-judge show that our translations surpass existing resources, resulting in more accurate downstream model assessment. We release both the framework and the improved benchmarks to facilitate robust and reproducible multilingual AI development.

1 Introduction

Large Language Models have demonstrated rapid progress in machine translation, now outperforming classical tools such as Google Translate and DeepL (Deutsch et al., 2025). This advancement benefits multilingual model development by addressing data and benchmark scarcity in non-English settings. Recent work shows that sampling multiple translation candidates at higher temperatures enhances translation quality, with test-time scaling methods like Best-of-N (Stiennon et al., 2020) and Fusion-of-N (Khairi et al., 2025a) demonstrating improved performance across multilingual domains.

A significant gap in recent research is the limited exploration of benchmark translation quality. Multilingual benchmarks are critical for assessing model development efficiency, yet receive less attention than data development. Existing benchmarks remain imperfect due to reliance on older LLMs or oversimplified methods – most were translated using classical MT tools lacking instruction prompting capabilities, or older models like GPT-4 whose multilingual capabilities lag behind current frontier models (Lai et al., 2023).

We investigate available benchmarks and identify common translation flaws stemming from translating questions and answers separately, creating contextual and grammatical mismatches. We focus on Eastern and Southern European languages because: (1) they exhibit complex grammatical features (extensive case systems, grammatical gender, aspect-based verbs) that are sensitive to contextual misalignment; (2) they represent mid-resource languages where frontier models still lack adequate localization; (3) unlike lower-resource languages, translated benchmarks exist for comparison. We translate MMLU, Hellaswag, ARC, and Winogrande into Ukrainian, Romanian, Slovak, Lithuanian, Bulgarian, Turkish, Greek and Estonian.

In this paper, we address three key research questions:

- **RQ1**: Are current translated multilingual benchmarks robust and reliable?
- **RQ2**: To what extent do test-time compute scaling methods improve machine translation quality?
- **RQ3**: How can we translate benchmarks and datasets into other languages while efficiently integrating their unique linguistic features into the evaluated tasks?

To address these challenges, we present a novel automated translation framework supporting four

methods across various model types, including open-weight models. The framework facilitates machine translation of datasets and benchmarks with minimal manual supervision and maximum configurability, offering multiple methods to balance costs and time investment based on language performance in selected MT models. Our framework incorporates: (1) classic one-prompt translation with optional self-correction (SC) as a lightweight solution; (2) Best-of-N with LLM-as-a-judge scoring to provide reward signals; (3) Universal Self-Improvement (USI) with multi-prompt sampling and unification of translation candidates into refined outputs; (4) our newly proposed Translation Ranking (T-RANK), which employs multi-prompt candidate sampling and multi-round competitive ranking to enhance error detection and achieve superior translation quality.

We evaluate these test-time scaling methods across several mid-resource languages, translate popular benchmarks into these languages, and compare our results with existing translations. Our findings demonstrate that the integrated methods produce higher-quality translations with more accurate evaluation results and improved text quality. Moreover, these methods yield performance improvements on established machine translation benchmarks such as WMT24++ and FLORES.

We summarize our key contributions as follows:

- We provide a comprehensive analysis of multilingual benchmark quality and identify the root causes of current translation deficiencies.
- We propose a novel, fully automated framework that enables efficient translation of datasets and benchmarks with maximum flexibility. The pipeline is fully configurable, allowing practitioners to optimize the balance between translation quality and cost in multilingual development.
- We release several benchmarks translated into Eastern and Southern European languages, including Ukrainian, Romanian, Slovak, Lithuanian, Bulgarian, Turkish, Greek and Estonian. We find out, that translation quality influences benchmark evaluation results significantly, and our proposed methods yield more reliable and accurate multilingual evaluation.

These findings enable more efficient multilingual model development by improving both data and benchmark acquisition, thereby facilitating more time- and cost-effective training of locally adapted LLMs.

2 Background and Related Work

2.1 Foundation for LLM-Based Efficient Translation

The rapid expansion of multilingual applications for large language models (LLMs) necessitates scalable and high-quality translation frameworks. Recent advancements propose innovative methodologies to enhance LLM-based translation through adaptive prompting and self-refinement strategies.

The Adaptive Few-shot Prompting (AFSP) framework addresses prompt sensitivity in machine translation by dynamically selecting suitable translation demonstrations. AFSP retrieves semantically similar examples from parallel corpora based on LLM-generated embeddings, then generates and reranks multiple translation candidates to select the most contextually appropriate translation (Tang et al., 2025).

The TEaR (Translate, Estimate, and Refine) framework introduces systematic self-refinement for LLM-based translation. TEaR operates by translating the source text, estimating translation quality, and refining based on identified errors. This iterative process enables LLMs to self-correct and improve translation quality, with cross-model experiments demonstrating that general-purpose LLMs can perform both translation and evaluation simultaneously (Feng et al., 2024).

In addition to these frameworks, several test-time compute scaling methods show potential for enhancing LLM translation capabilities. Although originally designed for general reasoning tasks such as mathematics and coding, these methods demonstrate promising applications in translation:

- **Best-of-N Sampling**: This method generates multiple translation outputs and selects the best one based on predefined criteria. It leverages the diversity in LLM outputs due to temperature variations, increasing the likelihood of obtaining high-quality translations (Stiennon et al., 2020).
- **Universal Self-Consistency (USC)**: USC extends the concept of self-consistency by enabling LLMs to select the most consistent answer among multiple candidates without relying on answer extraction processes. This ap-

proach is particularly effective for open-ended generation tasks, improving performance in mathematical reasoning, code generation, and summarization (Chen et al., 2023).

- **Fusion-of-N**: Rather than selecting the single best answer as in Best-of-N, Fusion-of-N synthesizes the most informative elements from multiple candidates into a single final output. The method uses an LLM judge to aggregate diverse strengths from the candidate pool. Fusion-of-N outperforms Best-of-N and shows promising gains on multilingual tasks, including machine translation (Khairi et al., 2025a).

Findings from Khairi et al. (2025b) also confirm, that sampling multiple candidates with higher tempterature with self-improvement and refined selection process leads to significant improvements in many multilingual domains, including machine translation.

Results from the WMT24++ benchmark (Deutsch et al., 2025) further corroborate the efficacy of these methodologies, demonstrating that state-of-the-art LLMs outperform traditional machine translation tools across various language pairs. The research suggests that integrating test-time compute strategies can significantly enhance translation quality, particularly for low-resource languages. A clear trend emerges whereby newer and larger LLMs consistently outperform existing tools and earlier model iterations, indicating the potential for continued improvements in translation quality as more advanced models are released.

