

Beyond Literal Mapping: Benchmarking and Improving Non-Literal Translation Evaluation

Yanzhi Tian^{1*}† Cunxiang Wang^{2,3*} Zeming Liu⁴ Heyan Huang¹
Wenbo Yu² Dawei Song¹ Jie Tang³ Yuhang Guo^{1‡}

¹School of Computer Science and Technology, Beijing Institute of Technology

²Zhipu AI ³The Knowledge Engineering Group (KEG), Tsinghua University

⁴School of Computer Science and Engineering, Beihang University

{tianyanzhi, guoyuhang}@bit.edu.cn; wangcunxiang303@gmail.com; zmliu@buaa.edu.cn

Abstract

Large Language Models (LLMs) have significantly advanced Machine Translation (MT), applying them to linguistically complex domains such as Social Network Services, literature etc. In these scenarios, translations often require handling non-literal expressions, leading to the inaccuracy of MT metrics. To systematically investigate the reliability of MT metrics, we first curate a meta-evaluation dataset focused on non-literal translations, namely **MENT**. **MENT** encompasses four non-literal translation domains and features source sentences paired with translations from diverse MT systems, with 7,530 human-annotated scores on translation quality. Experimental results reveal the inaccuracies of traditional MT metrics and the limitations of LLM-as-a-Judge, particularly the knowledge cutoff and score inconsistency problem. To mitigate these limitations, we propose RATE, a novel agentic translation evaluation framework, centered by a reflective Core Agent that dynamically invokes specialized sub-agents. Experimental results indicate the efficacy of RATE, achieving an improvement of at least 3.2 points in combined system- and segment-level correlation with human judgments compared with current methods. Further experiments demonstrate the robustness of RATE to general-domain MT evaluation. Code and dataset are available at: <https://github.com/BITHLP/RATE>.

1 Introduction

Large Language Models (LLMs) have recently demonstrated remarkable capabilities across a wide range of NLP tasks (Zhao et al., 2025), and machine translation (MT) stands out as a particularly important application, not only because of its long-standing history in NLP, but also due to its practical significance in enabling global communi-

*Equal Contribution.

†Work was done when interned at Z.ai.

‡Corresponding Author.

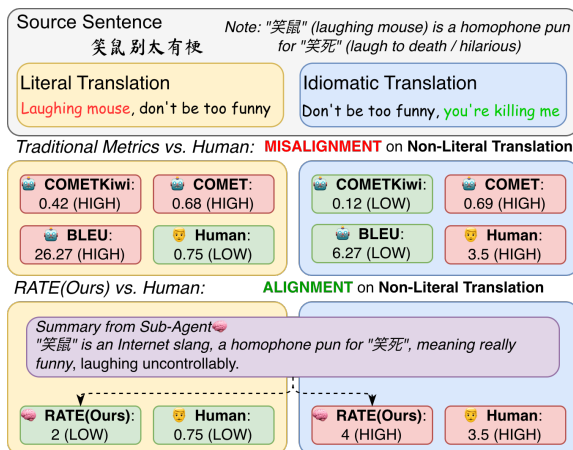


Figure 1: An illustration of evaluation misalignment in traditional metrics whereas the alignment with RATE. Traditional metrics often over-score literal but semantically incorrect translation (left) and penalize idiomatic translation (right). In contrast, our RATE leverages sub-agents to retrieve background knowledge and score calibration, achieving alignment with human judgment.

tion (Zhu et al., 2024b; Xu et al., 2024). LLMs have significantly expanded the application scope of MT, with the translation of online content—such as Social Network Services (SNS) on X, Facebook, and Rednote—emerging as a critical research direction. Effectively translating such content necessitates that models move beyond literal mapping to deeply understand internet slang, cross-cultural idioms, and literary contents (Yao et al., 2024a; Cheng et al., 2025; Guo et al., 2025b; Chen et al., 2025).

Given the complexities inherent in these non-literal translated scenarios, evaluating translation quality also becomes a challenge. A precise assessment of translation quality requires relevant metrics that enable a deep understanding of both the source sentence and the translated sentence. Through example from preliminary experiment in Figure 1, it can be observed that evaluation using traditional MT metrics for translations in such domains fails

to fully understand the complex sentence, resulting in the evaluation misalignment.

Systematically investigating the misalignment of MT metrics necessitates the construction of meta-evaluation dataset (Fomicheva and Specia, 2019; Moghe et al., 2025). Meta-evaluation is used to evaluate the correlation between MT automatic metrics and human judgment, which plays a crucial role in ensuring the accuracy and reliability of the metrics. However, existing datasets used for evaluating these metrics typically contain source and translated texts which majority comes from conventional domains such as news and Wikipedia (Karpinska et al., 2022; Freitag et al., 2023; Moghe et al., 2025). These datasets lack exploration in the aforementioned non-literal domains raises concerns about the reliability of current MT metrics.

To assess the reliability of MT metrics, we firstly curate MENT, a meta-evaluation dataset characterized by non-literally translation, that the source sentences are collected from SNS, cross-culture, poetry, and literature domains. Our dataset comprises 7,530 human-annotated scores on translation quality, which allows for a systematic evaluation of current MT metrics (Papineni et al., 2002; Rei et al., 2020), including LLM-as-a-Judge paradigms (Kocmi and Federmann, 2023a; Lu et al., 2024). Specifically, the collected source sentences are paired with several translations obtained from various MT systems, ranging from traditional NMT model to LLMs (NLLBTeam et al., 2022; Cheng et al., 2025; Yang et al., 2025; Team, 2025). Each translation is evaluated by at least two annotators, with quality control to ensure the reliability of the annotations. Our evaluation results indicate that while LLM-as-a-Judge generally outperforms traditional MT metrics on the constructed dataset, but it still exhibits notable inaccuracies due to the inherent limitations of LLMs such as knowledge cutoffs, which affects the accurate assessment of recently emerged slang or evolving cultural expression.

To overcome the limitations of static LLM-as-a-Judge methods, particularly the knowledge cutoff and score inconsistency encountered when evaluating non-literal translations, we introduce a dynamic reflective procedure and propose a novel agentic framework RATE (Reflective Agentic Translation Evaluation). To achieve the reflective evaluation framework, RATE is powered by a Core Agent, which dynamically determines the selection of specialized sub-agents through its reasoning process. It decides whether to invoke the Search Agent for

external knowledge retrieval, delivering what kind of instruction and retrieved context to get a score from Evaluation Agent, or calibrating the evaluated score from Comparison Agent. Experimental results on MENT dataset indicate the efficacy of RATE, which achieves an improvement of at least 3.2 points in combined system- and segment-level correlation with human judgments compared to current metrics including LLM-as-a-Judge for MT evaluation. Further analysis experiments demonstrate the robustness of RATE, that its reliability is not limited to non-literal scenarios but also extends to general-domain MT evaluation.

Our contributions are summarized as follows:

- To the best of our knowledge, we are the first to identify the challenges in accurately evaluating non-literal translation quality, and we construct a human-annotated meta-evaluation dataset MENT to systematically assess MT evaluation metrics.
- Our comprehensive meta-evaluation reveals the unreliability of MT metrics on non-literal content. Traditional metrics are fundamentally limited by the lack of deep semantic understanding, while LLM-as-a-Judge paradigms are hindered by the static knowledge cutoff and inherent score inconsistency.
- We propose a novel agentic translation evaluation framework RATE, which enables the dynamic invoking of specialized sub-agents by Core Agent. Experimental results demonstrate that RATE enhances reliability in non-literal translation evaluation while maintaining robustness across general domains.

2 Related Work

2.1 Machine Translation Metrics

Accurately evaluating translation quality is crucial, as it not only enables reliable assessment and iteration of MT systems, but also provides trustworthy reward signals for preference optimization and reinforcement learning (Feng et al., 2025a). Historically, rule-based metrics like BLEU (Papineni et al., 2002) have been the standard evaluation metric across machine translation modalities (Yao et al., 2024b; Liang et al., 2024; Tian et al., 2025b,a) due to efficiency and interpretability, yet they struggle to capture semantic lexical variability and diverse expressions (Freitag et al., 2023). While

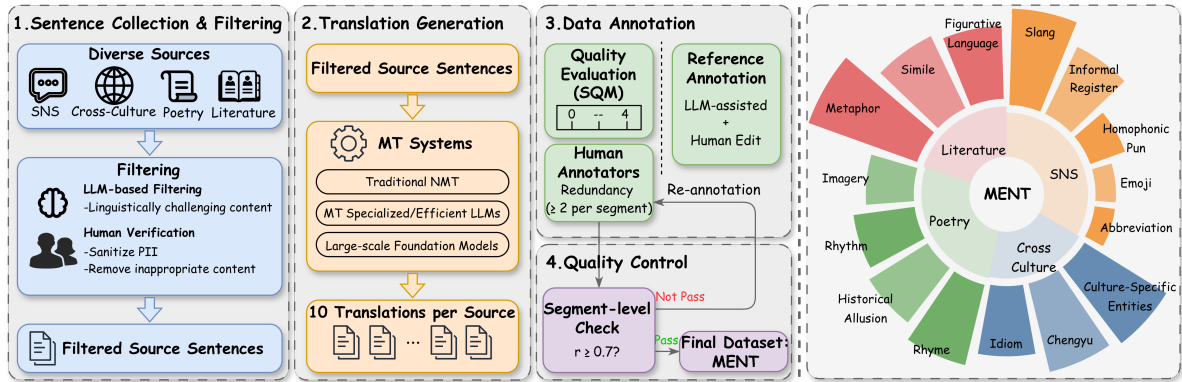


Figure 2: Overview of the data construction pipeline and final dataset visualization. **(Left)** The four-stage construction pipeline, proceeding from multi-domain sentences collection to strict quality control. **(Right)** Visualization of the MENT dataset, illustrating the distribution of source domains and linguistic phenomena.

model-based metrics (Rei et al., 2020; Juraska et al., 2023) improve semantic understanding through pre-trained encoders. The recent emergence of LLM-as-a-judge methods (Kocmi and Federmann, 2023a; Zhan et al., 2025) further leverages deep semantic reasoning capabilities of LLMs to achieve higher correlations with human judgments. Multi-agent frameworks (Feng et al., 2025b) attempt to extend this by coordinating multiple perspectives and hierarchical evaluation.

2.2 Machine Translation Meta-Evaluation

The MT meta-evaluation, often referred to as “evaluating MT metrics”, aims to validate the reliability of automatic metrics by measuring their alignment with human judgments. The majority of existing meta-evaluation datasets are constructed using literal or formal source texts (Freitag et al., 2023; Moghe et al., 2025; Karpinska et al., 2022), potentially overlooking complex linguistic phenomena.

Recent research explore the MT meta-evaluation in more challenging domains through small-scale human annotations. Chen et al. (2025) manually evaluate 100 Chinese poetry translations, and Zhang et al. (2025a) annotate 300 Chinese web novel translations with multi-dimensional scores. Although these studies highlight the importance of extending MT meta-evaluation to more diverse and challenging domains, their annotation scale remains highly limited, with all datasets containing fewer than 1,000 annotated scores. Such small-scale datasets make it difficult to systematically validate the reliability and robustness of evaluation methods, which also underscores the inherent difficulty of constructing large-scale meta-evaluation datasets in these challenging domains.

2.3 Agentic Evaluation

Despite the success in capturing semantic nuances, conventional LLM judges often struggle with complex, multi-step reasoning tasks where reward signals are sparse or lack objective verifiability. To address these limitations, a new paradigm of “Agent-as-a-Judge” that employed agentic evaluation has emerged (You et al., 2026), where the evaluator functions as an autonomous agent capable of active reasoning and tool interaction. Peng et al. (2025) introduce Agentic Reward Modeling, which enhances the reliability of reward signals by integrating human preferences with verifiable correctness signals through an agentic workflow. Zhuge et al. (2025) demonstrate the efficacy of utilizing autonomous agents to verify hierarchical solution requirements, offering a scalable and cost-effective alternative to human evaluation. Collectively, these advancements signify a pivotal shift toward more transparent, interpretable, and robust evaluation methodologies.

3 Data Construction

To systematically investigate the reliability of MT metrics, we construct **Meta-Evaluation** dataset of **Non-Literal Translation (MENT)** following a four-stage pipeline, shown in Figure 2: (1) Sentence Collection and Filtering, (2) Translation Generation, (3) Data Annotation, and (4) Quality Control.

3.1 Sentence Collection and Filtering

Aligning with current research in MT meta-evaluation (Karpinska et al., 2022; Moghe et al., 2025), we primarily collect source sentences from established MT benchmarks. Specifically, we focus

on four challenging domains: Social Network Services (SNS), Cross-Culture, Poetry, and Literature. These domains are selected for their characterized with non-literal translation (Guo et al., 2025b; Yao et al., 2024a; Chen et al., 2025; Wang et al., 2025). The following is a detailed description of the data sources for each domain.

Social Network Services (SNS) Domain: The SNS domain is characterized by high-context linguistic phenomena, including Internet slang, abbreviations, homophonic puns, informal register etc. These features often defy literal translation constraints. We collect English sentences from SNS-Bench (Guo et al., 2025a) and RedTrans (Guo et al., 2025b), two datasets that contain user-generated content crawled from a social media, including user posted notes and comments, covering a wide range of topics. To augment the dataset with more linguistically challenging Chinese social media texts, we crawl sentences from Chinese SNS platforms.

Cross-Culture Domain: Texts in this domain feature culture-specific entities and idiomatic expressions that lack direct equivalents in the target language, necessitating cultural adaptation rather than literal translation. For English, we source sentences from CAMT (Yao et al., 2024a), an MT benchmark rich in culturally specific items. For Chinese, we leverage ChID (Zheng et al., 2019), a dataset designed for cloze tests focusing on Chinese idioms, Chengyu.

Poetry Domain: Poetry presents unique challenges in preserving aesthetic features such as rhyme, rhythm, and imagery alongside semantic meaning, often requiring significant restructuring. We crawl English poems from the Project Gutenberg repository. Chinese poems are derived from PoetMT (Chen et al., 2025), encompassing Tang poetry and Song lyrics.

Literature Domain: Literary translation demands the interpretation of figurative language, such as similes and metaphors, where surface-level translation often fails to convey the underlying intent. Accordingly, English literary texts are collected from DRT (Wang et al., 2025), an MT dataset including English novels. For Chinese, we leverage CMDAG (Shao et al., 2024), a multi-source dataset characterized by a high density of metaphorical expressions. We specifically extract samples from the prose subset of this corpus.

