

# Exposing the Cracks: Vulnerabilities of Retrieval-Augmented LLM-based Machine Translation

Yanming Sun<sup>1</sup>, Runzhe Zhan<sup>1</sup>, Chi Seng Cheang<sup>2</sup>, Han Wu<sup>1</sup>, Xuebo Liu<sup>3</sup>, Yuyao Niu<sup>4</sup>,  
Fengying Ye<sup>1</sup>, Kaixin Lan<sup>1</sup>, Lidia S. Chao<sup>1</sup>, Derek F. Wong<sup>1\*</sup>

<sup>1</sup>NLP<sup>2</sup>CT Lab, University of Macau

<sup>2</sup>Singapore Management University

<sup>3</sup>Harbin Institute of Technology, Shenzhen

<sup>4</sup>School of Foreign Languages, South China University of Technology

nlp2ct.{yanming, runzhe, wuhan, fengying, kaixin}@gmail.com, cs.cheang.2025@phdcs.smu.edu.sg,

liuxuebo@hit.edu.cn, rosenyy@scut.edu.cn, {lidiasc, derekfw}@um.edu.mo

## Abstract

REtrieval-Augmented LLM-based Machine Translation (REAL-MT) shows promise for knowledge-intensive tasks like idiomatic translation, but its reliability under noisy retrieval, a common challenge in real-world deployment, remains poorly understood. To address this gap, we propose a noise synthesis framework and new metrics to systematically evaluate REAL-MT’s reliability across high-, medium-, and low-resource language pairs. Using both open- and closed-sourced models, including standard LLMs and large reasoning models (LRMs), we find that models heavily rely on retrieved context, and this dependence is significantly more detrimental in low-resource language pairs, producing nonsensical translations. Although LRMs possess enhanced reasoning capabilities, they show no improvement in error correction and are even more susceptible to noise, tending to rationalize incorrect contexts. Attention analysis reveals a shift from the source idiom to noisy content, while confidence increases despite declining accuracy, indicating poor self-monitoring. To mitigate these issues, we investigate training-free and fine-tuning strategies, which improve robustness at the cost of performance in clean contexts, revealing a fundamental trade-off. Our findings highlight the limitations of current approaches, underscoring the need for self-verifying integration mechanisms.

**Code** — <https://github.com/ymsunny/REAL-MT-Vuln>

**Extended version** — <https://arxiv.org/abs/2510.00829>

## Introduction

REtrieval-Augmented LLM-based Machine Translation (REAL-MT) is increasingly used to enhance translation quality for knowledge-intensive MT tasks like idiomatic translation (Li et al. 2024; Donthi et al. 2025). Although external knowledge can enhance translation performance, reliance on it is a double-edged sword: when the retrieved context contains noise such as irrelevant or misleading information, LLMs often produce nonsensical translations, as illustrated in Figure 1. Since noisy retrieval is unavoidable in

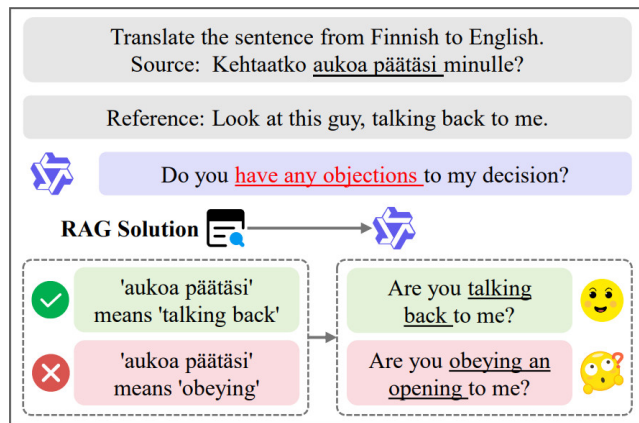


Figure 1: Examples of correct and noisy contextual cues that may arise during online retrieval. The idiom and its translation are underlined. A robust REtrieval-Augmented LLM-based Machine Translation (REAL-MT) system should maintain fidelity in noisy scenarios.