Takeaway 1: Large language models outperform traditional machine translation systems like Google Translate and DeepL through targeted prompts that address domain-specific and language-specific translation challenges.

Collectively, these methodologies underscore the importance of adaptive prompting, self-refinement, and test-time scaling strategies in enhancing LLM-based translation systems. By integrating these approaches, it is possible to develop more robust, efficient, and high-quality translation frameworks capable of addressing the challenges posed by low-resource languages while minimizing the need for manual human intervention.

2.2 Shortcomings of Existing Multilingual Resources

The rapid development of large language models (LLMs) has highlighted the need for robust multilingual benchmarks. However, translating existing benchmarks often introduces issues that compromise accurate assessment of LLM capabilities across languages.

A prominent example is the MuBench benchmark dataset introduced by Han et al. (2025), comprising widely used benchmarks (Hellaswag (Zellers et al., 2019), ARC (Clark et al., 2018), Winogrande (Sakaguchi et al., 2019), MMLU (Hendrycks et al., 2021), and others) translated into 61 languages with 3.9M samples. While this initiative represents an important step toward multilingual evaluation, the translation process relies on an automated pipeline with quality control measures including semantic consistency evaluation, translation purity assessment, and cultural sensitivity checks. However, despite human expert evaluation of 30k samples across 15 languages combined with LLM-as-a-Judge validation, the predominantly automated approach raises concerns regarding translation quality for the remaining samples, as machine translation methods may fail to capture nuanced linguistic and cultural contexts in all 61 languages. Furthermore, variations in quality assurance depth across languages could introduce semantic inaccuracies that skew evaluation results.

Beyond general quality concerns, certain benchmarks present inherent translation challenges. For example, Winogrande, MMLU and Hellaswag may contain answer options with gender-specific adjective endings for sentence completion tasks. In languages with gender-specific adjectives, translating these options can inadvertently reveal correct answers or mislead towards picking an incorrect option, allowing models to succeed through language proficiency rather than reasoning capability. Such linguistic features can compromise evaluation integrity through unintended answer leakage.

The Global-MMLU project (Singh et al., 2024) represents a substantial effort to advance multilingual evaluation by translating the MMLU benchmark into 42 languages. The process combines machine translation with crowdsourced human verification; however, only approximately 20% of machine-translated texts underwent manual correction. Moreover, questions and answers were translated separately, resulting in observable gram-

matical inconsistencies in translated entries for languages such as Ukrainian. This context-agnostic translation approach can produce inconsistencies or grammatical errors that affect benchmark coherence and accuracy, particularly for lower-resource languages. Additionally, relying on tools like Google Translate without thorough human validation may overlook language-specific nuances essential for accurate evaluation.

The Okapi framework (Lai et al., 2023) introduces a novel approach to multilingual instruction tuning by leveraging Reinforcement Learning from Human Feedback (RLHF). Unlike traditional methods relying solely on supervised fine-tuning (SFT), Okapi combines SFT with RLHF to align model outputs more closely with human preferences across diverse languages. Moreover, some of the popular benchmarks like MMLU, Hellaswag and ARC were also translated to 26 languages to enable multilingual evaluation. However, Okapi's translation process utilized the GPT-4 model series and the translation could be further improved by employing test-time compute scaling methods to enhance translation quality. Additionally, the framework does not explicitly address language-specific grammatical features that may impact benchmark coherence and accuracy.

Takeaway 2: Translating questions and answer options within the same prompt context is essential for sentence completion tasks, as it preserves semantic relationships and prevents contextual misleading during evaluation.

Recent studies, including WMT24++ (Deutsch et al., 2025), demonstrate that state-of-the-art LLMs outperform traditional machine translation tools such as DeepL or Google Translate across all evaluated language pairs. This finding suggests that reliance on older translation tools may be insufficient for ensuring high-quality translations in multilingual datasets and benchmarks. Furthermore, practical constraints characteristic of popular frameworks (such as low API rate limits or inability to use advanced prompting techniques) can hinder large-scale translation efforts, making it more difficult to produce comprehensive and accurate datasets.

In summary, while initiatives such as MuBench, Global-MMLU, and Okapi represent important progress toward multilingual LLM evaluation, the methods employed in translating benchmarks exhibit limitations that hinder improvement. Dependence on automatic translation without comprehensive human verification, challenges arising from language-specific grammatical structures, and the limited capabilities of translation tools collectively highlight the need for more sophisticated approaches that respect linguistic features. Furthermore, current frontier language models are already outperfoming classical tools, which were used to translate some of widely used multilingual benchmarks. Future research should advance automated methodologies that reduce reliance on manual curation while simultaneously improving both the quality and accessibility of multilingual evaluation resources.

3 Building an Efficient Automated Translation Framework

In this section we present a novel automated translation framework, which utilizes Large Language Models with features adapted for both dataset and benchmark (QA/test) formats. Moreover, the framework also allows flexibility in methods and their inner configurable parameters to optimize cost- and time-effectiveness.

3.1 Motivation

As highlighted in previous sections, we observed the lack of research and solutions towards scaled automated translation adapted for custom benchmark formats without the loss of the translation quality. For this framework, we aim to tackle the following key problems in LLM-assisted data translation:

- **Support for Diverse Data Formats**: Many datasets have complex, nested structures that complicate processing. For LLM training, flat string fields without hierarchical nesting ensure easier ingestion and more efficient workflows.
- **Preservation of Benchmark QA Structure**: Benchmark evaluation reliability depends on maintaining coherence between questions and answer choices. Preserving these relationships during translation ensures benchmarks remain faithful to the original task structure.
- **Adapting to Varying Language Resource Availability**: High-resource languages (German, Spanish, Hindi) achieve strong translation quality with simple zero-shot prompt-

ing, while mid- and low-resource languages require sophisticated, language-specific pipelines. Users need flexibility to configure translation approaches based on each language's characteristics.

- **Addressing Language-Specific Phenomena**: EEU languages possess unique grammatical features requiring careful handling. Language-specific prompt engineering with few-shot examples and LLM-powered verification stages can improve translation fidelity by identifying and correcting language-specific errors.

3.2 Methodology

In our final framework, we propose two configuration modes – Dataset and Benchmark. They use different prompts and data format handling, since dataset is more straightforward, but benchmark has a complex connection between question and answers, which needs to be preserved. Finally, we propose four translation methods:

- **SC (Self-Check)**: 0-shot simple translation with optional additional check from another LLM (different chat).
- **Best-of-N sampling**: sampling N translation candidates and prompting the model to score them, then picking the one with the highest score.
- **USI (Universal Self-Improvement)**: sampling N translations, asking the model to combine them into the best one according to predefined evaluation criteria.
- **T-RANK (Translation Ranking)**: sampling N translations, asking the model to rank them according to their quality and criteria eligibility, then correcting the best candidate.