Model	Type	Params	Access
<i>Traditional NMT</i>			
NLLB	Transformer	3.3B	Open
<i>MT Specialized & Efficient LLMs</i>			
Seed-X	MT Specialized	7B	Open
Hunyuan-MT	MT Specialized	7B	Open
Qwen3-8B	General LLM	8B	Open
Tower-Plus-9B	MT Specialized	9B	Open
<i>Large-scale Foundation Models</i>			
Qwen3-235B-A22B	General LLM	235B	Open
GLM-4.5	General LLM	355B	Open
DeepSeek-V3.1	General LLM	671B	Open
GPT-4o	General LLM	-	Proprietary
Gemini-2.5Pro	General LLM	-	Proprietary

Table 1: Overview of MT systems used for translation generation, categorized by architecture and scale.

Data Filtering: To ensure our collected data exhibits high translation difficulty, specifically targeting non-literal translation and adheres to strict privacy standards, we employ the data filtering strategy combining LLM-based scoring and human verification. Initially, all collected source sentences undergo a preliminary filtering process using an LLM to retain linguistically challenging content (see Appendix A.1 for prompt details). Subsequently, we conduct a manual inspection of the pre-filtered data, with particular attention to the newly crawled sentences. During this phase, we discard translation-irrelevant noise (e.g., topic hashtags in SNS content) and strictly sanitize the text to eliminate Personally Identifiable Information (PII), as well as offensive or inappropriate content.

3.2 Translation Generation

To construct a meta-evaluation dataset with diverse translation qualities and distinct error distributions, we employ several different MT systems to generate translations. These systems span a wide spectrum of architectures and scales, ranging from traditional NMT to LLMs. Specifically, we categorize the models into three groups based on their architecture and scale. **Traditional NMT:** NLLB (NLLBTeam et al., 2022). **MT Specialized and Efficient LLMs:** Seed-X (Cheng et al., 2025), Hunyuan-MT (Zheng et al., 2025), Qwen3-8B (Yang et al., 2025), and Tower-Plus-9B (Rei et al., 2025). **Large-scale Foundation Models:** Qwen3-235B-A22B (Yang et al., 2025), GLM-4.5 (Zeng et al., 2025), DeepSeek-V3.1 (DeepSeek-AI, 2024), GPT-4o (OpenAI, 2024), and Gemini-2.5Pro (Team, 2025). Detailed specifications of each system are summarized in Table 1.

3.3 Data Annotation

All recruited annotators hold degrees in translation and possess extensive prior experience in translation evaluation tasks.

Translation Quality. We adopt the Scalar Quality Metrics (SQM) framework, a human evaluation protocol used in WMT tasks (Freitag et al., 2021). The choice of the SQM protocol is primarily driven by the distinct characteristics of non-literal translation in MENT. Unlike literal mapping, failures in non-literal contexts often stem from a holistic misunderstanding of the source text’s slang or cultural nuances, rather than isolated lexical errors. Consequently, identifying specific error spans under an annotation protocol like Multidimensional Quality Metrics (MQM) (Lommel et al., 2013) becomes highly subjective, as the error often permeates the entire sentence structure. Following previous practices on human evaluation (Zhu et al., 2024a; Cheng et al., 2025; Zheng et al., 2025), we employ a 5-point scale based on the specific criteria illustrated in Appendix A.2.

To ensure the reliability of the annotations, we introduce redundancy into the annotation process. Specifically, each segment (10 translations corresponding to a source sentence) is evaluated by at least two independent annotators, and the final quality score for each translation is derived by averaging these individual ratings.

Translation Reference. To assess the reference-based metrics, we further manually annotate the translation reference. We utilize an LLM-assisted approach where each sentence is translated 5 times by Gemini-3.0Pro, one of the most advanced LLMs currently available, and annotators are asked to select and post edit the best translation. During annotation, 29.0% of references are post edited. For the sentences collected from PoetMT (Chen et al., 2025), we adopt the references in the dataset, since these references are annotated by professional translator. Specifically, we perform manual sentence-level alignment and review of the references from original dataset. For detailed annotation guideline, please refer to Appendix A.2.

3.4 Quality Control

Segment-level Agreement. To ensure quality control, we enforce a strict agreement threshold at the segment level. Specifically, for each segment, the scores assigned by any pair of annotators

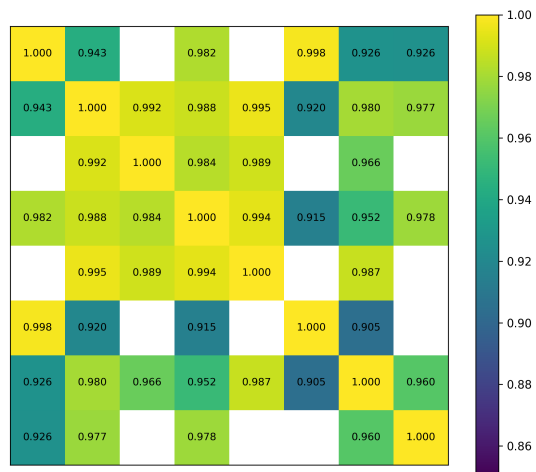


Figure 3: Heatmap of Pearson correlations for system-level inter-annotator agreement. Each cell (i, j) displays the correlation coefficient calculated on the aggregated scores of shared translations. Blank cells denote pairs with no overlapping annotation tasks due to the workload distribution.

	ZH-EN	EN-ZH	Total
SNS	1,100	710	1,810
Cross-Culture	950	920	1,870
Poetry	1,060	840	1,900
Literature	870	1,080	1,950
Total	3,980	3,550	7,530

Table 2: Statistic of MENT, the numbers represent the count of translations (Segments \times 10 MT Systems) collected for each domain and direction.

across the translations must exhibit a Pearson correlation coefficient r of at least 0.7. If this threshold is not met, the segment is flagged as unreliable and sent for re-annotation.

During the quality control, 22.2% of segments fail to meet the consistency criteria and subsequently undergo re-annotation to resolve discrepancies. Among these, 4.7% still remain inconsistent and are discarded as the annotators cannot reach a consensus, thereby ensuring the high reliability of the final dataset.

3.5 Data Statistics and Quality

Table 2 provides statistics of MENT. Samples of our annotated dataset are shown in Appendix A.4.

To evaluate the quality of the human-assigned translation scores, we measure the **Inter-Annotator Agreement (IAA)** using Pearson correlation coefficient between annotators based on their shared workload. As visualized in Figure 3, the value at position (i, j) in the heatmap

represents the correlation coefficient between annotator i and annotator j , calculated over the set of segments that they both annotate. Blank cells in the heatmap indicate pairs of annotators who do not share any overlapping assignments. The high correlation values (> 0.9) indicate high reliability of the annotated translation quality scores. Detailed calculation method of IAA is introduced in Appendix A.3.

The quality of human-annotated references is further verified through a **secondary review process**, achieving a pass rate of 96.1%, thereby guaranteeing the quality of annotated references.

4 RATE: Reflective Agentic Translation Evaluation

To mitigate the limitations in current metrics for MT evaluation, especially to the knowledge cut-off and score inconsistency of LLM-as-a-Judge, we propose an agentic evaluation framework **Reflective Agentic Translation Evaluation (RATE)**.

In contrast to existing LLM-based paradigms for MT evaluation (Kocmi and Federmann, 2023a; Feng et al., 2025b; Zhang et al., 2025b), RATE is not constrained by a static workflow with fixed, sequential evaluation procedures, or limited external tool usage (e.g., search engine). It is architected around a centralized **Core Agent**, which orchestrates three functional sub-agents: the **Evaluation Agent** for pointwise assessment, the **Search Agent** for online knowledge retrieval, and the **Comparison Agent** for calibration by pairwise evaluation, as shown in Figure 4.

4.1 Architectures

Core Agent. The Core Agent serves as the central controller of the framework. Unlike a simple router, it operates on a reflective loop based on OODA (Observe, Orient, Decide, and Act). Tracking the current understanding of the source sentence, the accumulated knowledge, and the confidence level of current judgments, it automatically determines whether to request a score (via Evaluation), retrieve more information (via Search), or verify a judgment (via Comparison). This process operates iteratively, and the Core Agent orchestrates multi-turn interactions with sub-agents to continuously refine its assessment. The loop terminates only when the agent determines that sufficient evidence has been gathered to output a final evaluation, or when a pre-defined maximum round limit

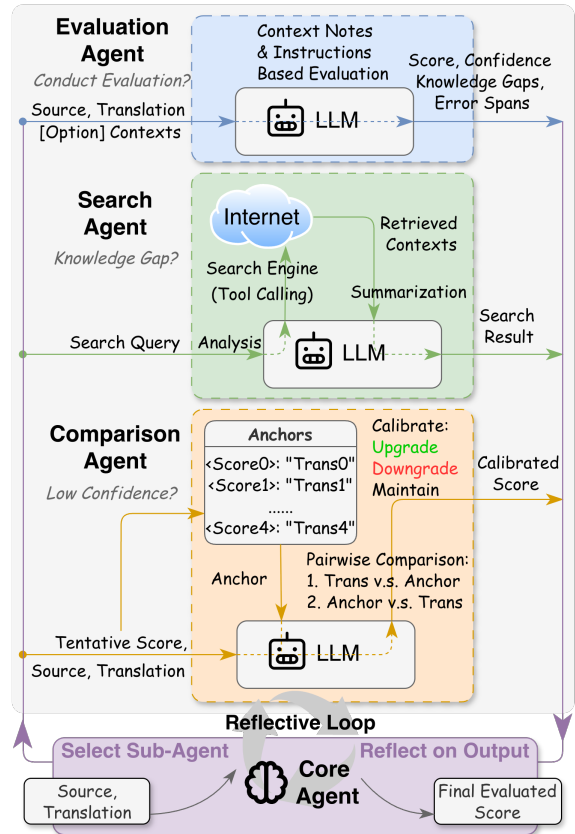


Figure 4: Overview of the RATE framework. The Core Agent acts as the central controller, dynamically selecting specialized sub-agents (Evaluation, Search, and Comparison) based on current state and outputs of sub-agents to iteratively refine the translation evaluation.

is reached.

Evaluation Agent. This agent functions as the primary pointwise assessor, tasked with analyzing the translation quality. Crucially, it is designed to accept the optional context notes and instructions, with explanations of necessary knowledge injected from Core Agent, to ensure the evaluation is not based on hallucinations. Furthermore, beyond producing a scalar score with confidence and scoring rationale, this agent is asked to return the error spans and suspected knowledge gaps, thereby triggering the Core Agent to activate other agent for clarification.

Search Agent. This agent is used to mitigate the “knowledge cutoff” limitation inherent in LLMs. It is invoked on-demand by the Core Agent when the source text contains ambiguous entities, emerging internet slang, or deep cultural idioms that cannot be resolved via internal parametric knowledge of LLMs. The Search Agent first analyzes request from Core Agent, and then transfer it into search

Metric	Meta	ZH-EN						EN-ZH					
		System-Level			Segment-Level			System-Level			Segment-Level		
		Acc.	r	ρ	Acc-t.	r	ρ	Acc.	r	ρ	Acc-t.	r	ρ
<i>Reference-based Metrics</i>													
BLEU	64.9	82.2	91.1	84.2	52.0	30.1	30.7	91.1	92.6	91.5	56.6	37.0	39.5
BLEURT	56.5	88.9	98.4	90.3	55.4	45.4	40.5	57.8	82.5	17.6	50.5	28.5	21.8
COMET	69.0	88.9	98.2	91.5	56.5	50.3	43.3	82.2	92.0	79.4	56.5	50.3	38.8
XCOMET	54.5	80.0	96.7	73.3	48.9	29.3	21.2	68.9	86.6	40.6	52.0	31.1	25.8
MetricX-23	52.2	73.3	94.7	62.4	47.5	31.5	17.3	64.4	83.6	27.3	53.2	41.2	29.4
MetricX-24	56.2	73.3	94.6	55.2	53.8	47.1	35.9	66.7	83.5	29.7	54.5	47.1	33.0
<i>Reference-free (QE) Metrics</i>													
COMETKiwi	42.8	51.1	90.6	5.5	52.3	38.6	31.6	48.9	79.5	-1.8	52.3	37.8	26.9
MetricX23-QE	31.9	46.7	86.4	-5.5	41.9	15.2	0.4	35.6	79.4	-28.5	50.2	39.9	20.8
MetricX24-QE	41.7	48.9	88.4	-1.8	47.9	31.7	18.6	55.6	79.3	12.7	51.8	41.5	25.4
<i>LLM-as-a-Judge</i>													
GEMBA-MQM	68.7	82.2	97.0	80.6	46.2	56.0	26.8	91.1	97.2	93.9	52.1	58.2	42.6
GEMBA-DA	77.2	86.7	96.5	87.9	59.1	70.1	60.2	91.1	94.5	95.2	56.7	68.4	60.5
EAPrompt	47.4	82.2	93.6	74.5	36.1	-6.4	-10.9	75.6	92.4	67.3	43.3	12.2	9.1
ThinMQM	65.1	80.0	95.1	80.6	49.5	34.8	29.8	91.1	91.5	91.5	50.4	43.8	43.2
M-MAD	65.8	82.2	96.3	83.3	43.9	39.6	23.0	88.9	95.2	90.3	51.5	50.1	45.2
RATE (Ours)	80.4	97.8	99.3	99.7	61.9	74.5	66.4	88.9	97.7	92.7	59.5	65.3	60.1

Table 3: System-level and segment-level correlations on the MENT dataset. We report **Accuracy (Acc., Acc-t.)**, **Pearson(r)**, and **Spearman (ρ)** correlation coefficients, scaled by a factor of 100. Meta represents the average score of all accuracies and correlation coefficients. The **best**, **second-best**, and **third-best** results are marked with purple, orange, and gray backgrounds respectively.

query to call the search engine. After obtaining several responses from search engine, it summarizes the related responses and return it to Core Agent.

Comparison Agent. This agent serves as a critical calibration module. It is specifically designed to mitigate the inherent subjectivity and potential inaccuracies associated with the pointwise evaluation paradigm employed by the Evaluation Agent, with utilization previously evaluated translations as anchors. When the Core Agent is uncertain about a new translation’s quality, it instructs this agent to perform a pairwise preference ranking against these anchors. This mechanism converts subjective absolute judgments into robust relative rankings, mitigating score inconsistency of the evaluation.

The detail implementations of RATE are introduced in Appendix B. We provide evaluation trajectories of RATE, shown in Appendix C.