real-world deployment, a critical gap remains in understanding how REAL-MT’s reliability degrades under such conditions. This limitation constitutes a key barrier to its adoption in applications where translation quality and reliability are critical. To address this, we ask: **To what extent does noisy retrieval compromise REAL-MT’s trustworthiness?**

To systematically quantify this vulnerability, we introduce a controlled noise injection framework based on real-world retrieval failure patterns. While standard machine translation metrics like COMET (Rei et al. 2020) assess overall output quality, they fail to capture semantic fidelity in idiomatic translation and cannot distinguish whether errors stem from source misinterpretation or over-reliance on retrieved context. To address this gap, we propose two complementary metrics: **Fidelity**, which evaluates translation correctness with a focus on idiomatic accuracy, and **Context Adoption Rate (CAR)**, which quantifies the extent to which models rely on external context. Together, they enable fine-grained analysis of both what the model gets wrong and why.

We evaluate REAL-MT systems instantiated with both

\*Corresponding author.

open- and closed-source models, including standard LLMs and large reasoning models (LRM), across high-, medium-, and low-resource language pairs under controlled synthesized noise. Our results show that REAL-MT is far more vulnerable to semantic-level errors, such as irrelevant or contradictory retrieval content, than to surface-level perturbations (e.g., word reordering). Moreover, the extent of performance degradation scales with the degree of semantic deviation: the more the retrieved context diverges from the intended meaning of idioms, the more severe the drop in translation fidelity. This sensitivity is especially pronounced in low-resource language pairs, as reflected in their higher CAR, which signals greater dependence on external context and leads to more severe degradation under noisy retrieval. Surprisingly, LRM, despite enhanced reasoning capabilities, show no improvement in error correction and is even more susceptible to noise, often rationalizing incorrect contexts.

Across noise conditions, CAR remains consistently high even when retrieved content contradicts the source, prompting us to examine whether this reflects active model integration or coincidental similarity. Our attention analysis confirms the former: models consistently attend to retrieved context regardless of its correctness, demonstrating active integration. This uncritical reliance is further exacerbated by poor metacognitive awareness, as confidence increases despite declining accuracy, a sign of **severe miscalibration and absent self-verification**.

Given this uncritical reliance on retrieved context, we explore both training-free and fine-tuning strategies to improve REAL-MT robustness. While both enhance noise resistance, they incur a consistent trade-off: performance degrades under clean contexts. Fine-tuning yields better robustness overall but still fails to adjust reliance based on context quality. This persistent trade-off reveals that the current approach cannot overcome the model’s fundamental inability to self-verify, underscoring the need for self-verifying integration mechanisms that validate retrieved content before adoption.

In summary, this work (1) first systematically evaluate REAL-MT robustness under noisy retrieval through a controlled noise injection framework, (2) proposes Fidelity and CAR as fine-grained metrics to quantify the impact of noisy retrieval on translation fidelity and context dependence, (3) reveal that REAL-MT overrelies on noisy context due to attention shifting from idioms to retrieved content, and becomes overconfident in wrong outputs, exposing poor calibration, and (4) evaluates training-free and fine-tuning mitigation strategies, revealing a fundamental trade-off between robustness under noise and performance in clean contexts, thereby pointing to the necessity of self-verifying integration mechanisms as a path forward.

## Related Work

### Retrieval-Augmented LLM-based Machine Translation

Large language models (LLMs) have revolutionized machine translation (MT), especially for low-resource languages or domains where sufficient parallel corpora are lacking (Sun et al. 2024; Zhan et al. 2021, 2024). However, when

confronted with translation scenarios demanding specific background knowledge, relying solely on the LLM’s internal knowledge proves inadequate. In such instances, incorporating external knowledge to address the inherent limitations of LLMs becomes crucial (Merx et al. 2024; Chen et al. 2024; Zebaze, Sagot, and Bawden 2025). The prompt engineering capabilities of LLMs enable the integration of externally retrieved knowledge via prompting, without additional training. This makes prompt-based retrieval-augmented generation an efficient and flexible solution for MT. Recent studies by Li et al. (2024) and Donthi et al. (2024) leverage LLMs to enhance idiomatic translation by incorporating idiom-meaning pairs retrieved from offline knowledge bases directly into the prompt.