We now examine the translation methods employed in our framework, highlighting their advantages and optimal use cases.

Default translation with optional self-check (SC). This method involves a simple 0-shot prompting of LLM to translate a text. The user has an option to include a self-check stage: after translation, the model (in a new chat with no history) is prompted to evaluate and correct the result with respect to the original text content. This method is suitable for large text translation into high-resource languages, as it is less costly to perform, due to the sufficient translation capabilities in high-resource languages; however, since the model is prompted to look out for potential errors it might hallucinate them. Additionally, there exists an option to include a few-shot prompt to provide an example of which language-specific points the model should consider during translation.

Best-of-N sampling (BoN). Drawing from test-time compute scaling methods, we implement Best-of-N sampling without a reward model to maintain a training-free framework. We sample N translations at higher temperature (0.7) for diversity, then prompt the LLM to score candidates 1-10 based on specified criteria (Appendix A.6), selecting the highest-scored translation. While cost-effective and language-agnostic, this method yields lower quality than T-RANK and USI, as LLMs exhibit limitations in numerical scoring and positional bias, favoring earlier candidates despite identifying obvious errors effectively.

Universal Self-Improvement (USI). Building on Universal Self-Consistency and Fusion-of-N, this method operates on the principle that the most consistent translation is not necessarily the best. While Fusion-of-N improves translation quality substantially, the absence of precise translation-specific metrics limits the model's ability to efficiently identify each candidate's strengths for optimal merging, potentially resulting in fusion outputs that incorporate errors. Therefore, we adapt this approach to cultivate self-improvement specifically for machine translation. We sample N candidate translations using higher temperature (0.7 recommended), then present them to an evaluator LLM with instructions to combine the candidates into the best version according to specified criteria. This method proves cost-efficient and time-efficient, requiring only $N + 1$ model calls per entry, while successfully addressing and correcting language-specific features, making it particularly suitable for low-resource languages. We illustrate the Universal Self-Improvement workflow in more detail in Figure 1.

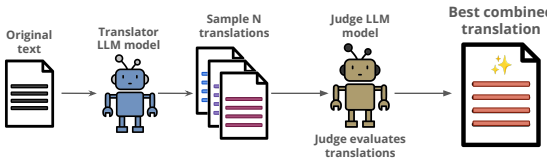

Figure 1: Universal Self-Improvement (USI) method workflow outlook

Translation Ranking (T-RANK). Building upon prior work in translation ranking, we propose an adaptive approach that refines this concept. This method begins by sampling N diverse translations with high temperature. A judge model then evaluates these candidates by ranking them according to predetermined quality criteria, with the top-ranked translation selected as the final output, optionally with refinements. Using this method, we observe more sophisticated reasoning from non-reasoning models, as they successfully identify subtle errors that other methods fail to correct. Moreover, evaluation through comparative ranking rather than numerical scoring appears to facilitate more detailed and rigorous assessment of translation quality.

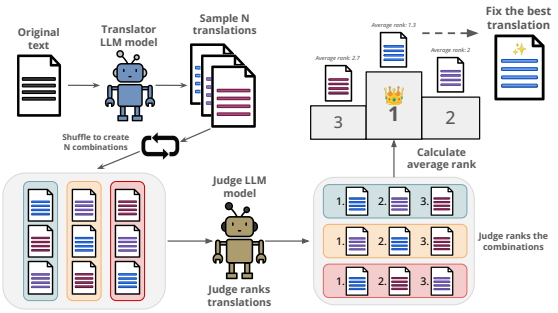

Figure 2: Translation Ranking (T-RANK) method workflow outlook

Previous research on LLM-based evaluation has identified several risks and biases inherent to such methodologies. One notable issue is positional bias: LLMs tend to assign higher scores to candidates presented earlier in the sequence. Furthermore, LLM judges often focus their evaluation on the first candidate, identifying specific flaws, and subsequently bias their assessment of following candidates by primarily searching for similar errors, potentially overlooking novel issues in later candidates.

To reduce these prejudices, we introduce a multiple-round sequential ranking strategy, in which the candidate translations are systematically moved across different positional orders. In particular, each translation candidate is presented once in every possible position across all rounds. This process can be visualized as the construction of a sequential grid, where with each new round, candidates are moved one position further till the end (as visualized in Figure 2). After N rounds for N candidates, each candidate has been presented in each position exactly once. Lastly, we show all the candidates to the judge model again and ask to correct and refine the selected translation candidate if needed - this improves model's ability to notice potential translation flaws by looking at strong and weak sides of other variants. We noticed an improved behavior and reasoning when using this method and in particular, showing all candidates when doing the final correction as shown in Appendix A.1. This approach shows a promising improvement in reliability of the evaluation by reducing positional biases while controlling evaluation costs and computational efficiency with a total of $2N + 1$ model calls. Optionally, once can combine best candidate correction during ranking, however, we found evaluator model's reasoning to be more attentive in a separated setup.

We provide example prompts for all proposed methods in Appendix A.6.

Takeaway 3: While Best-of-N and USI sampling improve translation quality, the T-RANK method's competitive ranking approach proves more efficient at identifying subtle translation errors.

3.3 Benchmark Translation for Eastern European Languages

We have translated several popular benchmarks into Eastern European languages using our proposed framework. The selected benchmarks include MMLU, Hellaswag, ARC, and Winogrande, which are widely used for evaluating LLM capabilities across various tasks. We focused on Eastern and Southern European languages such as Ukrainian, Romanian, Slovak, Lithuanian, Greek, Bulgarian, Turkish and Estonian due to their complex grammatical structures and mid-resource status. These languages present unique challenges for machine translation, making them ideal candidates for testing the effectiveness of our framework. MMLU was translated using GPT-4o-mini-2024-07-18 model checkpoint from OpenAI and all mentioned benchmarks were translated to Ukrainian using the same model. Remaining benchmark translations to other languages were performed using Gemini-2.0-Flash model from Google. In the next section and in Appendix A.3, we present a comprehensive evaluation of the translation quality achieved through our framework, comparing it with existing translations and assessing its impact on benchmark performance.

4 Evaluation Results

4.1 Machine Translation Benchmarks

We utilize two datasets to evaluate the translation quality of our proposed methods: FLORES and WMT24++.

FLORES. The FLORES benchmark evaluates machine translation systems across multilingual scenarios. FLORES-101 provides 3,001 (1,012 devtest split) sentences from English Wikipedia professionally translated into 101 languages, enabling evaluation of many-to-many translation systems, particularly for low-resource language pairs (Guzmán et al., 2019; Goyal et al., 2022).