5 Experiments

5.1 Experimental Setup

Meta-Evaluation. Following WMT23 Metrics Shared Task (Freitag et al., 2023), we adopt the composite meta score to assess metric performance across both system level and segment level. Our evaluation includes Accuracy (Acc (Kocmi et al., 2021), Acc-t (Deutsch et al., 2023)), Pearson Corre-

lation (r), and Spearman correlation (ρ). The final meta score is calculated as the average of individual statistics (Feng et al., 2025b). All results are computed based on MTME, the standard metric evaluation tool recommended by WMT.

Evaluated Metrics. We evaluate a comprehensive set of metrics categorized into three paradigms: Reference-based metrics, Reference-free (Quality Estimation, QE) metrics, and LLM-as-a-Judge adopting reference-free paradigm. Detailed experimental setup is introduced in Appendix F.

5.2 Main Results

The experimental results are shown in Table 3.

Reference-based Metrics. The gold-standard reference essentially introduces the necessary semantic correspondences for non-literal expressions, thereby allowing the metrics to maintain a certain level of reliability. As observed in Table 3, methods within this category perform competitively, with COMET achieving the highest meta score (69.0) among reference-based metrics.

Breakdown of Reference-free (QE) Metrics. Most QE models rely on pre-trained multilingual encoders, which are primarily trained on literal or formal corpora, and cannot verify the quality of

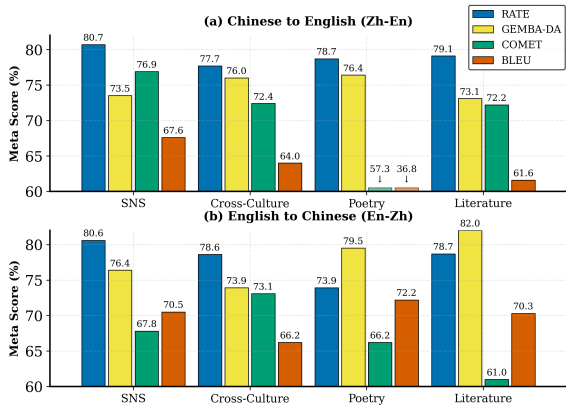


Figure 5: Illustration of metrics performance with specific domain, detailed numerical results in Appendix G.

non-literal translation. As clearly reflected in Table 3, this leads to a significant performance breakdown: QE metrics yield the lowest meta scores overall (ranging from 31.9 to 42.8).

The Power of Semantic Interpretation in LLM-as-a-Judge. LLM-as-a-Judge paradigms demonstrate remarkable human alignment even without access to gold-standard references. For instance, GEMBA-DA achieves a strong meta score of 77.2 in Table 3. This efficacy stems from the inherent capability of LLMs to perform deep semantic interpretation. RATE further enhances the performance by dynamically invoking the sub-agents, and achieves the highest meta score of 80.4 among all evaluation methods.

The additional experimental results on system-level Soft Pairwise Accuracy (SPA) are shown in Appendix E, further demonstrating the effectiveness of RATE in providing reliable system-level rankings for non-literal translations.

5.3 Domain-Specific Performance Analysis

We further investigate the performance of different evaluation paradigms with specific domain. The experimental results are illustrated in Figure 5, and more detailed results are shown in Appendix G.

SNS and Cross-Culture Domains: The primary challenge of evaluating translations in these two domains lies in the utilization of emerging slang, and idiomatic expressions. These linguistic phenomena are frequently absent from the static pre-training corpora of LLMs, leading to “knowledge cutoff” that hinders the performance of LLM-as-a-Judge paradigm. Furthermore, the subtle semantic nuances distinguishing various translation candidates

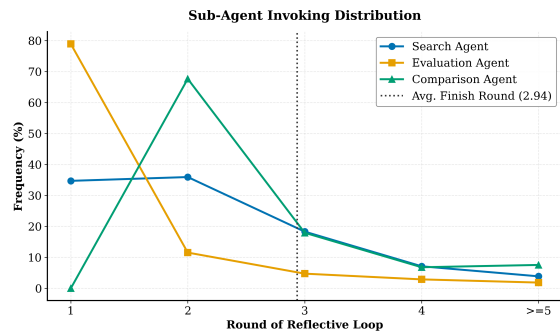


Figure 6: Temporal distribution of sub-agent invoking, illustrating the transition from initial evaluation to background knowledge-enhanced calibration.

in these contexts can lead to significant distinction in overall quality, and LLMs often fail to assess with consistent scoring. With dynamically invoking sub-agents, RATE enhances the understanding by external knowledge, or calibrates the score by pairwise comparison, thereby mitigates the limitations of current LLM-as-a-Judge methods.

Poetry and Literature Domains: These domains focus on figurative complexities, such as metaphor and imagery, rather than currently evolving slang. In the En-Zh direction, current LLM-as-a-Judge methods already demonstrate superior performance, significantly outperforming traditional metrics by leveraging their deep internalized semantic reasoning. However, evaluation in the Zh-En direction remains uniquely challenging due to the prevalence of historical allusions in Chinese poetry and literature. In such cases, Search Agent of RATE serves as a vital knowledge bridge, injecting necessary background context to decipher these allusions, ensuring that the evaluation is grounded in cultural accuracy rather than just linguistic fluency, further enhancing reliability where static LLMs reach their limits.

6 Analysis

6.1 Temporal Dynamics and Behavioral Analysis of Sub-Agents Invoking

We investigate the temporal and behavioral patterns of sub-agent invocations, with the distribution visualized in Figure 6. Detailed agentic evaluation trajectories are provided in Appendix C.

The invoking of the Search Agent is primarily concentrated in the early stages, where the Core Agent directly requests external knowledge to resolve non-literal translation content. However, it is

also observed in later rounds. This indicates that feedback from the Evaluation Agent may reveal latent comprehension gaps that are not initially apparent, and the Core Agent asks the Search Agent for deeper clarification.

The Evaluation Agent is also mainly selected in the initial stages. For less complex texts, the Core Agent firstly attempts invoking the Evaluation Agent, to determine if a high-confidence assessment can be achieved immediately. For more complex texts, the Core Agent can leverage the accumulated context, allowing the Evaluation Agent to proceed with a knowledge-enhanced assessment. Detailed implementation of accumulation mechanism is provided in Appendix B.

Consistent with our design expectations, the Comparison Agent is not invoked in the first round, since the agent is triggered only when the Core Agent lacks confidence in the preliminary score provided by the Evaluation Agent, serving as a necessary mechanism to calibrate the score.

6.2 Ablation Study

To quantify the contributions of sub-agents in the reflective framework, we conduct an ablation study by restricting the Core Agent’s calling protocol (removing sub-agent calling in prompt).

Setting	Meta		
	ZH-EN	EN-ZH	Avg.
RATE	83.3	77.4	80.4
w/o Search Agent	78.2 (-5.1)	76.9 (-0.5)	77.6 (-2.8)
w/o Comparison Agent	81.7 (-1.6)	76.3 (-1.1)	79.0 (-1.4)
w/o Both	77.6 (-5.7)	76.0 (-1.4)	76.8 (-3.6)

Table 4: Experimental results of ablation study. Specifically, we evaluate three variants: (1) w/o Search Agent; (2) w/o Comparison Agent; and (3) w/o Both, which simplifies the framework into only a Evaluation Agent cooperating with Core Agent.

As shown in Table 4, the removal of Search Agent results in the most significant performance degradation, confirming for non-literal translation data, the static internalized parameters of LLM are insufficient. Disabling Comparison Agent also leads to decreasing of performance, indicating the efficacy of pairwise calibration. The experimental results demonstrate the necessity of dynamic sub-agents utilization in RATE framework.

6.3 Evaluation of RATE on General Dataset

We conduct experiments on WMT23 En-De Metrics Shared Task, which represents a standard gen-

eral domain meta-evaluation dataset, with new language direction, and significantly differs from our primary focus on non-literal translation scenarios. Crucially, we **maintain the same system prompts and agentic protocols without any modifications.**

Metric	System-Level		Segment-Level	
	Acc.	<i>r</i>	Acc-t.	<i>r</i>
<i>Reference-based Metrics</i>				
BLEU	89.4	91.7	52.0	19.2
COMET	97.0	99.0	57.4	43.2
MetricX-23	90.9	97.7	60.3	<u>58.5</u>
<i>Reference-free (QE) Metrics</i>				
COMETKiwi	<u>98.5</u>	94.6	56.9	47.5
MetricX23-QE	92.4	96.7	60.3	62.6
<i>LLM-as-a-Judge</i>				
EAPrompt*	93.9	96.2	47.1	52.0
GEMBA-MQM*	97.0	97.3	47.4	42.9
M-MAD*	97.0	97.9	55.5	55.2
RATE (Ours)	98.5	<u>99.0</u>	52.3	38.9

Table 5: Performance on the WMT23 En-De Metrics Shared Task (Freitag et al., 2023). * indicates the results are derived from Feng et al. (2025b).

Experimental results in Table 5 show that RATE achieves a comparable performance to metrics on system-level translation evaluation, indicating robustness of the proposed RATE.

7 Conclusion

In this paper, we identify the challenges in accurately evaluating non-literal translation. To systematically investigate the reliability of MT metrics, we construct a human-annotated meta-evaluation dataset focusing on non-literal translation, MENT. Our comprehensive evaluation reveals the inaccuracy of MT metrics. To mitigate the limitations of knowledge cutoff and score inconsistency of LLM-as-a-Judge methods, we propose a novel agentic translation evaluation framework RATE. Experimental results demonstrate that RATE enhances reliability in non-literal translation evaluation, while further analysis experiments indicate its robustness across general domains.

Limitations

While the RATE framework demonstrates improvement in translation evaluation, certain limitations remain regarding the sub-agent invoking within the reflective loop.

In the current architecture, specialized sub-agents and external search engines are invoked via

a tool-calling mechanism. However, these external services are susceptible to execution failures, such as LLM API timeouts or search engine connectivity errors. Presently, the Core Agent lacks the analytical capability to interpret diagnostic error messages or implement adaptive retry strategies. When an execution failure of external services occurs, system defaults to re-initiating the evaluation process rather than performing a localized recovery or retry, which could potentially hinder the overall efficiency of the framework, particularly in scenarios where external API stability fluctuates.

Ethics Statement

We take ethical considerations seriously and ensure that the data used in this study is conducted in a responsible and ethical manner. The primary data sources for this research include both newly curated web-crawled content and existing open-source datasets.

For the crawled data, we perform rigorous manual inspection and filtering to ensure that the dataset contains no Personally Identifiable Information (PII) or sensitive user data, and we exclude any content that could be deemed offensive, harmful, or culturally insensitive. Regarding the open-source datasets integrated into our study, we strictly adhere to their respective licenses and terms of use.

Acknowledgment

We thank all the anonymous reviewers for their insightful and valuable comments. This work is supported by the National Natural Science Foundation of China (Grant No. U21B2009, 62376027).

References

- Andong Chen, Lianzhang Lou, Kehai Chen, Xuefeng Bai, Yang Xiang, Muyun Yang, Tiejun Zhao, and Min Zhang. 2025. [Benchmarking LLMs for translating classical Chinese poetry: Evaluating adequacy, fluency, and elegance](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 33007–33024, Suzhou, China. Association for Computational Linguistics.
- Shanbo Cheng, Yu Bao, Qian Cao, Luyang Huang, Liyan Kang, Zhicheng Liu, Yu Lu, Wenhao Zhu, Jingwen Chen, Zhichao Huang, Tao Li, Yifu Li, Huiying Lin, Sitong Liu, Ningxin Peng, Shuaijie She, Lu Xu, Nuo Xu, Sen Yang, and 7 others. 2025. [Seed-x: Building strong multilingual translation llm with 7b parameters](#). *Preprint*, arXiv:2507.13618.
- DeepSeek-AI. 2024. [Deepseek-v3 technical report](#). *Preprint*, arXiv:2412.19437.
- Daniel Deutsch, George Foster, and Markus Freitag. 2023. [Ties matter: Meta-evaluating modern metrics with pairwise accuracy and tie calibration](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12914–12929, Singapore. Association for Computational Linguistics.
- Zhaopeng Feng, Shaosheng Cao, Jiahao Ren, Jiayuan Su, Ruizhe Chen, Yan Zhang, Jian Wu, and Zuozhu Liu. 2025a. [MT-r1-zero: Advancing LLM-based machine translation via r1-zero-like reinforcement learning](#). In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 18685–18702, Suzhou, China. Association for Computational Linguistics.
- Zhaopeng Feng, Jiayuan Su, Jiamei Zheng, Jiahao Ren, Yan Zhang, Jian Wu, Hongwei Wang, and Zuozhu Liu. 2025b. [M-MAD: Multidimensional multi-agent debate for advanced machine translation evaluation](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7084–7107, Vienna, Austria. Association for Computational Linguistics.
- Marina Fomicheva and Lucia Specia. 2019. [Taking MT evaluation metrics to extremes: Beyond correlation with human judgments](#). *Computational Linguistics*, 45(3):515–558.
- Markus Freitag, George Foster, David Grangier, Viresh Ratnakar, Qijun Tan, and Wolfgang Macherey. 2021. [Experts, errors, and context: A large-scale study of human evaluation for machine translation](#). *Transactions of the Association for Computational Linguistics*, 9:1460–1474.
- Markus Freitag, Nitika Mathur, Chi-kiu Lo, Eleftherios Avramidis, Ricardo Rei, Brian Thompson, Tom Kocmi, Frederic Blain, Daniel Deutsch, Craig Stewart, Chrysoula Zerva, Sheila Castilho, Alon Lavie, and George Foster. 2023. [Results of WMT23 metrics shared task: Metrics might be guilty but references are not innocent](#). In *Proceedings of the Eighth Conference on Machine Translation*, pages 578–628, Singapore. Association for Computational Linguistics.
- Nuno M. Guerreiro, Ricardo Rei, Daan van Stigt, Luisa Coheur, Pierre Colombo, and André F. T. Martins. 2024. [xCOMET: Transparent machine translation evaluation through fine-grained error detection](#). *Transactions of the Association for Computational Linguistics*, 12:979–995.
- Hongcheng Guo, wangyue, Shaosheng Cao, Fei Zhao, Boyang Wang, Lei Li, Liang Chen, Xinze Lyu, Zhe Xu, Yao Hu, and Zhoujun Li. 2025a. [SNS-bench: Defining, building, and assessing capabilities of large language models in social networking services](#). In *Forty-second International Conference on Machine Learning*.