WMT24++. The WMT24++ benchmark covers 55 languages and dialects with human-written references and post-edits across four domains: literary, news, social, and speech. This diversity enables comprehensive evaluation of translation systems across high-resource and low-resource languages (Deutsch et al., 2025).

Both FLORES and WMT24++ serve as standard benchmarks in the machine translation community due to their broad language coverage and high-quality human translations, enabling meaningful comparisons across models and approaches.

We evaluate our proposed methods on English-Ukrainian translation using the COMET (Crosslingual Optimized Metric for Evaluation of Translation) metric. COMET leverages multilingual pretrained models to assess translations by comparing source text, hypothesis, and reference, demonstrating higher correlation with human judgments than traditional metrics like BLEU or chrF++ (Rei et al., 2020). We report COMET system-level scores (aggregated across all translations) in Table 1, using the Unbabel/XCOMET-XL model for reference-based quality estimation (Guerreiro et al., 2023). For mentioned tasks we have used GPT-4o-mini-2024-07-18 model for translation.

From the presented results, we observe that USI and T-RANK demonstrate clear advantages over other methods; however, it remains unclear whether T-RANK consistently outperforms USI. This raises the question of whether a correlation exists between translation cost or effort and quality. We must account for the fact that COMET evaluations are not entirely reliable, even when reference translations are provided. Notably, both datasets used for evaluation (WMT24++ and FLORES) primarily contain short texts that do not adequately test complex grammatical structures, particularly for Ukrainian.

Method	WMT24++	FLORES
Baseline	0.827	0.937
SC (with check)	0.821	0.937
Best-of-N (n=5)	0.843	0.943
USI (n=5)	0.843	**0.945**
T-RANK (p=5)	**0.845**	0.940

Table 1: COMET (reference-based) translation quality scores for WMT24++ and FLORES EN→UK pair (GPT-4o-mini). n denotes number of sampled candidates, p denotes number of different translation prompts (for $p > 1$ we use $n = 1$).

Moreover, the choice of a correct translation often depends on personal preferences and stylistic conventions; a single source text may have multiple valid translation candidates, all grammatically correct. Therefore, reference translations in commonly used machine translation benchmarks cannot be considered absolute gold standards, and achieving the highest COMET scores does not guarantee that a method is error-free. This motivates the need for additional comparative tests under different evaluation paradigms, such as Quality Estimation (QE) or reference-free machine translation evaluation. As discussed in appendix A.5, USI is more suitable for short and simple dataset translation, whereas T-RANK shows better performance when translating benchmarks, especially when they have complex question structure, which might get "lost in translation" for languages explored in this paper.

For reference-free QE, we can directly use COMET to score translations without requiring an ideal reference translation. We compared USI and T-RANK methods using the QE setup on FLORES and MMLU datasets, which contain longer and more complex texts. In this work, we use Unbabel/wmt23-cometkiwi-da-xl model for reference-free QE COMET evaluation Rei et al. (2023) and the detailed evaluations are provided in appendix A.5.

4.2 Multilingual Benchmark Translation Quality

Moreover, we also use LLM-as-a-judge to compare MMLU translations between standard industry-recognized source (Global-MMLU) and our most efficient translation method (T-RANK). Using Gemini-2.5-Flash with high reasoning mode as a judge, we showcase a clear advantage in translation quality using our proposed method, as presented

in Table 2. We observe that our translations are preferred also in other available languages like Romanian and Lithuanian, as shown in Appendix A.4.

Translation	Wins	Draws	Losses
Global-MMLU-UK	2016	3276	8750
T-RANK (ours)	**8750**	3276	2016

Table 2: Comparison of translation quality between Global-MMLU and our proposed T-RANK method for MMLU benchmark in Ukrainian language using LLM-as-a-judge.

We also compare evaluation results of mid-sized models like Gemma 3 4/12B, Qwen 3 8B and Llama 3.1 8B on our improved translations versus existing benchmark translations. As shown in Table 3, we observe higher scores on our translations, providing additional evidence of enhanced translation quality. Below, we provide average differece between evaluating on our translations and existing ones (positive difference means results are higher on our translations); more results for other languages are presented in Appendix A.3.

Benchmark	AVG Δ
ARC-Challenge	+2.35%
Hellaswag	+1.63%
MMLU	+0.94%
Winogrande	+3.42%

Table 3: Average improvement per benchmark across all languages

5 Conclusion

We present a novel automated translation framework designed to enable rapid, high-quality translation of datasets and benchmarks while minimizing human intervention. Our integrated methods demonstrate substantial improvements on WMT24++ and FLORES benchmarks as measured by COMET scores, with T-RANK and USI achieving the strongest performance through test-time compute strategies. The quality of our translated benchmarks is further validated through COMET Quality Estimation scores and LLM-as-a-judge evaluations, which confirm meaningful improvements over existing translations. Notably, evaluation results of mid-sized models such as Gemma 3, Qwen 3 and Llama 3.1 on our improved translations yield higher scores compared to existing benchmark translations, providing additional evidence of enhanced translation quality. These findings demonstrate that our framework effectively balances translation accuracy, computational efficiency, and scalability, offering a practical solution for creating high-quality multilingual evaluation resources. Future work should explore adaptive method selection based on translation difficulty, integration of dedicated quality models, and comprehensive evaluation across open-weight models to further enhance the framework's capabilities on languages beyond Europe.

Limitations

Our study has several limitations that warrant consideration. First, we employ LLM-based scoring rather than dedicated translation quality models like COMET for the Best-of-N selection process, which may affect the reliability of candidate ranking. Second, our approach applies uniform methods across all entries without automatically estimating translation difficulty per input, though adaptive method selection based on text complexity could potentially improve efficiency and quality. Machine translation benchmarks demonstrate that the most advanced and computationally expensive methods do not always yield superior results for shorter text sequences, particularly those not used in question-answering contexts. We defer to users of our framework to select methods based on their specific objectives and resource constraints. Third, we rely primarily on closed-source models for translation, with limited testing of open-weight alternatives. We hypothesize that our proposed methods would yield greater benefits for open-weight models, which typically demonstrate weaker performance in zero-shot translation settings, though this requires validation through comprehensive evaluation. Finally, while we focus on Eastern and Southern European languages, further research is needed to assess the generalizability of our framework across a broader range of low-resource languages with diverse linguistic characteristics.

Ethics Statement

This work aims to improve multilingual benchmark quality to enable more equitable evaluation of language models across diverse languages. We acknowledge several ethical considerations. First, our reliance on closed-source models raises repro-

ducibility concerns; we plan to extend evaluation to open-weight alternatives to promote accessibility. Second, while focusing on Eastern European languages addresses an important gap, many lower-resource languages remain underrepresented, and we encourage extending these methods to broader linguistic contexts. Third, we recognize that improved benchmarks may still be imperfect or subject to translation model bias; however, accurate multilingual evaluation is crucial for developing culturally aware and safer AI systems. Our automated translation approach avoids potential labor exploitation concerns associated with human translation.

The code and translated benchmarks produced in this work will be released under permissible open licenses to enable community validation and reproducibility. Generative AI systems were employed exclusively for language assistance, including paraphrasing, spell-checking, and stylistic refinement of the authors' original content, as well as for generating boilerplate code. All core research contributions, experimental design, and findings represent the original work of the authors. Used benchmarks and datasets are publicly available under premissive licensing; any ethical considerations related to their use are discussed in the original publications.

References

Xinyun Chen, Renat Aksitov, Uri Alon, Jie Ren, Kefan Xiao, Pengcheng Yin, Sushant Prakash, Charles Sutton, Xuezhi Wang, and Denny Zhou. 2023. Universal self-consistency for large language model generation.

Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. 2018. Think you have solved question answering? try arc, the ai2 reasoning challenge. *Preprint*, arXiv:1803.05457.

Daniel Deutsch, Eleftheria Briakou, Isaac Caswell, Mara Finkelstein, Rebecca Galor, Juraj Juraska, Geza Kovacs, Alison Lui, Ricardo Rei, Jason Riesa, Shruti Rijhwani, Parker Riley, Elizabeth Salesky, Firas Trabelsi, Stephanie Winkler, Biao Zhang, and Markus Freitag. 2025. Wmt24++: Expanding the language coverage of wmt24 to 55 languages & dialects. *Preprint*, arXiv:2502.12404.

Zhaopeng Feng, Yan Zhang, Hao Li, Bei Wu, Jiayu Liao, Wenqiang Liu, Jun Lang, Yang Feng, Jian Wu, and Zuozhu Liu. 2024. Tear: Improving llm-based machine translation with systematic self-refinement.

Naman Goyal, Cynthia Gao, Vishrav Chaudhary, Peng-Jen Chen, Guillaume Wenzek, Da Ju, Sanjana Krishnan, Marc'Aurelio Ranzato, Francisco Guzmán, and Angela Fan. 2022. The flores-101 evaluation benchmark for low-resource and multilingual machine translation. *Transactions of the Association for Computational Linguistics*, 10:522–538.

Nuno M. Guerreiro, Ricardo Rei, Daan van Stigt, Luisa Coheur, Pierre Colombo, and André F. T. Martins. 2023. xcomet: Transparent machine translation evaluation through fine-grained error detection. *Preprint*, arXiv:2310.10482.

Francisco Guzmán, Peng-Jen Chen, Myle Ott, Juan Pino, Guillaume Lample, Philipp Koehn, Vishrav Chaudhary, and Marc'Aurelio Ranzato. 2019. The flores evaluation datasets for low-resource machine translation: Nepali-english and sinhala-english. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 6098–6111. Association for Computational Linguistics.

Wenhan Han, Yifan Zhang, Zhixun Chen, Binbin Liu, Haobin Lin, Bingni Zhang, Taifeng Wang, Mykola Pechenizkiy, Meng Fang, and Yin Zheng. 2025. Mubench: Assessment of multilingual capabilities of large language models across 61 languages. *Preprint*, arXiv:2506.19468.

Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. Measuring massive multitask language understanding. *Preprint*, arXiv:2009.03300.

Ammar Khairi, Daniel D'souza, Marzieh Fadaee, and Julia Kreutzer. 2025a. Making, not taking, the best of n. *Preprint*, arXiv:2510.00931.

Ammar Khairi, Daniel D'souza, Ye Shen, Julia Kreutzer, and Sara Hooker. 2025b. When life gives you samples: The benefits of scaling up inference compute for multilingual llms. *Preprint*, arXiv:2506.20544.

Viet Dac Lai, Chien Van Nguyen, Nghia Trung Ngo, Thuat Nguyen, Franck Dernoncourt, Ryan A. Rossi, and Thien Huu Nguyen. 2023. Okapi: Instruction-tuned large language models in multiple languages with reinforcement learning from human feedback. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 318–327, Singapore. Association for Computational Linguistics.

Ricardo Rei, Nuno M. Guerreiro, José Pombal, Daan van Stigt, Marcos Treviso, Luisa Coheur, José G. C. de Souza, and André F. T. Martins. 2023. Scaling up cometkiwi: Unbabel-ist 2023 submission for the quality estimation shared task. *Preprint*, arXiv:2309.11925.

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.

Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2019. Winogrande: An adversarial winograd schema challenge at scale. *Preprint*, arXiv:1907.10641.

Shivalika Singh, Angelika Romanou, Clémentine Fourrier, David I Adelani, Jian Gang Ngui, Daniel Vila-Suero, Peerat Limkonchotiwat, Kelly Marchisio, Wei Qi Leong, Yosephine Susanto, and 1 others. 2024. Global mmlu: Understanding and addressing cultural and linguistic biases in multilingual evaluation.

Nisan Stiennon, Long Ouyang, Jeffrey Wu, Daniel Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul F Christiano. 2020. Learning to summarize with human feedback. *Advances in Neural Information Processing Systems*, 33:3008–3021.

Lei Tang, Jinghui Qin, Wenxuan Ye, Hao Tan, and Zhijing Yang. 2025. Adaptive few-shot prompting for machine translation with pre-trained language models.

Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. Hellaswag: Can a machine really finish your sentence? *Preprint*, arXiv:1905.07830.

A Appendix

A.1 Translation Examples

In this subsection we will provide some translation examples showcasing common errors and how our proposed methods address them. Figure 3 demonstrates how T-RANK identifies and corrects translation errors through competitive ranking. Unlike all other methods, we observe improved ability of the judge model to notice translation flaws when evaluating the translation in competitive manner. Figure 4 illustrates a common issue where translating questions and answers separately can leak the correct answer through grammatical gender markers. Figure 5 shows the correct approach where context is preserved during translation.

A.2 Benchmark Translation Statistics

In this subsection we provide statistics of selected benchmark translation, used splits and total number of samples used. We also provide information on models and methods used for final translation.

A.3 Detailed Evaluation Results

In this section we provide additional evaluation results on our own and existing translated benchmarks for analyzed languages. Tables 6 through 14 present performance comparisons across Ukrainian, Romanian, Slovak, Lithuanian, Bulgarian, Turkish, Greek and Estonian languages respectively, showing consistent improvements of our translations over existing benchmarks. Where available, we compare against Okapi/MuBench/Global-MMLU translations as baselines. For Bulgarian language we also compare against our own older translation done with the help of professional translators. With these results we confirm, that translation quality influences the evaluation results across various models with some exceptions to Winogrande. We note that professionally translated Winogrande to Bulgarian shows better results than our automated translation, which indicates that for some languages and benchmarks human intervention is still required to achieve the best quality, though we believe that with our methods the need for such intervention is significantly reduced.