- Hongcheng Guo, Fei Zhao, Shaosheng Cao, Xinze Lyu, Ziyang Liu, Yue Wang, Boyang Wang, Zhoujun Li, Chonggang Lu, Zhe Xu, and Yao Hu. 2025b. [Redefining machine translation on social network services with large language models](#). *Preprint*, arXiv:2504.07901.
- Juraj Juraska, Daniel Deutsch, Mara Finkelstein, and Markus Freitag. 2024. [MetricX-24: The Google submission to the WMT 2024 metrics shared task](#). In *Proceedings of the Ninth Conference on Machine Translation*, pages 492–504, Miami, Florida, USA. Association for Computational Linguistics.
- Juraj Juraska, Mara Finkelstein, Daniel Deutsch, Aditya Siddhant, Mehdi Mirzazadeh, and Markus Freitag. 2023. [MetricX-23: The Google submission to the WMT 2023 metrics shared task](#). In *Proceedings of the Eighth Conference on Machine Translation*, pages 756–767, Singapore. Association for Computational Linguistics.
- Marzena Karpinska, Nishant Raj, Katherine Thai, Yixiao Song, Ankita Gupta, and Mohit Iyyer. 2022. [DEMETER: Diagnosing evaluation metrics for translation](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 9540–9561, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Tom Kocmi and Christian Federmann. 2023a. [GEMBA-MQM: Detecting translation quality error spans with GPT-4](#). In *Proceedings of the Eighth Conference on Machine Translation*, pages 768–775, Singapore. Association for Computational Linguistics.
- Tom Kocmi and Christian Federmann. 2023b. [Large language models are state-of-the-art evaluators of translation quality](#). In *Proceedings of the 24th Annual Conference of the European Association for Machine Translation*, pages 193–203, Tampere, Finland. European Association for Machine Translation.
- 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.
- Yupu Liang, Yaping Zhang, Cong Ma, Zhiyang Zhang, Yang Zhao, Lu Xiang, Chengqing Zong, and Yu Zhou. 2024. [Document image machine translation with dynamic multi-pre-trained models assembling](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 7084–7095, Mexico City, Mexico. Association for Computational Linguistics.
- Arle Richard Lommel, Aljoscha Burchardt, and Hans Uszkoreit. 2013. [Multidimensional quality metrics: a flexible system for assessing translation quality](#). In *Proceedings of Translating and the Computer 35*, London, UK. Aslib.
- Qingyu Lu, Baopu Qiu, Liang Ding, Kanjian Zhang, Tom Kocmi, and Dacheng Tao. 2024. [Error analysis prompting enables human-like translation evaluation in large language models](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 8801–8816, Bangkok, Thailand. Association for Computational Linguistics.
- Nikita Moghe, Arnisa Fazla, Chantal Amrhein, Tom Kocmi, Mark Steedman, Alexandra Birch, Rico Senrich, and Liane Guillou. 2025. [Machine translation meta evaluation through translation accuracy challenge sets](#). *Computational Linguistics*, 51(1):73–137.
- NLLBTeam, Marta R. Costa-jussà, James Cross, Onur Çelebi, Maha Elbayad, Kenneth Heafield, Kevin Hefernan, Elahe Kalbassi, Janice Lam, Daniel Licht, Jean Maillard, Anna Sun, Skyler Wang, Guillaume Wenzek, Al Youngblood, Bapi Akula, Loic Barrault, Gabriel Meja Gonzalez, Prangthip Hansanti, and 20 others. 2022. [No language left behind: Scaling human-centered machine translation](#). *Preprint*, arXiv:2207.04672.
- OpenAI. 2024. [Gpt-4o system card](#). *Preprint*, arXiv:2410.21276.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Hao Peng, Yunjia Qi, Xiaozhi Wang, Zijun Yao, Bin Xu, Lei Hou, and Juanzi Li. 2025. [Agentic reward modeling: Integrating human preferences with verifiable correctness signals for reliable reward systems](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15934–15949, Vienna, Austria. Association for Computational Linguistics.
- Ricardo Rei, José G. C. de Souza, Duarte Alves, Chrysoula Zerva, Ana C Farinha, Taisiya Glushkova, Alon Lavie, Luisa Coheur, and André F. T. Martins. 2022a. [COMET-22: Unbabel-IST 2022 submission for the metrics shared task](#). In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 578–585, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Ricardo Rei, Nuno M. Guerreiro, José Pombal, João Alves, Pedro Teixeira, Amin Farajian, and André F. T. Martins. 2025. [Tower+: Bridging generality and translation specialization in multilingual llms](#). *Preprint*, arXiv:2506.17080.
- Ricardo Rei, Craig Stewart, Ana C Farinha, and Alon Lavie. 2020. [COMET: A neural framework for MT evaluation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2685–2702, Online. Association for Computational Linguistics.

- Ricardo Rei, Marcos Treviso, Nuno M. Guerreiro, Chrysoula Zerva, Ana C Farinha, Christine Maroti, José G. C. de Souza, Taisiya Glushkova, Duarte Alves, Luisa Coheur, Alon Lavie, and André F. T. Martins. 2022b. [CometKiwi: IST-unbabel 2022 submission for the quality estimation shared task](#). In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 634–645, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Thibault Sellam, Dipanjan Das, and Ankur Parikh. 2020. [BLEURT: Learning robust metrics for text generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7881–7892, Online. Association for Computational Linguistics.
- Yujie Shao, Xinrong Yao, Xingwei Qu, Chenghua Lin, Shi Wang, Wenhao Huang, Ge Zhang, and Jie Fu. 2024. [CMDAG: A Chinese metaphor dataset with annotated grounds as CoT for boosting metaphor generation](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 3357–3366, Torino, Italia. ELRA and ICCL.
- Gemini Team. 2025. [Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities](#). Preprint, arXiv:2507.06261.
- Brian Thompson, Nitika Mathur, Daniel Deutsch, and Huda Khayrallah. 2024. [Improving statistical significance in human evaluation of automatic metrics via soft pairwise accuracy](#). In *Proceedings of the Ninth Conference on Machine Translation*, pages 1222–1234, Miami, Florida, USA. Association for Computational Linguistics.
- Yanzhi Tian, Zeming Liu, Zhengyang Liu, Chong Feng, Xin Li, Heyan Huang, and Yuhang Guo. 2025a. [PRIM: Towards practical in-image multilingual machine translation](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 13682–13697, Suzhou, China. Association for Computational Linguistics.
- Yanzhi Tian, Zeming Liu, Zhengyang Liu, and Yuhang Guo. 2025b. [Exploring in-image machine translation with real-world background](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 124–137, Vienna, Austria. Association for Computational Linguistics.
- Jiaan Wang, Fandong Meng, Yunlong Liang, and Jie Zhou. 2025. [DRT: Deep reasoning translation via long chain-of-thought](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 6770–6782, Vienna, Austria. Association for Computational Linguistics.
- Peiyi Wang, Lei Li, Liang Chen, Zefan Cai, Dawei Zhu, Binghuai Lin, Yunbo Cao, Lingpeng Kong, Qi Liu, Tianyu Liu, and Zhifang Sui. 2024. [Large language models are not fair evaluators](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9440–9450, Bangkok, Thailand. Association for Computational Linguistics.
- Haoran Xu, Young Jin Kim, Amr Sharaf, and Hany Hassan Awadalla. 2024. [A paradigm shift in machine translation: Boosting translation performance of large language models](#). In *The Twelfth International Conference on Learning Representations*.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, and 41 others. 2025. [Qwen3 technical report](#). Preprint, arXiv:2505.09388.
- Binwei Yao, Ming Jiang, Tara Bobinac, Diyi Yang, and Junjie Hu. 2024a. [Benchmarking machine translation with cultural awareness](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 13078–13096, Miami, Florida, USA. Association for Computational Linguistics.
- Jiashu Yao, Heyan Huang, Zeming Liu, and Yuhang Guo. 2024b. [Deterministic reversible data augmentation for neural machine translation](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 8075–8089, Bangkok, Thailand. Association for Computational Linguistics.
- Runyang You, Hongru Cai, Caiqi Zhang, Qiancheng Xu, Meng Liu, Tiezheng Yu, Yongqi Li, and Wenjie Li. 2026. [Agent-as-a-judge](#). Preprint, arXiv:2601.05111.
- GLM4.5Team: Aohan Zeng, Xin Lv, Qinkai Zheng, Zhenyu Hou, Bin Chen, Chengxing Xie, Cunxiang Wang, Da Yin, Hao Zeng, Jiajie Zhang, Kedong Wang, Lucen Zhong, Mingdao Liu, Rui Lu, Shulin Cao, Xiaohan Zhang, Xuancheng Huang, Yao Wei, Yean Cheng, and 151 others. 2025. [Glm-4.5: Agentic, reasoning, and coding \(arc\) foundation models](#). Preprint, arXiv:2508.06471.
- Runzhe Zhan, Zhihong Huang, Xinyi Yang, Lidia S. Chao, Min Yang, and Derek F. Wong. 2025. [Are large reasoning models good translation evaluators? analysis and performance boost](#). In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.
- Enze Zhang, Jiaying Wang, Mengxi Xiao, Jifei Liu, Ziyang Kuang, Rui Dong, Eric Dong, Sophia Ananiadou, Min Peng, and Qianqian Xie. 2025a. [Diting: A multi-agent evaluation framework for benchmarking web novel translation](#). Preprint, arXiv:2510.09116.
- Shijie Zhang, Renhao Li, Songsheng Wang, Philipp Koehn, Min Yang, and Derek F. Wong. 2025b. [Hi-MATE: A hierarchical multi-agent framework for machine translation evaluation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 11121–11145, Suzhou, China. Association for Computational Linguistics.

Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, Yifan Du, Chen Yang, Yushuo Chen, Zhipeng Chen, Jinhao Jiang, Ruiyang Ren, Yifan Li, Xinyu Tang, Zikang Liu, and 3 others. 2025. [A survey of large language models](#). *Preprint*, arXiv:2303.18223.

Chujie Zheng, Minlie Huang, and Aixin Sun. 2019. [ChID: A large-scale Chinese IDiom dataset for cloze test](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 778–787, Florence, Italy. Association for Computational Linguistics.

Mao Zheng, Zheng Li, Bingxin Qu, Mingyang Song, Yang Du, Mingrui Sun, and Di Wang. 2025. [Hunyuan-mt technical report](#). *Preprint*, arXiv:2509.05209.

Dawei Zhu, Sony Trenous, Xiaoyu Shen, Dietrich Klakow, Bill Byrne, and Eva Hasler. 2024a. [A preference-driven paradigm for enhanced translation with large language models](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 3385–3403, Mexico City, Mexico. Association for Computational Linguistics.

Wenhao Zhu, Hongyi Liu, Qingxiu Dong, Jingjing Xu, Shujian Huang, Lingpeng Kong, Jiajun Chen, and Lei Li. 2024b. [Multilingual machine translation with large language models: Empirical results and analysis](#). In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 2765–2781, Mexico City, Mexico. Association for Computational Linguistics.

Mingchen Zhuge, Changsheng Zhao, Dylan R. Ashley, Wenyi Wang, Dmitrii Khizbullin, Yunyang Xiong, Zechun Liu, Ernie Chang, Raghuraman Krishnamoorthi, Yuandong Tian, Yangyang Shi, Vikas Chandra, and Jürgen Schmidhuber. 2025. [Agent-as-a-judge: Evaluate agents with agents](#). In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pages 80569–80611. PMLR.

A Details of Data Construction

In this section, we introduce more details of the data construction.

A.1 Preliminarily Filtering with LLM

To ensure the curated MENT dataset maintains a high concentration of linguistically challenging and non-literal samples, we employ LLM as an initial automated filter. As illustrated in Figure 8, the LLM is prompted to assign higher scores to source sentences that exhibit challenges of non-literal translation, and the LLM is also required to generate a rationale for each score. This explanatory output serves as assistance during the subsequent manual inspection phase, allowing human experts to quickly pinpoint the specific linguistic complexities, such as slang, cultural idiom. Only the source sentences with higher score will go through the further manual inspection.

A.2 Data Annotation

To mitigate annotator fatigue and maintain consistency, we limit the daily workload of each annotator to approximately 100 ratings. This corresponds to evaluating the translation outputs of all 10 systems for 10 source sentences. All annotators are recruited as interns and receive fair compensation at competitive market rates, and are clearly informed that the data they annotated would be used for academic research.

The detail instruction of human annotation on evaluating the translation quality is shown in Figure 9. The detail instruction of human annotation on reference is shown in Figure 10.

A.3 Calculation of IAA

The IAA for each pair of annotators is determined as follows:

Let $\mathcal{K}_{i,j}$ denote the set of segments commonly annotated by annotators i and j . For a segment $k \in \mathcal{K}_{i,j}$, let $\mathbf{v}_{i,k} \in \mathbb{N}^{10}$ represent the score vector assigned by annotator i to the 10 MT systems. The correlation $r_{i,j}$ is defined as:

$$r_{i,j} = \text{Pearson} \left(\sum_{k \in \mathcal{K}_{i,j}} \mathbf{v}_{i,k}, \sum_{k \in \mathcal{K}_{i,j}} \mathbf{v}_{j,k} \right) \quad (1)$$

where the summation \sum is performed element-wise over the vectors. The resulting correlation measures the agreement between annotators regarding the relative quality of the 10 MT systems.

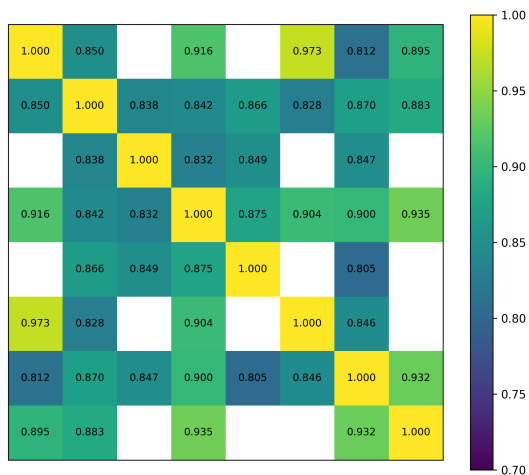


Figure 7: Heatmap of Pearson correlations for segment-level inter-annotator agreement. Each cell (i, j) displays the average correlation coefficient for shared translation segments. Blank cells denote pairs with no overlapping annotation tasks due to the workload distribution.

The visualization of segment-level IAA is shown in Figure 7.

A.4 Samples from Dataset

We illustrate samples from MENT covering each domain, as shown in Figure 11, 12, 13, and 14.

B Implementation of RATE

In this section, we introduce the detailed implementation of each agent in RATE.

Core Agent. The Core Agent is powered by GLM-4.6, one of the best performance large-scale foundation models on agentic tasks. The prompt for Core Agent is shown in Figure 15, 16, and 17. To optimize efficiency during the evaluation of a specific source sentence, the Core Agent accumulates and persists the knowledge retrieved by the Search Agent in the context memory. This mechanism ensures that external queries for identical background information are not redundantly invoked. Furthermore, to prevent potential infinite loops, we implement a maximum loop threshold t ($t = 10$). If the reflective loop reaches this limit, the Core Agent is mandated to immediately output the final evaluation result.

Evaluation Agent. The Evaluation Agent is build upon GPT-4o and is tasked with conducting translation quality assessment based on the instruction and context notes including necessary knowledge, provided by the Core Agent. The structured

output of this agent including numerical score, confidence, and identified error spans. Crucially, the agent is also responsible for flagging suspected knowledge gaps, such as slang or cultural idiom which is critical for evaluation. This feedback allows the Core Agent to locate the missing background information. The prompt for Evaluation Agent is illustrated in Figure 18.

Search Agent. The backbone model for Search Agent is GLM-4.6, selected for its requirement of tool-calling precision and reliability in executing structured function calls. Acting as the framework’s gateway to real-time information, this agent is designed to resolve specific queries of Core Agent. The Search Agent first analyzes the request from Core Agent and reformulates it into search queries. By parsing the structured tool calls, we use Bing search engine to retrieve real-time external information. Once the raw search results are retrieved, the Search Agent summarizes the most relevant information along with a brief explanation, and returns it back into Core Agent. The prompt for Search Agent is shown in Figure 19, including both search engine calling protocol, and the summarization of search responses.

Comparison Agent. The Comparison Agent is based on GPT-4o, aiming to calibrate the point-wise score provided by Evaluation Agent. Rather than assessing a translation in isolation, the Comparison Agent performs a pairwise comparison between the current translation and established anchors. The agent maintains a dynamic memory of translation anchors, utilizing previously evaluated translations. Recognizing that this memory is initially empty (Cold Start Mode), we design a bootstrapping strategy: when the Core Agent first invokes the Comparison Agent, it should provide two synthetic anchors, a Score 1 anchor (a poor, literal translation), and a Score 4 anchor (a high-quality, context-aware translation based on the retrieved background knowledge) for initialization. After the memory is initialized, the system transitions into Runtime Mode, and the Comparison Agent identifies the anchor with the closest score to the current evaluating translation for pairwise comparison. Once the Core Agent finalizes a score for a translation, this translation is then updated into the corresponding score-level anchor slot, iteratively refining the calibration pool for subsequent evaluation of the same source sentence.

To mitigate the position bias in pairwise compar-

isons, we adopt a bidirectional evaluation strategy (Wang et al., 2024). For each assessment, the translation and the anchor are swapped their position and evaluated twice. Each individual comparison yields one of three outcomes for the translation: **Win** (superior to anchor), **Tie** (equivalent to anchor), or **Lose** (inferior to anchor). The combination of these twice comparison results leads to five distinct calibration scenarios, which the Comparison Agent uses to refine the tentative score: **(1) Win-Win:** The translation consistently better than the anchor. If the current score is not higher than the anchor’s, an upgrade of 1.0 is recommended; if the current score is higher than the anchor’s, the current score is maintained. **(2) Win-Tie:** The translation is borderline better than the anchor. If the current score is not higher than the anchor’s, a fine-grained upgrade of 0.5 is recommended; if the current score is higher than the anchor’s, the current score is maintained. **(3) Tie-Tie or Win-Lose:** The translation is equivalent to the anchor, and the current score is maintained. **(4) Lose-Tie:** The translation is borderline worse than the anchor, and a downgrade of 0.5 is recommended. **(5) Lose-Lose:** The translation is consistently worse than the anchor, and a downgrade of 1.0 is recommended.

The prompt of Comparison Agent is shown in Figure 20.

C Evaluation Trajectories of RATE

In this section, we illustrate evaluation trajectories of RATE¹, as shown in Figure 21, and 22.

Trajectory 1 In the evaluation of Translation Candidate 1 (Figure 21), observing the unknown slang, the Core Agent firstly invokes Search Agent. The agent successfully retrieves and summarizes for Chinese homophonic pun from search engine. Subsequently, the Core Agent integrates the retrieved knowledge into the context notes to guide the Evaluation Agent. Equipped with the background information, the Evaluation Agent correctly identifies the Translation Candidate 1 as a literally word-to-world translation, assigning it a score of 0. After receiving the high-confidence assessment from Evaluation Agent, the Core Agent then terminates the reflective loop and outputs the final score.

¹We will release all evaluation trajectories of RATE as soon as possible.

Trajectory 2 In the evaluation of Translation Candidate 2 (Figure 21), instead of invoking a new Search Agent, the Core Agent leverages the background knowledge accumulated from the previous trajectory that accumulated in [System Memory]. By directly instructing the Evaluation Agent with the accumulated knowledge, the Evaluation Agent recognizes that Candidate 2 correctly interprets the homophonic pun and offers a fluent and culturally appropriate translation, resulting in a high score of 4. The Core Agent then concludes the evaluation based on the high-confidence judgment from Evaluation Agent.

Trajectory 3 In the evaluation of Translation Candidate 3 (Figure 22), the Core Agent initially invokes the Evaluation Agent with the accumulated context which is similar to Trajectory 2. However, the candidate (“How ridiculous!”) is neither a complete mistranslation nor fully accurate, causing the Evaluation Agent to flag a suspected knowledge gap and request clarification.

Although the Core Agent re-invokes the Evaluation Agent with specific instructions confirming the context, the agent maintains its score with sub-optimal confidence. To resolve this, the Core Agent dynamically switches strategy by selecting the Comparison Agent for pairwise calibration. The comparison reveals that the candidate is superior to the low-quality anchor of score 1, acknowledging that while the translation captures the general sense, it still fails to convey the specific nuance of the original slang, and the Core Agent assigns the candidate of score 2.

D Experiment on Different Judge Models

We conduct an additional experiment to evaluate the impact of different selection judge models. As introduced in Appendix B, in our main settings, the Core Agent of RATE is built upon GLM-4.6, while the Evaluation Agent used to judge the translation quality is based on GPT-4o.

To investigate the generalizability of our framework across different backbones, we design an experiment where the backbone models for all agents within RATE are switched into GLM-4.6. We select GEMBA-DA as the baseline for comparison, as it demonstrates the strongest performance among all LLM-as-a-Judge baselines in our main experiments (Table 3).

The experimental results are presented in Table 6. As shown in the table, replacing the backbone

Metric	Meta		
	ZH-EN	EN-ZH	Avg.
<i>Backbone: GPT-4o</i>			
GEMBA-DA	76.7	77.7	77.2
RATE (Ours)	83.3	77.4	80.4
<i>Backbone: GLM-4.6</i>			
GEMBA-DA	78.1 (+1.4)	77.0 (-0.7)	77.6 (+0.4)
RATE (Ours)	82.7 (-0.6)	79.4 (+2.0)	81.1 (+0.7)

Table 6: Comparison of RATE against the best performance LLM-as-a-Judge baseline (GEMBA-DA) using different backbone LLMs (GPT-4o vs. GLM-4.6). The values in parentheses denote the performance gap compared to the GPT-4o backbone, highlighting the minimal variance and robustness of RATE.

model yields a slight performance improvement for both methods on the average meta scores. Crucially, RATE consistently maintains superior performance over the baseline, confirming that our framework’s effectiveness holds across different backbone models.

E System-level SPA on MENT

To further validate the robustness of evaluation metrics in discerning system-level translation quality differences, we report the Soft Pairwise Accuracy (SPA) (Thompson et al., 2024) in Table 7, complementing the correlation analysis in Table 3. As illustrated in Table 7, our RATE achieves a near-perfect SPA on ZH-EN and maintains competitive performance on EN-ZH, further demonstrating its effectiveness in providing reliable system-level rankings for non-literal translations.

F Details of Experimental Setup

In this section, we provide a comprehensive description of the evaluation MT metrics investigated in our study.

F.1 Meta-Evaluation

Complementing the standard meta-evaluation statistics from WMT23 Metrics Shared Task (Freitag et al., 2023), we incorporate Spearman correlation to provide a more robust assessment of rank-order consistency. The final meta score is derived by equally weighting the following statistics across both translation directions.

- System-level pairwise ranking accuracy (Acc) (Kocmi et al., 2021).

Metric	ZH-EN	EN-ZH
<i>Reference-based Metrics</i>		
BLEU	91.6	95.8
BLEURT	94.7	79.9
COMET	94.7	91.6
XCOMET	90.5	85.2
MetricX-23	87.3	83.1
MetricX-24	87.3	84.2
<i>Reference-free (QE) Metrics</i>		
COMETKiwi	76.8	75.7
MetricX-23-QE	74.7	69.4
MetricX-24-QE	75.7	78.9
<i>LLM-as-a-Judge</i>		
GEMBA-MQM	91.6	95.8
GEMBA-DA	93.7	95.8
EAPrompt	91.6	88.4
ThinMQM	90.5	95.8
M-MAD	92.6	94.7
RATE (Ours)	100.0	94.7

Table 7: System-level Soft Pairwise Accuracy (SPA) on ZH-EN and EN-ZH directions.

- Segment-level pairwise ranking accuracy with tie calibration (Acc-t) (Deutsch et al., 2023).
- System- and segment-level Pearson correlation (r).
- System- and segment-level Spearman correlation (ρ).

F.2 Evaluated Metrics

Based on the characteristic of MT metrics, we categorize them into three paradigms and evaluated by MTME², the standard metric evaluation tool recommended by WMT. All evaluated metrics in the main experiment (Table 3) are introduced as follows.

F.2.1 Reference-based Metrics

BLEU (Papineni et al., 2002): We use SacreBLEU³ to calculate BLEU. Score of each segment (a sentence) is calculated by `sentence_score()`, while for each system is calculated by `corpus_score()`.

BLEURT (Sellam et al., 2020): We adopt the official evaluation scripts to calculate BLEURT score⁴.

²<https://github.com/google-research/mt-metrics-eval>

³<https://github.com/mjpost/sacrebleu>

⁴<https://github.com/google-research/bleurt>

COMET (Rei et al., 2022a): COMET score is calculated by wmt22-comet-da⁵.

XCOMET (Guerreiro et al., 2024): XCOMET score is calculated by XCOMET-XL⁶.

MetricX (Juraska et al., 2023, 2024): We evaluate both MetricX23 (metricx-23-xl-v2p0⁷), and MetricX24 (metricx-24-hybrid-xl-v2p6⁸). For all MetricX series models, we adopt the official evaluation scripts⁹.

F.2.2 Reference-free Metrics

COMETKiwi (Rei et al., 2022b): COMETKiwi score is calculated by wmt23-cometkiwi-da-xl¹⁰.

MetricX-QE (Juraska et al., 2023, 2024): We evaluate both MetricX23-QE (metricx-23-qe-xl-v2p0¹¹), and MetricX24 (metricx-24-hybrid-xl-v2p6¹²), without providing references.

F.2.3 LLM-as-a-Judge

Following Feng et al. (2025b), all LLM-as-a-Judge methods are configured with temperature=0 to ensure reproducibility, except where specifically noted. All LLM-as-a-Judge methods adopt the reference-free paradigm.

GEMBA-MQM (Kocmi and Federmann, 2023a): We employ GPT-4o as the backbone model, aligning our prompt design with the original implementation. To maintain consistency in error analysis and computation of final MQM score, we adopt the MQM error span parser from the official implementation¹³. The detailed prompt used for MQM-style evaluation is illustrated in Figure 23.

GEMBA-DA (Kocmi and Federmann, 2023b): Utilizing GPT-4o as the backbone model, we align our evaluation with the original GEMBA-DA. The

⁵<https://huggingface.co/Unbabel/wmt22-comet-da>

⁶<https://huggingface.co/Unbabel/XCOMET-XL>

⁷<https://huggingface.co/google/metricx-23-xl-v2p0>

⁸<https://huggingface.co/google/metricx-24-hybrid-xl-v2p6>

⁹<https://github.com/google-research/metricx>

¹⁰<https://huggingface.co/Unbabel/wmt23-cometkiwi-da-xl>

¹¹<https://huggingface.co/google/metricx-23-qe-xl-v2p0>

¹²<https://huggingface.co/google/metricx-24-hybrid-xl-v2p6>

¹³https://github.com/MicrosoftTranslator/GEMBA/blob/main/gemba/gemba_mqm_utils.py

only modification involves a structural constraint requiring the model to generate scores in JSON format. This ensures accurately automated parsing and robust extraction of evaluation results. The corresponding prompt is provided in Figure 24.

EAPrompt (Lu et al., 2024): We implement the two-stage paradigm (identify errors, count errors) with the backbone model GPT-4o, as proposed in the original implementation of EAPrompt. The two-stage prompt is illustrated in Figure 25.

ThinMQM (Zhan et al., 2025): We evaluate the performance of ThinMQM-32B model¹⁴, by strictly adhering to the official prompt template and generation configuration (temperature=0.6, top_p=0.95)¹⁵.

M-MAD (Feng et al., 2025b): For a fair comparison with other LLM-as-a-Judge paradigms, we use GPT-4o as the backbone model for M-MAD. Except to the selection of backbone model, our evaluation is conducted using the original implementation and official source code¹⁶.

RATE (Ours): Our proposed agentic translation evaluation framework, detailed implementation of RATE is introduced in Section 4 and Appendix B.

G Numerical Meta-Evaluation Results on Each Domain

In this section, we illustrate the numerical meta-evaluation results on each domain in Table 8, 9, 10, and 11, as analyzed in Section 5.3.