A.4 MMLU LLM-as-a-Judge Evaluation Results

In this section we provide additional LLM-as-a-judge comparisons for available languages. Table 15 shows comparison for Romanian and Table 16 presents results for Lithuanian.

A.5 Machine Translation Benchmark Evaluation Results

In this section, we test GPT-4o-mini and Gemini-2.0-Flash model on WMT and FLORES translation benchmarks, using both reference-based and reference-free COMET evaluation setups. As mentioned before, for reference-based evaluation we use Unbabel/XCOMET-XL model and for reference-free QE evaluation we use Unbabel/wmt23-cometkiwi-da-xl model. Tables 17 through 20 present GPT-4o-mini results, while Tables 21 and 22 show Gemini-2.0-Flash performance.

We see, that on both datasets and evaluation types, USI method gives the best results. We also see, that using sampling single candidate from various prompts in both English and target language (2 English and 2 target languages prompts) gives additional boost in quality over single-prompt generation.

In this case, for Gemini-2.0-Flash we observe that T-RANK multi-prompt method can outperform USI on most language pairs, although the difference is quite small. We believe that both methods have their strong and weak points – while for datasets with shorter texts we recommends using USI, for benchmark translation we observed im-

Candidate 1

Question:

Ви спостерігаєте вид, який народжує лише одне потомство за раз і має відносно тривалий життєвий цикл для свого розміру тіла. Що з наступного, ймовірно, також вірно для цього організму?

Answers:

1. Він живе в новозаселеному середовищі.
2. Це водний організм.
3. Він потребує відносно високої батьківської турботи про потомство.
4. Вік, в якому потомство може самостійно народжувати, відносно молодий.

Candidate 2

Question:

Ви спостерігаєте вид, який народжує лише одне потомство за раз і має відносно тривалий термін життя для свого розміру. Що з наведеного, ймовірно, також вірно для цього організму?

Answers:

1. Він живе в новозаселеному середовищі.
2. Це водний організм.
3. Він потребує відносно великої батьківської турботи про потомство.
4. Вік, в якому потомство саме може народжувати, є відносно молодим.

Candidate 3

Question:

Ви спостерігаєте вид, який народжує лише одне потомство за раз і має відносно тривалий термін життя для свого розміру. Що з наведеного, ймовірно, також вірно для цього організму?

Answers:

1. Він живе в нещодавно колонізованому середовищі.
2. Це водний організм.
3. Він потребує відносно великої батьківської опіки за потомством.
4. Вік, в якому потомство може самостійно народжувати, відносно молодий.

Selected Translation

Question:

Ви спостерігаєте вид, який народжує лише одне потомство за раз і має відносно тривалий термін життя для свого розміру. Що з наведеного, ймовірно, також вірно для цього організму?

Answers:

1. Він живе в нещодавно колонізованому середовищі.
2. Це водний організм.
3. Він потребує відносно великої батьківської опіки за потомством.
4. Вік, в якому потомство може самостійно народжувати, відносно молодий.

Reasoning

The final selected translation is generally accurate and conveys the original meaning well. However, there are minor improvements needed in terms of word choice and grammatical structure. The phrase 'відносно тривалий термін життя' could be simplified to 'відносно довгий термін життя' for better clarity. Additionally, the phrase 'батьківської опіки за потомством' can be improved to 'батьківської опіки над потомством' for grammatical correctness. Overall, the translation is good but requires these minor adjustments for optimal clarity and correctness.

Final Corrected Translation

Question:

Ви спостерігаєте вид, який народжує лише одне потомство за раз і має відносно довгий термін життя для свого розміру. Що з наведеного, ймовірно, також вірно для цього організму?

Answers:

1. Він живе в нещодавно колонізованому середовищі.
2. Це водний організм.
3. Він потребує відносно великої батьківської опіки над потомством.
4. Вік, в якому потомство може самостійно народжувати, відносно молодий.

Figure 3: Translation Ranking (T-RANK) Correction Example

proved quality from using T-RANK method specifically, as shown in first subsection of appendix. Moreover, for Gemini-2.0-Flash model the ranking capabilities seem to be stronger, which may also contribute to better performance of T-RANK method. This supports our idea, that different methods may be more suitable for different models and use-cases.

A.6 Translation Prompts

Table 23 provides the complete prompt templates used for all translation methods described in this

```
Вони хвилювалися, що вино зіпсує ліжко та ковдру, але _ не було зіпсовано.
What does the blank _ refer to?
Option A: ковдра
Option B: ліжко
Answer with A or B.
Answer: ліжко
```

Figure 4: Winogrande Answer Leakage Example

```
Вони хвилювалися, що вино зіпсує ліжко та ковдру, але _ не бу(-в/-ла/-ло/-ли) зіпсовано.
What does the blank _ refer to?
Option A: ковдра
Option B: ліжко
Answer with A or B.
Answer: ліжко ?
```

Figure 5: Winogrande Correct Translation Example

Benchmark	**Train**	**Val**	**Dev**	**Test**	**All**
MMLU	–	1531	285	14042	15858
Winogrande	–	1267	–	–	–
ARC-Challenge	1119	299	–	1172	2,291
Hellaswag	–	–	–	10042	10042
Total	1119	3097	285	25556	**29757**

Table 4: Number of examples in each benchmark split (– indicates split not available or not used).

paper.

Lang/Bench	UKR	SK	RO	LT	EST	BG	GR	TR
MMLU	4o-mini(T-RANK)	4o-mini(T-RANK)	4o-mini(T-RANK)	4o-mini(USI)	4o-mini(USI)	4o-mini(T-RANK)	4o-mini(USI)	4o-mini(USI)
Hellaswag	4o-mini	4o-mini	4o-mini	gemini-2.0-fl	gemini-2.0-fl	gemini-2.0-fl	gemini-2.0-fl	gemini-2.0-fl
ARC	4o-mini	gemini-2.0-fl	gemini-2.0-fl	gemini-2.0-fl	gemini-2.0-fl	gemini-2.0-fl	gemini-2.0-fl	gemini-2.0-fl
Winogrande	4o-mini	gemini-2.0-fl	gemini-2.0-fl	gemini-2.0-fl	gemini-2.0-fl	gemini-2.0-fl	gemini-2.0-fl	gemini-2.0-fl

Table 5: Selected model and method for final translation of benchmarks, if not specified, USI method was used. 4o-mini = GPT-4o-mini-2024-07-18, gemini-2.0-fl = Gemini-2.0-Flash.

	Hellaswag			Winogrande			ARC-Ch.			MMLU		
Model	O	Ot	Δ	O	Ot	Δ	O	Ot	Δ	O	Ot	Δ
Gemma-3-12B-IT	0.687	0.625	0.062	0.580	0.619	-0.039	0.517	0.470	0.048	0.614	0.603	0.012
Llama-3.1-8B	0.570	0.540	0.030	0.549	0.495	0.055	0.431	0.416	0.015	0.497	0.489	0.009
Gemma-3-4B-IT	0.578	0.528	0.050	0.540	0.501	0.039	0.453	0.417	0.037	0.453	0.444	0.010
Qwen3-8B-IT	0.540	0.520	0.020	0.569	0.546	0.023	0.448	0.404	0.044	0.619	0.599	0.020

Table 6: Ukrainian. O=Ours, Ot=Other (Okapi/MuBench/Global-MMLU)

	Hellaswag			Winogrande			ARC-Ch.			MMLU		
Model	O	Ot	Δ	O	Ot	Δ	O	Ot	Δ	O	Ot	Δ
Gemma-3-12B-IT	0.678	0.681	-0.002	0.661	0.589	0.072	0.486	0.454	0.032	0.628	0.614	0.014
Gemma-3-4B-IT	0.567	0.567	0.000	0.575	0.531	0.045	0.412	0.394	0.018	0.479	0.473	0.006
Llama-3.1-8B-IT	0.576	0.581	-0.005	0.616	0.521	0.095	0.358	0.357	0.002	0.529	0.519	0.009
Qwen3-8B-IT	0.538	0.537	0.001	0.588	0.566	0.023	0.368	0.359	0.008	0.650	0.632	0.018

Table 7: Romanian. O=Ours, Ot=Other (Okapi/MuBench/Global-MMLU)

	Winogrande			ARC-Challenge		
Model	Ours	Other	Δ	Ours	Other	Δ
Gemma-3-12B-IT	0.610	0.589	0.021	0.539	0.509	0.030
Gemma-3-4B-IT	0.554	0.517	0.037	0.448	0.444	0.004
Llama-3.1-8B-IT	0.547	0.515	0.032	0.443	0.403	0.040
Qwen3-8B-IT	0.555	0.559	-0.004	0.428	0.427	0.001

Table 8: Slovak. Other=MuBench/Okapi

	Winogrande			MMLU		
Model	Ours	Other	Δ	Ours	Other	Δ
Gemma-3-12B-IT	0.631	0.557	0.073	0.585	0.573	0.013
Gemma-3-4B-IT	0.530	0.506	0.023	0.433	0.410	0.023
Llama-3.1-8B-IT	0.527	0.498	0.029	0.431	0.417	0.014
Qwen3-8B-IT	0.551	0.530	0.021	0.553	0.541	0.012

Table 9: Lithuanian. Other=MuBench/Global-MMLU

	Winogrande		
Model	Ours	Other	Δ
Gemma-3-12B-IT	0.615	0.566	0.049
Gemma-3-4B-IT	0.541	0.516	0.025
Llama-3.1-8B-IT	0.543	0.509	0.034
Qwen3-8B-IT	0.517	0.503	0.014

Table 10: Estonian. Other=MuBench

	Winogrande			MMLU		
Model	Ours	Other	Δ	Ours	Other	Δ
Gemma-3-12B-IT	0.635	0.571	+0.065	0.587	0.581	+0.007
Gemma-3-4B-IT	0.583	0.514	+0.070	0.447	0.437	+0.010
Llama-3.1-8B-IT	0.573	0.518	+0.055	0.494	0.480	+0.014
Qwen3-8B-IT	0.545	0.571	-0.027	0.588	0.569	+0.018

Table 11: Turkish. Other=MuBench/Global-MMLU

	Winogrande			MMLU		
Model	Ours	Other	Δ	Ours	Other	Δ
Gemma-3-12B-IT	0.657	0.601	+0.056	0.593	0.580	+0.013
Gemma-3-4B-IT	0.598	0.518	+0.080	0.428	0.421	+0.007
Llama-3.1-8B-IT	0.593	0.520	+0.073	0.465	0.446	+0.020
Qwen3-8B-IT	0.581	0.528	+0.053	0.563	0.553	+0.011

Table 12: Greek. Other=MuBench/Global-MMLU

	Hellaswag			Winogrande			ARC-Challenge			MMLU		
Model	Ours	Other	Δ	Ours	Other	Δ	Ours	Other	Δ	Ours	Other	Δ
Gemma-3-12B-IT	0.708	0.690	+0.018	0.643	0.677	-0.035	0.526	0.495	+0.031	0.605	0.619	-0.013
Gemma-3-4B-IT	0.590	0.569	+0.022	0.588	0.595	-0.007	0.430	0.410	+0.020	0.454	0.469	-0.015
Llama-3.1-8B-IT	0.539	0.528	+0.011	0.575	0.605	-0.031	0.394	0.367	+0.027	0.481	0.487	-0.006
Qwen3-8B-IT	0.536	0.518	+0.018	0.602	0.629	-0.027	0.414	0.394	+0.020	0.619	0.617	+0.002

Table 13: Bulgarian. Other=INSAIT

	Winogrande		
Model	Ours	Other	Δ
Gemma-3-12B-IT	0.643	0.597	+0.046
Gemma-3-4B-IT	0.588	0.506	+0.082
Llama-3.1-8B-IT	0.575	0.505	+0.069
Qwen3-8B-IT	0.602	0.559	+0.043

Table 14: Bulgarian. Other=MuBench

Translation	Wins	Draws	Losses
Global-MMLU-RO	2646	3020	8376
T-RANK (ours)	**8376**	3020	2646

Table 15: Comparison of translation quality between Global-MMLU and our proposed T-RANK method for MMLU benchmark in Romanian language using Gemini-2.5-Flash as a judge.

Translation	Wins	Draws	Losses
Global-MMLU-LT	3478	3181	7382
USI (ours)	**7382**	3181	3478

Table 16: Comparison of translation quality between Global-MMLU and our proposed USI method for MMLU benchmark in Lithuanian language using Gemini-2.5-Flash as a judge.