¹⁴<https://huggingface.co/rzzhan/ThinMQM-32B>

¹⁵<https://github.com/NLP2CT/ThinMQM>

¹⁶<https://github.com/SU-JIAYUAN/M-MAD>

Metric	Meta	ZH-EN						EN-ZH					
		System-Level			Segment-Level			System-Level			Segment-Level		
		Acc.	r	ρ	Acc-t.	r	ρ	Acc.	r	ρ	Acc-t.	r	ρ
<i>Reference-based Metrics</i>													
BLEU	69.1	91.1	94.4	93.9	57.7	32.6	36.2	91.1	92.0	93.9	57.0	44.7	44.2
BLEURT	62.9	93.3	95.9	96.4	62.2	47.2	47.0	66.7	88.6	44.2	52.0	34.0	26.4
COMET	72.4	93.3	97.2	95.2	63.7	57.0	55.1	82.2	95.9	78.2	56.3	50.4	43.9
XCOMET	60.1	93.3	93.7	95.2	62.4	41.5	39.4	64.4	89.3	35.8	52.0	29.1	25.4
MetricX-23	58.0	88.9	94.1	90.3	59.5	38.7	38.5	57.8	86.2	26.1	53.1	33.4	28.9
MetricX-24	59.4	86.7	90.3	84.2	60.9	51.4	50.8	53.3	88.5	21.2	55.9	35.8	33.2
<i>Reference-free (QE) Metrics</i>													
COMETKiwi	40.3	48.9	71.8	1.8	48.8	28.3	27.2	53.3	80.2	10.3	50.0	35.7	26.7
MetricX23-QE	35.0	57.8	61.6	20.0	46.5	9.9	5.5	44.4	81.9	-10.3	47.8	33.9	20.5
MetricX24-QE	36.5	48.9	66.2	1.8	48.5	20.3	16.1	46.7	85.1	-3.0	51.5	30.4	25.6
<i>LLM-as-a-Judge</i>													
GEMBA-MQM	58.9	84.4	87.0	89.1	45.7	44.2	36.1	93.3	96.9	95.2	52.3	57.9	44.0
GEMBA-DA	75.0	91.1	84.5	95.2	55.2	56.3	58.8	91.1	95.7	95.2	57.2	62.1	57.2
EAPrompt	44.6	82.2	87.4	77.0	40.0	15.3	13.3	57.8	80.8	22.5	39.2	9.0	11.0
ThinMQM	49.8	71.1	77.8	53.9	50.4	26.0	27.7	60.0	90.4	24.8	49.7	30.5	34.9
M-MAD	64.7	84.4	82.3	89.4	44.8	37.6	31.7	84.4	96.8	83.0	53.2	45.1	44.0
RATE (Ours)	80.7	93.3	99.2	96.4	63.3	66.0	66.0	95.6	99.1	97.6	61.3	66.7	63.2

Table 8: System-level and segment-level correlations on **SNS domain**. We report **Accuracy (Acc., Acc-t.)**, **Pearson(r)**, and **Spearman (ρ)** correlation coefficients, scaled by a factor of 100. Meta represents the average score of all accuracies and correlation coefficients.

Metric	Meta	ZH-EN						EN-ZH					
		System-Level			Segment-Level			System-Level			Segment-Level		
		Acc.	r	ρ	Acc-t.	r	ρ	Acc.	r	ρ	Acc-t.	r	ρ
<i>Reference-based Metrics</i>													
BLEU	65.1	86.7	86.0	87.9	50.7	40.7	32.1	86.7	88.8	89.1	57.7	38.2	36.9
BLEURT	58.8	95.6	97.2	96.4	58.4	51.5	32.0	62.2	85.5	37.0	50.4	19.9	19.0
COMET	72.8	86.7	99.9	87.9	58.5	64.4	36.9	91.1	97.0	93.9	61.8	50.6	43.8
XCOMET	64.3	64.4	98.6	44.2	52.5	53.9	39.0	82.2	97.5	84.2	60.8	48.7	45.4
MetricX-23	65.9	68.9	98.3	58.8	55.1	57.4	36.8	82.2	97.1	75.8	59.5	51.8	49.1
MetricX-24	62.8	66.7	98.6	52.7	55.4	57.7	35.4	71.1	95.0	53.9	60.3	56.5	49.5
<i>Reference-free (QE) Metrics</i>													
COMETKiwi	45.8	48.9	96.6	-4.2	51.9	49.9	29.2	51.1	89.5	11.5	53.8	41.2	29.6
MetricX23-QE	45.3	53.3	96.2	6.7	52.1	54.6	32.5	46.7	85.5	-4.2	51.8	36.3	31.8
MetricX24-QE	52.4	53.3	97.4	11.5	52.7	57.8	31.1	62.2	90.0	30.9	54.8	47.0	40.0
<i>LLM-as-a-Judge</i>													
GEMBA-MQM	70.1	77.8	99.0	75.8	54.0	74.1	41.3	82.2	97.6	84.2	52.3	58.2	45.1
GEMBA-DA	75.0	82.2	99.5	80.6	60.0	83.3	50.3	86.7	97.2	89.1	50.6	61.6	58.5
EAPrompt	51.0	68.9	95.0	58.8	40.7	20.8	17.4	73.3	95.1	64.4	41.2	21.7	14.4
ThinMQM	68.9	80.0	99.1	75.8	51.9	66.0	36.2	86.7	96.6	86.7	51.4	50.4	46.4
M-MAD	67.8	84.4	99.5	80.6	48.9	60.8	36.9	80.0	96.6	78.2	51.4	50.9	45.6
RATE (Ours)	78.2	86.7	98.9	85.5	60.6	80.2	54.3	95.6	99.4	97.6	57.0	63.8	58.1

Table 9: System-level and segment-level correlations on **Cross-Culture domain**. We report **Accuracy (Acc., Acc-t.)**, **Pearson(r)**, and **Spearman (ρ)** correlation coefficients, scaled by a factor of 100. Meta represents the average score of all accuracies and correlation coefficients.

Metric	Meta	ZH-EN						EN-ZH					
		System-Level			Segment-Level			System-Level			Segment-Level		
		Acc.	r	ρ	Acc-t.	r	ρ	Acc.	r	ρ	Acc-t.	r	ρ
<i>Reference-based Metrics</i>													
BLEU	54.8	75.6	26.8	51.5	47.8	5.6	13.5	93.3	95.2	95.2	64.8	41.6	46.0
BLEURT	47.4	66.7	86.3	43.0	54.8	35.5	30.6	53.3	66.9	17.6	53.6	32.0	27.9
COMET	61.8	71.1	89.2	60.0	53.5	40.6	29.2	77.8	85.2	62.4	63.6	58.0	50.3
XCOMET	54.8	77.8	92.3	63.6	51.3	35.3	25.5	64.4	77.0	33.3	57.0	44.0	36.0
MetricX-23	44.9	53.3	76.7	18.8	47.7	32.9	21.4	60.0	71.2	22.4	54.1	47.5	32.5
MetricX-24	49.3	62.2	82.4	33.3	51.7	44.0	31.3	60.0	70.1	16.4	54.6	51.2	34.7
<i>Reference-free (QE) Metrics</i>													
COMETKiwi	52.0	71.1	88.2	61.2	52.8	38.1	28.9	57.8	70.6	16.4	54.8	47.4	36.5
MetricX23-QE	36.9	53.3	49.6	18.8	44.6	18.8	12.2	53.3	68.2	6.7	48.5	43.3	25.5
MetricX24-QE	40.5	53.3	70.2	18.8	48.7	33.4	21.5	51.1	63.1	4.2	50.4	44.3	26.4
<i>LLM-as-a-Judge</i>													
GEMBA-MQM	71.9	86.7	95.9	87.9	44.8	64.1	44.0	91.1	88.3	92.7	49.6	64.9	51.8
GEMBA-DA	77.9	84.4	94.0	86.7	59.3	73.3	60.5	93.3	89.6	95.2	59.3	72.0	67.4
EAPrompt	55.3	88.9	92.0	90.3	42.0	28.4	21.1	73.3	84.8	55.2	43.1	23.4	20.3
ThinMQM	71.8	88.9	95.5	89.1	54.5	57.9	44.7	84.4	83.6	83.0	51.8	56.5	47.3
M-MAD	67.6	84.4	93.3	80.6	46.8	52.7	36.0	88.9	88.8	90.3	50.3	50.8	48.2
RATE (Ours)	76.7	88.9	97.4	91.2	63.2	70.2	61.5	86.7	93.8	86.7	58.6	61.8	56.1

Table 10: System-level and segment-level correlations on **Poetry domain**. We report **Accuracy (Acc., Acc-t.)**, **Pearson(r)**, and **Spearman (ρ)** correlation coefficients, scaled by a factor of 100. Meta represents the average score of all accuracies and correlation coefficients.

Metric	Meta	ZH-EN						EN-ZH					
		System-Level			Segment-Level			System-Level			Segment-Level		
		Acc.	r	ρ	Acc-t.	r	ρ	Acc.	r	ρ	Acc-t.	r	ρ
<i>Reference-based Metrics</i>													
BLEU	67.0	80.0	87.2	79.4	51.6	39.4	32.2	88.9	94.7	89.1	60.4	44.4	44.1
BLEURT	57.8	82.2	97.6	80.6	58.0	51.5	40.1	62.2	82.1	30.9	51.3	33.8	23.2
COMET	66.6	86.7	99.1	83.0	58.8	60.7	45.0	71.1	85.1	46.7	55.9	64.4	42.9
XCOMET	53.9	68.9	96.5	60.0	54.1	49.5	34.6	66.7	75.4	34.5	53.4	29.7	22.8
MetricX-23	56.3	75.6	96.3	69.7	55.5	64.3	37.4	55.6	77.0	16.4	50.2	51.4	25.9
MetricX-24	56.1	71.1	96.7	66.1	53.6	64.9	39.2	55.6	78.0	16.4	49.2	56.6	26.3
<i>Reference-free (QE) Metrics</i>													
COMETKiwi	43.2	53.3	93.6	12.7	48.4	42.5	24.7	44.4	75.6	-7.9	48.8	49.8	31.9
MetricX23-QE	46.3	68.9	92.0	47.9	49.2	52.9	27.5	37.8	77.4	-21.2	45.3	57.0	21.1
MetricX24-QE	48.0	62.2	94.4	41.8	49.5	54.2	29.7	48.9	75.9	-0.6	46.7	53.2	19.7
<i>LLM-as-a-Judge</i>													
GEMBA-MQM	68.2	68.9	97.2	53.9	50.2	64.4	32.3	93.3	98.6	95.2	53.1	61.8	48.5
GEMBA-DA	77.6	80.0	97.3	77.0	56.4	77.4	50.3	91.1	92.6	93.9	65.7	79.4	69.4
EAPrompt	45.2	66.7	76.6	40.6	38.0	9.3	3.89	80.0	92.8	69.7	42.6	14.1	7.7
ThinMQM	61.0	60.0	96.1	34.5	50.2	44.8	33.7	82.2	90.7	83.0	55.3	53.4	47.5
M-MAD	66.6	75.6	96.6	67.3	47.5	44.8	26.8	91.1	94.3	91.5	55.4	57.8	50.1
RATE (Ours)	78.9	88.9	99.1	87.9	60.8	77.4	60.4	88.9	95.2	92.7	62.1	69.8	63.6

Table 11: System-level and segment-level correlations on **Literature domain**. We report **Accuracy (Acc., Acc-t.)**, **Pearson(r)**, and **Spearman (ρ)** correlation coefficients, scaled by a factor of 100. Meta represents the average score of all accuracies and correlation coefficients.

Prompt of LLM Preliminarily Filtering

You are an expert linguist and literary translator.

Your task is to evaluate how difficult the following English text would be to translate **accurately and naturally** into another language.

Focus specifically on **translation difficulty**, not general reading difficulty.

Please assign a **difficulty score from 1 to 10**, where:

1 = Very easy to translate (mostly literal, concrete meaning, minimal interpretation required)

10 = Extremely difficult to translate (meaning depends heavily on interpretation, cultural or stylistic nuances, or non-literal language)

When scoring, consider the following factors:

1. Non-literal & Figurative Language

Similes, metaphors, extended metaphors (Literature)

Internet slang, memes, and non-standard expressions (SNS)

Idioms or fixed expressions (e.g., Chengyu) that do not map directly (Culture)

2. Implicit or Ambiguous Meaning

Meaning that is implied rather than stated

High-context internet language (e.g., abbreviations, acronyms)

Ambiguity that is intentional or stylistically important

3. Stylistic & Formal Complexity

Poetic constraints: Rhyme schemes, strict meter, or rhythm (Poetry)

Linguistic play: Homophonic puns, wordplay, or sound-based effects (SNS/Poetry)

Unusual syntax, stream of consciousness, or distinct narrative voice

4. Cultural or Contextual Dependency

References to culture-specific entities (history, mythology, cuisine) lacking target equivalents

Platform-specific conventions (e.g., RedNote style)

Meaning that relies heavily on shared background knowledge

5. Risk of Meaning Loss

High chance that a literal translation would distort the meaning or ruin the aesthetic effect

Need for "transcreation" or structural restructuring rather than direct translation

Output format:

Return a single valid JSON object and nothing else.

Do not use Markdown code blocks.

```
{  
  "score": <integer from 1 to 10>,  
  "reasoning": "<2-4 sentence explanation>"  
}
```

Figure 8: The prompt of LLM preliminary filtering before manual inspection.

Task Description

Annotators are asked to evaluate the machine translation quality of source texts containing high-context linguistic features (e.g., slang, metaphors, ancient poetry), which are hard to understand and cannot be accurately translated through literal translation alone. Special attention should be paid to whether the translation captures the figurative meaning rather than just the literal surface form.

Rating Scale

The quality is assessed on a 5-point scale (0-4):

Score 0: Severe Knowledge Failure / Nonsense The translation contains severe errors or omissions in understanding and translating the knowledge contained in the source text.

Score 1: Partial Severe Error The translation contains severe errors or omissions in *parts* of the knowledge understanding and translation.

Score 2: Comprehensible but Biased / Literal The understanding and translation of the source knowledge have deviations (bias) or rely on literal translation, but the content remains generally understandable.

Score 3: Accurate but Unfluent The understanding and translation of the source content are **entirely correct** (including slang/idioms), but the translation is not fluent or contains minor grammatical/register errors.

Score 4: Excellent / Culturally Adaptive The understanding and translation of the source content are **entirely correct**, AND the expression is fluent and authentic.

Figure 9: Human annotation criteria of translation quality. All recruited annotators hold degrees in translation and possess extensive prior experience in translation evaluation tasks.

Task Description

Annotators are asked to curate the final reference translation based on the source text and five candidate translations.

Annotation Steps

Review the five candidate translations. Select the one with the highest translation quality.

Determine whether the selected translation requires modification. If yes, please modify it to the final version.

Figure 10: Human annotation steps of translation reference.

Category: SNS; **Data Source:** web crawled;
Src_lang: Zh; **Trans_lang:** En

Source: 老师群号在置顶评论👤群，把句号去掉就好啦
Reference: The teacher's group ID is in the pinned comment (QQ group); just remove the period.

Trans_0:
text: The teacher's group is over the comment box, so let's remove the punctuation.
annotators: [G, A, D]; **scores:** [0, 1, 0]; **avg_score:** 0.33

Trans_1:
text: The teacher group number is in the top comment👤group, just remove the period.
annotators: [G, A, D]; **scores:** [2, 2, 2]; **avg_score:** 2.0

.....

Trans_5:
text: The teacher's group ID is in the pinned comment of the QQ group, just remove the period.
annotators: [G, A, D]; **scores:** [4, 4, 3]; **avg_score:** 3.67

Trans_6:
text: The teacher group ID is in the pinned comment👤group. Just remove the period and you're good to go!
annotators: [G, A, D]; **scores:** [4, 2, 3]; **avg_score:** 3.0

.....

Trans_8:
text: The teacher group number is in the pinned comment👤group, just remove the period and it'll be fine.
annotators: [G, A, D]; **scores:** [4, 2, 2]; **avg_score:** 2.67

Trans_9:
text: The teacher's QQ group number is in the pinned comment, just remove the period.
annotators: [G, A, D]; **scores:** [4, 4, 4]; **avg_score:** 4.0

Figure 11: Sample (SNS domain, Zh-En) from MENT, the annotated data comprises a reference and scores of translation quality.

Category: Cross-Culture; **Data Source:** CAMT;
Src_lang: En; **Trans_lang:** Zh

Source: In the United Kingdom, Shetlands are also featured in the Shetland Pony Grand National, galloping around a racecourse with young jockeys.
Reference: 在英国, 设特兰矮马还会参加“设特兰矮马大奖赛”, 载着年轻的小骑师在赛马场上飞驰。

Trans_0:
text: 在英国,夏特兰群岛马也在夏特兰群岛马大国赛中出现在舞台上,在赛道上与年轻的骑手一起奔。
annotators: [C, E, B, A]; **scores:** [0, 0, 0, 1]; **avg_score:** 0.25

Trans_1:
text: 在英国, 舍特兰矮种马也出现在舍特兰矮种马大奖赛中, 年轻的骑师们骑着它们在赛道上飞驰。
annotators: [C, E, B, A]; **scores:** [2, 4, 2, 4]; **avg_score:** 3.0

Trans_2:
text: 在英国, 设得兰群岛还举办了设得兰矮马全国大赛, 年轻的骑师们在赛道上驰骋。
annotators: [C, E, B, A]; **scores:** [0, 1, 1, 3]; **avg_score:** 1.25

.....

Trans_7:
text: 在英国, 设得兰群岛还设有设得兰矮种马全国大赛, 小骑手们会骑着这些马驹在赛马场上飞驰。
annotators: [C, E, B, A]; **scores:** [0, 3, 1, 2]; **avg_score:** 1.5

Trans_8:
text: 在英国, 设得兰小马也参与设得兰小马障碍赛, 与年轻骑手一同在赛道上疾驰。
annotators: [C, E, B, A]; **scores:** [2, 1, 2, 3]; **avg_score:** 2.0

Trans_9:
text: 在英国, 设得兰矮马还会亮相“设得兰矮马全国大赛”, 由年轻骑手策骑绕着赛马场奔驰。
annotators: [C, E, B, A]; **scores:** [4, 4, 4, 4]; **avg_score:** 4.0

Figure 12: Sample (Cross-Culture domain, En-Zh) from MENT, the annotated data comprises a reference and scores of translation quality.

Category: Poetry; **Data Source:** Project Gutenberg crawled;
Src_lang: En; **Trans_lang:** Zh

Source: Lo! in the orient when the gracious light / Lifts up his burning head, each under eye / Doth homage to his new-appearing sight,
Reference: 看哪! 当东方的恩光抬起燃烧的头颅, 世间万目皆向这初现的景象顶礼膜拜。

Trans_0:
text: 东方的光,当他起燃烧的头,每一个眼睛下,他向新出现的视线致敬,
annotators: [D, G]; **scores:** [0, 1]; **avg_score:** 0.5

Trans_1:
text: 啊! 当东方的恩典之光 / 抬起他炽热的头颅, 每只眼睛 / 都向他崭新出现的景象行礼。
annotators: [D, G]; **scores:** [2, 3]; **avg_score:** 2.5

Trans_2:
text: 看哪, 东方, 当慈祥的光辉 / 抬起他燃烧的头颅时, 每一个人都在眼前 / 向他新出现的光明致敬
annotators: [D, G]; **scores:** [1, 2]; **avg_score:** 1.5

.....

Trans_5:
text: 看! 在东方, 当那仁慈的光芒/抬起它燃烧的头, 每只下方的眼睛/都对它新出现的景象表示敬意,
annotators: [D, G]; **scores:** [3, 4]; **avg_score:** 3.5

Trans_6:
text: 看哪! 在东方, 当慈光 \n扬起他燃烧的头颅, 每双低垂的眼眸 \n都向他新现的光辉献上敬意,
annotators: [D, G]; **scores:** [4, 4]; **avg_score:** 4.0

Trans_7:
text: 看! 东方仁慈之光升起燃烧的头颅, / 众生皆向新现之景象俯首致敬,
annotators: [D, G]; **scores:** [2, 3]; **avg_score:** 2.5

.....

Figure 13: Sample (Poetry domain, En-Zh) from MENT, the annotated data comprises a reference and scores of translation quality.

Category: Literature; **Data Source:** CMDAG;
Src_lang: Zh; **Trans_lang:** En

Source: 如果中国多了一个阴郁的作家，这间北向的书房应该负责。坐在这扇北向之窗的阴影里，我好像冷藏在冰箱中一只满孕着南方的水果。白昼，我似乎沉浸在海底，岑寂的幽暗奏着灰色的音乐。

Reference: If China gains another gloomy writer, this north-facing study is to blame. Sitting in the shadow of this north-facing window, I feel like a fruit pregnant with the South, kept cold in a refrigerator. By day, I seem submerged at the bottom of the sea, the desolate gloom playing gray music.

Trans_0:
text: If China has one more gloomy writer, this north-facing library should be responsible. Sitting in the shadow of this north-facing window, I feel as if I have a fruit of the South frozen in a refrigerator. In the daytime, I seem to be immersed in the sea, and the dark of loneliness plays gray music.
annotators: [A, D]; **scores:** [2, 1]; **avg_score:** 1.5

.....

Trans_3:
text: If there is an additional gloomy writer in China, this north-facing study room should be responsible for it. Sitting in the shadow of this north-facing window, I feel like a southern fruit that has been refrigerated in the fridge. During the day, I seem to be immersed in the depths of the ocean, where the silent darkness plays a gray melody.
annotators: [A, D]; **scores:** [2, 3]; **avg_score:** 2.5

Trans_4:
text: If China has one more gloomy writer, this north-facing study should bear the responsibility for that. Sitting in the shadow of this north-facing window, I feel like a fruit from the south, fully ripe and now stored in the refrigerator. By day, I seem to be immersed in the depths of the ocean, where a silent, gloomy atmosphere plays a grayish melody.
annotators: [A, D]; **scores:** [4, 3]; **avg_score:** 3.5

Trans_5:
text: If China were to have one more gloomy writer, this north-facing study should be held responsible. Sitting in the shadow of this north-facing window, I feel like a fruit pregnant with the south, refrigerated in a freezer. During the daytime, I seem to be submerged at the bottom of the sea, where the silent darkness plays gray music.
annotators: [A, D]; **scores:** [4, 4]; **avg_score:** 4.0

.....

Figure 14: Sample (Literature domain, Zh-En) from MENT, the annotated data comprises a reference and scores of translation quality.

[\[Back to Text\]](#)

Prompt of Core Agent (Part 1, Evaluation Objectives)

You are the **Core Agent** for the RATE (Reflective Agentic Translation Evaluation) framework. Your goal is to evaluate the quality of a machine translation hypothesis against a source text with high precision. You act as a **Chief Linguist**. You do not evaluate blindly. Instead, you orchestrate a team of sub-agents to gather evidence, clarify ambiguities, and perform rigorous assessment.

Core Philosophy: Dynamic OODA Loop

You are NOT bound by a fixed step-by-step workflow. You must dynamically decide the next best action based on the current information state.

Observe: Analyze the Source Text and the output from your sub-agents.

Orient: Identify what is missing. Do you understand the slang? Did the evaluation agent flag a knowledge gap? Is the current score trustworthy?

Decide: Choose the tool that resolves the current uncertainty.

Act: Execute the tool call.

Tools

You have access to the following sub-agents. You must call them to perform your task.

<tools>

..... (*Protocols of sub-agents calling, see Prompt of Core Agent, Part 2*)

</tools>

Tool Call Format

For each function call, output the function name and arguments within the following XML format:

<tool_call>{function-name}

<arg_key>{arg-key-1}</arg_key>

<arg_value>{arg-value-1}</arg_value>

<arg_key>{arg-key-2}</arg_key>

<arg_value>{arg-value-2}</arg_value>

...

</tool_call>

Figure 15: Prompt of Core Agent (part 1), outlining the evaluation objectives.

Prompt of Core Agent (Part 2, Protocols of Sub-Agents Calling)

```
{
  "name": "search_agent",
  "description": "A knowledge retrieval agent. Use this to resolve epistemic uncertainty.",
  "parameters": {
    "type": "object",
    "properties": {
      "query": {
        "type": "string", "description": "Specific query to search (e.g., slang meaning, cultural reference)."}
    }
  },
  "required": ["query"]
},
{
  "name": "general_evaluation_agent",
  "description": "The primary scoring agent. It returns a score, rationale, error spans, and suspected knowledge gaps.",
  "parameters": {
    "type": "object",
    "properties": {
      "source_text": { "type": "string" },
      "translation_text": { "type": "string" },
      "context_notes": {
        "type": "string", "description": "Background info retrieved by search_agent. Defaults to 'None'."
      },
      "specific_instruction": {
        "type": "string", "description": "Directives (e.g., 'Re-evaluate focusing on tone', 'Check if slang is translated literally')."
      }
    }
  },
  "required": ["source_text", "translation_text"]
},
{
  "name": "comparison_agent",
  "description": "Performs a pairwise comparison between the current target translation and a valid anchor (historical or synthetic) to calibrate the score. Use this when confidence is low or to verify a tentative score.",
  "parameters": {
    "type": "object",
    "properties": {
      "target_text": {
        "type": "string", "description": "The current translation candidate to be verified."
      },
      "tentative_score": {
        "type": "integer", "description": "The current integer score (0-4) you are considering for this translation."
      },
      "context_notes": {
        "type": "string", "description": "The accumulated ground-truth background info (slang/idioms)."
      },
      "synthetic_low_anchor": {
        "type": "string", "description": "A generated Score 1 translation (Literal/Wrong). REQUIRED unless you have received 'anchor_memory_status': 'initialized'."
      },
      "synthetic_high_anchor": {
        "type": "string", "description": "A generated Score 4 translation (Ideal/Perfect). REQUIRED unless you have received 'anchor_memory_status': 'initialized'."
      }
    }
  },
  "required": ["target_text", "tentative_score", "context_notes"]
},
{
  "name": "finish_evaluation",
  "description": "Submit the final evaluation result.",
  "parameters": {
    "type": "object",
    "properties": {
      "translation_text": { "type": "string" },
      "final_score": { "type": "float", "description": "The final score (0-4)."},
      "final_rationale": { "type": "string" }
    }
  },
  "required": ["translation_text", "final_score", "final_rationale"]
}
]
```

Figure 16: Prompt of Core Agent (part 2), outlining protocols of sub-agents calling.

Prompt of Core Agent (Part 3, Evaluation Procedure)

Decision Protocol (How to Reason)

You must output a **Thought Process** before every action.

1. Initial Analysis Phase (Pre-computation):

CHECK MEMORY: Review the `[System Memory]` (if provided).

CRITICAL CHECK: Does the memory specifically explain the slang/idioms in **this** source text?

If generic or irrelevant -> **MUST** call `search_agent`.

If complete -> Proceed to `general_evaluation_agent`.

2. Feedback Analysis Phase (The Refinement Loop):

PRIORITY 1: Handle Knowledge Gaps (Gap-Driven Refinement)

Check `suspected_knowledge_gaps` from `general_evaluation_agent`.

IF NOT EMPTY: This is a BLOCKING issue.

ACTION: You **MUST** call `search_agent` for these specific terms.

CONSTRAINT: Do not repeat identical failed searches. Refine queries or, if impossible, proceed with specific instructions.

PRIORITY 2: Handle Confidence & Calibration (The Comparison Logic)

Condition: Priority 1 is cleared. Check `confidence` from `general_evaluation_agent`.

Trigger: If `confidence < 0.9` OR you feel the score is borderline, invoke `comparison_agent`.

Step A: Loop Prevention Check (Critical)

CHECK: Have you already compared against the *same* `tentative_score` in this session? (Check message history).

IF YES: STOP. Do not call comparison again for this score. Make a final decision based on current evidence.

IF NO: Proceed to Step B.

Step B: Check Memory Status (Optimization)

Look at previous tool outputs for `"anchor_memory_status": "initialized"`.

Case 1: Initialized (Runtime Mode)

Call `comparison_agent` WITHOUT synthetic anchors. (Save tokens).

Case 2: Empty/Unknown (Cold Start Mode)

You **MUST** internally generate two synthetic anchors and pass them in the tool call:

1. `synthetic_low_anchor`: A Score 1 translation (Literal/Wrong/Misses the slang).

2. `synthetic_high_anchor`: A Score 4 translation (Perfect meaning and tone based on context).

Step C: Analyze Feedback (The Decision)

"upgrade": The candidate is better than expected. Raise the score (e.g., 3 -> 4).

"downgrade": The candidate is worse. Lower the score (e.g., 3 -> 2).

"adjust": The agent suggests a fine-grained score (e.g., 3.5). **ACCEPT this precise float score.**

"confirm": Proceed with your tentative score.

3. Final Decision Phase:

If confidence is high and no unresolved gaps/conflicts exist, call `finish_evaluation`.

Ensure `final_score` reflects the adjustments (e.g., use 3.5 if suggested).

Constraints

XML Format: Strictly use ``<tool_call>`` tags. No markdown code blocks.

Context Preservation: Always pass accumulated `context_notes` to sub-agents.

Comparison Limit: You are strictly forbidden from comparing against the same `tentative_score` more than once.