Method	EN→UK	EN→SK	EN→RO	EN→LT	EN→ET	EN→BG	EN→TR	EN→EL
Baseline	0.726	0.741	0.882	0.741	0.788	0.751	0.718	0.802
SC	0.708	0.725	0.869	0.735	0.788	0.749	0.715	0.793
BoN n=5	0.743	0.756	0.895	0.767	0.809	0.788	0.736	0.819
USI n=5	0.750	0.759	0.899	0.766	0.806	0.791	0.733	0.817
T-RANK n=5	0.739	0.745	0.885	0.749	0.799	0.776	0.729	0.811
USI multi-prompt p=4	**0.755**	**0.764**	**0.898**	**0.771**	**0.809**	–	–	–
T-RANK multi-prompt p=4	0.742	0.756	0.883	0.753	0.797	–	–	–

Table 17: WMT QE (Quality Estimation) reference-free COMET scores for GPT-4o-mini

Method	EN→UK	EN→SK	EN→RO	EN→LT	EN→ET	EN→BG	EN→TR	EN→EL
Baseline	0.827	0.822	0.873	0.788	0.821	0.834	0.776	0.820
SC	0.821	0.817	0.869	0.790	0.822	0.000	0.834	0.821
BoN n=5	0.844	0.837	0.884	0.805	0.838	0.852	0.791	0.836
USI n=5	0.843	0.839	0.888	0.807	0.842	0.856	0.791	0.838
T-RANK n=5	0.840	0.834	0.885	0.803	0.836	0.855	0.793	0.837
USI multi-prompt p=4	**0.849**	**0.847**	**0.891**	**0.817**	**0.848**	–	–	–
T-RANK multi-prompt p=4	0.845	0.841	0.887	0.812	0.843	–	–	–

Table 18: WMT COMET reference-based scores for GPT-4o-mini

Method	EN→UK	EN→SK	EN→RO	EN→LT	EN→ET
Baseline	0.904	0.906	0.947	0.906	0.927
SC	0.902	0.908	0.950	0.911	0.935
BoN n=5	0.910	0.917	0.956	0.921	0.947
USI n=5	0.913	0.915	0.956	0.924	0.943
T-RANK n=5	0.899	0.906	0.952	0.909	0.932
USI multi-prompt p=4	**0.919**	**0.921**	**0.961**	**0.932**	**0.948**
T-RANK multi-prompt p=4	0.905	0.912	0.954	0.912	0.935

Table 19: FLORES QE (Quality Estimation) reference-free COMET scores for GPT-4o-mini

Method	EN→UK	EN→SK	EN→RO	EN→LT	EN→ET
Baseline	0.937	0.938	0.939	0.906	0.894
SC	0.937	0.937	0.936	0.911	0.902
BoN n=5	0.943	0.945	0.945	0.917	0.918
USI n=5	0.945	0.942	0.947	0.922	0.919
T-RANK n=5	0.938	0.937	0.943	0.910	0.903
USI multi-prompt p=4	**0.947**	**0.946**	**0.949**	**0.928**	**0.918**
T-RANK multi-prompt p=4	0.940	0.943	0.944	0.915	0.907

Table 20: FLORES COMET reference-based scores for GPT-4o-mini

Method	EN→UK	EN→SK	EN→RO	EN→LT	EN→ET	EN→BG	EN→TR	EN→EL
Baseline	0.688	0.703	0.777	0.704	0.737	0.708	0.683	0.719
SC	0.687	0.706	0.777	0.706	0.736	0.709	0.685	0.722
BoN n=5	0.687	0.707	0.780	0.710	0.743	0.713	0.689	0.723
USI multi-prompt p=4	**0.698**	**0.716**	**0.785**	**0.718**	**0.753**	**0.721**	**0.697**	**0.731**
T-RANK multi-prompt p=4	0.697	0.713	0.781	0.716	0.749	0.718	0.695	0.726

Table 21: WMT QE (Quality Estimation) reference-free COMET scores for Gemini-2.0-flash

Method	EN→UK	EN→SK	EN→RO	EN→LT	EN→ET	EN→BG	EN→TR	EN→EL
Baseline	0.828	0.829	0.867	0.805	0.847	0.840	0.778	0.821
SC	0.826	0.834	0.868	0.808	0.845	0.841	0.779	0.824
BoN n=5	0.828	0.836	0.870	0.814	0.852	0.843	0.785	0.823
USI multi-prompt p=4	**0.841**	0.843	0.881	0.826	0.862	0.856	0.797	0.836
T-RANK multi-prompt p=4	0.841	**0.849**	**0.882**	**0.827**	0.861	**0.859**	**0.800**	**0.838**

Table 22: WMT COMET reference-based scores for Gemini-2.0-flash

Method	Prompt Template
Base Translation	**Instructions:** Imagine you're part of a team at an international education center that's revamping its exams for a global audience. Your job is to translate an English question and its answer options into <target_language> so that students from <target_language> schools can be evaluated too. Just provide the final translation—leave out any extra comments or explanations. Use language which is authentic for <target_language> natives. Remember to keep the answer options connected to the question, using the same format as the original (a list for multiple choices or plain text for a single answer). Please do not translate valid code in any of the programming languages. **Original text:** {"Original_question": "<original_question>", "Original_answers": "<original_answers>"} **Output instructions:** Now, please give your final translation in <target_language> exactly in this format, with only the translated content: {"Question": "your_translated_question", "Answers": "translated_answers"}
USI Judge	My task is to translate BENCHMARK questions with answers from English to <target_language>. Your task is to evaluate if the response preserves the original question idea and to verify the correctness of declension and conjunction of words in the target language. The original text in English is: Question: <original_question>, Answers: <original_answers> I have generated the following responses: <responses> Combine the best features from responses to form the best response from grammatical and coherent points of view. Look for: (1) Quality of translation, including grammatical correctness; (2) Domain knowledge - were the terms correctly translated? Were coding terms preserved?; (3) Is the question text fully translated?; (4) Are all answer options fully translated?; (5) Are the words written correctly? Output only the selected response: Question: selected question, Answer: selected answers
T-RANK Ranking	My task is to translate BENCHMARK questions from English to <target_language>. Your task is to rank my translations and select the best one. **Ranking criteria:** (1) Quality of translation; (2) Domain knowledge; (3) Is the question fully translated?; (4) Are all answer options translated?; (5) Correct spelling?; (6) Correct declension and conjunction? **Original:** {"original_question": "<original_question>", "original_answers": <original_answers>} **Candidates:** <responses> **Instruction:** Select the best response (1st place = best). Correct if needed before output. **Output:** Reasoning, then: {"summary": "...", "final_ranks": {...}, "rankings_list": [...], "best_translation": {...}}
Best-of-N Scoring	My task is to translate questions from English to <target_language>. Score my translations 1-10 (10 = best). Original: Question: <original_question>, Answers: <original_answers> Scoring metrics: (1) Translation quality; (2) Question fully translated?; (3) All answers translated?; (4) Question idea preserved?; (5) Correct grammar? Responses: <responses> Output scores only: Response 1: score, Response 2: score, ..., Answers: [list of scores]

Table 23: Translation prompt templates for different methods