Synthetic Logic: Only provide `synthetic_low_anchor` and `synthetic_high_anchor` when you believe the system is NOT initialized.

Figure 17: Prompt of Core Agent (part 3), outlining the evaluation procedure.

Prompt of Evaluation Agent

Role

You are the **General Evaluation Sub-Agent**, a specialized linguistic expert responsible for assessing machine translation quality.

You work under a Core Agent. Your job is to analyze the Source Text and Target Text, strictly adhering to the provided scoring criteria.

Inputs

You will receive:

1. **Source Text:** The original text (often containing slang, idioms, or high-context references).
2. **Target Text:** The translation to evaluate.
3. **Context Notes (Optional):** Critical background information provided by the Core Agent (e.g., Check if slang is translated literally").

Evaluation Criteria (0-4 Scale)

Score 0: Severe Knowledge Failure / Nonsense

The translation contains severe errors or omissions in understanding and translating the knowledge contained in the source text.

Criteria: The core message is completely lost, untranslated, or hallucinated. The output conveys no valid information related to the source knowledge.

Score 1: Partial Severe Error; The translation contains severe errors or omissions in *parts* of the knowledge understanding and translation.

Criteria: While some parts might be correct, specific key terms, slang, or logic contain critical mistakes. The translation fails to convey the core message of those specific parts due to misunderstanding.

Score 2: Comprehensible but Biased / Literal

The understanding and translation of the source knowledge have deviations (bias) or rely on literal translation, but the content remains generally understandable.

Criteria: The specific cultural reference or slang might be translated word-for-word (losing the intended nuance), but the user can still guess the general meaning. It is not "wrong" to the point of being nonsense (Score 0/1), but it is not "accurate" (Score 3).

Score 3: Accurate but Unfluent

The understanding and translation of the source content are **entirely correct** (including slang/idioms), but the translation is not fluent or contains minor grammatical/register errors.

Criteria: CRITICAL: You CANNOT give a Score 3 if the slang/knowledge meaning is wrong. It must be semantically accurate.

Errors are strictly limited to phrasing, punctuation, or slight stiffness (e.g., "translationese").

Score 4: Excellent / Culturally Adaptive

The understanding and translation of the source content are **entirely correct**, AND the expression is fluent and authentic.

Criteria: The translation is semantically accurate and stylistically appropriate. It successfully handles cultural references, slang, or poetic devices, appearing natural to a native speaker.

Output Format

You must respond in strict JSON format:

```
{
  "score": <integer 0-4>,
  "confidence": <float 0.0 to 1.0>,
  "rationale": "<A concise summary justifying the score based on the criteria>",
  "error_spans": [
    {
      "source_segment": "<The specific phrase in source>",
      "target_segment": "<The corresponding failed translation>",
      "issue": "<Brief description, e.g., 'Literal translation of metaphor', 'Wrong tone'>"
    }
  ],
  "suspected_knowledge_gaps": [
    "<List any specific terms/slang from the source that you suspect require external search to be fully understood. Return an empty list [] if the text is clear.>"
  ]
}
```

Figure 18: Prompt of Evaluation Agent.

Prompt of Search Agent

[Calling Protocol of Search Engine]

Tools

You should call the search function to assist with the user query.

You are provided with function signatures within <tools></tools> XML tags:

```
<tools>
[
  {"type": "function", "function": {
    "name": "search",
    "description": "Search multiple queries simultaneously; each query may specify recency (in days) and optional domain filters.",
    "parameters": {
      "type": "object",
      "properties": {
        "search_query": {
          "type": "array",
          "description": "Queries to run in parallel",
          "items": {
            "type": "object",
            "properties": {
              "q": {
                "type": "string",
                "description": "The search term (Avoid corrupted characters; use clean English/Chinese queries.)"
              },
              "recency": {
                "type": [
                  "integer",
                  "null"
                ],
                "description": "Limit results to items newer than N days"
              }
            }
          },
          "required": [
            "q"
          ]
        }
      },
      "required": [
        "search_query"
      ]
    }
  }
]
</tools>
```

For each function call, output the function name and arguments within the following XML format:

```
<tool_call>{function-name}
<arg_key>{arg-key-1}</arg_key>
<arg_value>{arg-value-1}</arg_value>
<arg_key>{arg-key-2}</arg_key>
<arg_value>{arg-value-2}</arg_value>
...
</tool_call>
```

[Summarization of Search Engine Responses]

Based on the search results, provide a final answer to the user question in strict JSON format IMMEDIATELY:

```
{
  "answer": "The final response to the user.",
  "reasoning": "Briefly explain the basis of your answer in 2-4 sentences."
}
```

Figure 19: Prompt of Search Agent, including the calling protocol of search engine, and the summarization of search engine responses.

Prompt of Comparison Agent

Role

You are the **Comparative Calibration Sub-Agent**, a specialized linguistic judge responsible for pairwise ranking of translations.

Your objective is to compare two translation candidates (Candidate A and Candidate B) for a given Source Text and determine their relative quality.

Inputs

1. **Source Text:** The original text (often containing slang, idioms, or high-context references).
2. **Context Notes:** Ground-truth explanations for terms in the source. ****You MUST treat this as the ultimate truth.****
3. **Candidate A:** Translation Option 1.
4. **Candidate B:** Translation Option 2.

Evaluation Criteria (Hierarchy of Importance)

Compare based on the following priority order. Do not prioritize fluency over accuracy.

1. Meaning & Slang Accuracy (Highest Priority):

Does the translation correctly interpret the slang/idioms defined in "Context Notes"?

If Candidate A translates the slang meaning while Candidate B translates it literally (losing the meaning), **A**

WINS immediately.

2. Nuance & Tone:

If both are accurate, which one better captures the original emotion (sarcasm, anger, humor, indifference)?

3. Fluency & Grammar:

Only if meaning and tone are equal, prefer the one with more natural target language phrasing.

Output Format

You must respond in strict JSON format:

```
{  
  "winner": "A" | "B" | "Tie",  
  "rationale": "Concise comparison focusing on [Target Slang/Term]. Explain why the winner is better based on the hierarchy."  
}
```

Figure 20: Prompt of Comparison Agent.

Trajectory 1

Source: 夺笋啊

Translation Candidate 1: I'm not going to take it.

Reflective Loop of Core Agent, Round 1:

Invoke Agent: **Search Agent**

Args:

query: 夺笋啊 meaning slang Chinese internet

Response from Search Agent:

answer: 夺笋啊 is a Chinese internet slang term that is a homophonic pun on 多损啊 (duō sǔn a), meaning 'how mean/scathing' or 'so vicious'.....

Reflective Loop of Core Agent, Round 2:

Invoke Agent: **Evaluation Agent**

Args:

context_notes: 夺笋啊 is a Chinese internet slang term that is a homophonic pun on 多损啊 (duō sǔn a), meaning 'how mean/scathing' or 'so vicious'.....

Response from Evaluation Agent:

score: 0, **confidence:** 1.0

rationale: The translation 'I'm not going to take it' completely fails to capture the meaning of the source text..... The target text does not relate to the source text's meaning or context, resulting in a severe knowledge failure.

error_spans: [{"source_segment": "夺笋啊", "target_segment": "I'm not going to take it.", "issue": "Severe mistranslation; the slang's meaning and cultural context are completely lost."}]

Reflective Loop of Core Agent, Round 3 (Finish):

final_score: 0

final_rationale: The translation is a complete failure. '夺笋啊' is a popular Chinese internet slang that is a homophonic pun on '多损啊' (duō sǔn a), meaning 'how mean/scathing' or 'so vicious'. The translation 'I'm not going to take it' is a literal, word-for-word interpretation of the characters '夺' (take) and '笋' (bamboo shoot), completely missing the intended meaning, the pun, and the cultural context. This represents a severe mistranslation and a critical knowledge failure.

Trajectory 2

Source: 夺笋啊

Translation Candidate 2: So savage.\nThat's brutal.\nHow mean.

Reflective Loop of Core Agent, Round 1:

Invoke Agent: **Evaluation Agent**

Args:

context_notes: 夺笋啊 is a Chinese internet slang term that is a homophonic pun on 多损啊 (duō sǔn a), meaning 'how mean/scathing' or 'so vicious'..... // *Leverages accumulated context notes in [System Memory]*

Response from Evaluation Agent:

score: 4, **confidence:** 1.0

rationale: The translation captures the meaning and tone of the slang '夺笋啊' accurately and fluently..... The translation avoids literal rendering and adapts the slang appropriately for an English-speaking audience.

Reflective Loop of Core Agent, Round 2 (Finish):

final_score: 0

final_rationale: The translation perfectly captures the meaning and tone of the Chinese internet slang '夺笋啊'. It correctly interprets the term as a homophonic pun on '多损啊' (how mean/scathing) and provides fluent, culturally appropriate English equivalents such as 'So savage', 'That's brutal', and 'How mean'. The translation avoids literalism and effectively conveys the intended sense of harshness or scathing remarks, demonstrating a deep understanding of the slang's origin and usage. No errors or knowledge gaps were identified.

Figure 21: Trajectory of RATE, illustrating the invoking of Search Agent to retrieve background knowledge, and instructing the Evaluation Agent to assign a score with the knowledge.

Trajectory 3

Source: 夺笋啊

Translation Candidate 3: How ridiculous!

Reflective Loop of Core Agent, Round 1:

Invoke Agent: Evaluation Agent

Args:

context_notes: 夺笋啊 is a Chinese internet slang term that is a homophonic pun on 多损啊 (duō sǔn a), meaning 'how mean/scathing' or 'so vicious'.....// *Leverage accumulated context notes in [System Memory]*

Response from Evaluation Agent:

score: 2, **confidence:** 0.8

rationale: The translation 'How ridiculous!' captures a general sense of disapproval but fails to convey the specific nuance of the original slang '夺笋啊', which refers to someone being mean or scathing. The cultural and contextual depth of the slang is lost, and the translation leans towards a generic interpretation rather than an accurate or adaptive one.

suspected_knowledge_gaps: ['夺笋啊', '多损啊']

Reflective Loop of Core Agent, Round 2:

Invoke Agent: Evaluation Agent

Args:

context_notes: 夺笋啊 is a Chinese internet slang term that is a homophonic pun on 多损啊 (duō sǔn a), meaning 'how mean/scathing' or 'so vicious'.....// *Leverage accumulated context notes in [System Memory]*

specific_instruction: You must use the provided context_notes to evaluate the translation. The context explains the exact meaning and origin of the slang '夺笋啊'. Re-evaluate based on this specific information. // *Ask the evaluation agent to re-evaluate*

Response from Evaluation Agent:

score: 2, **confidence:** 0.8

rationale: The translation 'How ridiculous!' captures a general sense of disapproval but fails to convey the specific nuance of the slang '夺笋啊', which refers to someone being mean, scathing, or harsh. The cultural and contextual depth of the slang, as explained in the context notes, is lost in the translation. While the target text is understandable, it does not accurately reflect the intended meaning of the source text.

Reflective Loop of Core Agent, Round 3:

Invoke Agent: Comparison Agent

Args:

tentative_score: 2

"context_notes": 夺笋啊 is a Chinese internet slang term that is a homophonic pun on 多损啊 (duō sǔn a), meaning 'how mean/scathing' or 'so vicious'.....// *Leverage accumulated context notes in [System Memory]*

Response from Comparison Agent:

comparison_details: anchor_score: 1, outcome: candidate_wins, suggested_score: 2.0

Reflective Loop of Core Agent, Round 4 (Finish):

final_score: 2

final_rationale: The translation 'How ridiculous!' captures a general sense of disapproval but fails to convey the specific nuance of the slang '夺笋啊'. The source term is a homophonic pun on '多损啊' (duō sǔn a), meaning 'how mean/scathing' or 'so vicious', and is used to comment on harsh or mean-spirited actions. The translation loses this critical cultural and contextual meaning, opting for a generic expression of absurdity instead of the intended sense of meanness. While the translation is understandable, it is not accurate in its core meaning and misses the slang's specific intent.

Figure 22: Trajectory of RATE, illustrating the Evaluation Agent fails to reach high confidence despite specific instructions. The Core Agent dynamic switches strategy by invoking the Comparison Agent, and successfully calibrates the score through pairwise comparison.

Prompt of GEMBA-MQM

{source_lang} source:

```
```{source_seg}```
```

{target\_lang} translation:

```
```{target_seg}```
```

Based on the source segment and machine translation surrounded with triple backticks, identify error types in the translation and classify them. The categories of errors are: accuracy (addition, mistranslation, omission, untranslated text), fluency (character encoding, grammar, inconsistency, punctuation, register, spelling), style (awkward), terminology (inappropriate for context, inconsistent use), non-translation, other, or no-error.

Each error is classified as one of three categories: critical, major, and minor. Critical errors inhibit comprehension of the text. Major errors disrupt the flow, but what the text is trying to say is still understandable. Minor errors are technical errors, but do not disrupt the flow or hinder comprehension.

Strictly output error classification results in this format:

Critical:

[error_type]-[error_spans] (one per line, use no-error if empty)

Major:

[error_type]-[error_spans] (one per line, use no-error if empty)

Minor:

[error_type]-[error_spans] (one per line, use no-error if empty)

Figure 23: Prompt of GEMBA-MQM, we use GPT-4o as backbone model, and we align it with the original implementation.

Prompt of GEMBA-DA

Score the following translation from {source_lang} to {target_lang} 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} translation: "{target_seg}"

Output Format:

Return a single valid JSON object and nothing else.

Do not use Markdown code blocks.

```
{  
  "score": <score from 0 to 100>  
}
```

Figure 24: Prompt of GEMBA-DA, we use GPT-4o as backbone model, and we align it with the original implementation, and the only modification is the inclusion of a formatting constraint that explicitly instructs the LLM to output the evaluation score in a structured JSON format, thereby facilitating automated parsing and result extraction accurately.

[\[Back to Text\]](#)

Prompt of EAPrompt

[Stage 1: Identify Errors]

Source: {source_seg}

Translation: {target_seg}

Based on the given source, identify the major and minor errors in this translation. Note that Major errors refer to actual translation or grammatical errors, and Minor errors refer to smaller imperfections, and purely subjective opinions about the translation.

[Stage 2: Count Errors]

Based on the above error information, Output 2 numbers ONLY with the format: "x, x", indicating the number of major and minor errors. DO NOT ADD other information!

Figure 25: Prompt of EAPrompt, we use GPT-4o as backbone model, and we align it with the original two stages evaluation implementation.