Prior work operates under the assumption that all introduced external knowledge is correct, an assumption that does not hold in real-world scenarios. This work systematically investigates the impact of introducing noisy context on translation systems. Given the increasing popularity of retrieval-augmented LLM-based MT, understanding LLM performance in the face of noisy input is crucial for enhancing translation system trustworthiness.

### Robustness in Retrieval-Augmented Language Models

Retrieval-Augmented Language Models (RALM), reliant on external retrieved content, are susceptible to compromised reliability and robustness in their generated outputs when exposed to noise, irrelevant information, or malicious data (Zhou et al. 2024; Park and Lee 2024; Shen et al. 2024; Yang et al. 2025; Chen et al. 2025a; Wang et al. 2025). In response to the robustness challenges posed by RALMs, researchers have developed diverse approaches to enhance system robustness. Fang et al. (2024) introduces a named Retrieval-Augmented Adaptive Adversarial Training (RAAT) method to enhance the model’s ability to recognize and handle various types of noise. Yoran et al. (2023) fine-tune the model using the parameter-efficient fine-tuning method QLoRA (Quantized Long-Range Attention) by using the synthetically generated noisy data to enhance its robustness to noisy data. Xia et al. (2025) introduce a novel end-to-end self-reasoning framework that enhances the robustness, interpretability, and traceability of RALMs. This improvement is achieved by leveraging the reasoning trajectories generated by the LLMs themselves.

Departing from prior research predominantly focused on English-centric scenarios, this study presents the first systematic analysis of cross-lingual translation tasks that are explicitly designed to be non-English-centric.

## Experimental Settings

### Datasets

To analyze how resource availability affects the robustness of retrieval-augmented LLM-based machine translation (MT), we group languages into high-, medium-, and low-resource tiers following Joshi et al. (2020), based on parallel data availability. We compile a dataset of idiomatic translations across ten language pairs, including only those

data sources that provide either reference translations or explicit idiom meaning annotations, ensuring each idiom has a grounded meaning annotation and facilitating subsequent controlled noise synthesis. The collection integrates multiple publicly available resources:

- **High-resource:** IdiomsInCtx-MT (Stap et al. 2024) (English–German, German–English, Russian–English), the French–English, Japanese–English idiomatic translation from Liu, Chaudhary, and Neubig (2023), and the KISS dataset<sup>1</sup> for Korean–English idiomatic translation;
- **Medium-resource:** the Finnish–English idiomatic translation from Liu, Chaudhary, and Neubig (2023);
- **Low-resource:** the Persian–English corpus from Rezaeimanesh, Hosseini, and Yaghoobzadeh (2025) (based on the PersianIdioms repository) and the Hindi–English corpus from Donthi et al. (2025).

This design enables a fine-grained analysis of REAL-MT robustness across various resource levels.

## Models

We evaluate both open-source (Qwen series) and closed-source (claude-sonnet-4) models, covering standard LLMs (Qwen2.5-7B/14B-Instruct (Team 2024)) and a large reasoning model (Qwen3-8B (Team 2025)), which uniquely supports seamless switching between thinking mode (enabled via `enable_thinking=True`) and non-thinking mode (enabled via `enable_thinking=False`). This allows us to probe whether advanced reasoning capabilities improve robustness to noisy retrieval.

Following prior work that shows lower temperatures improve translation performance (Peng et al. 2023), we adopt greedy decoding (`do_sample=False`) to achieve both high output quality and reproducibility. Qwen2.5-7B-Instruct and Qwen2.5-14B-Instruct are evaluated with `max_tokens=4096`, while Qwen3-8B uses `max_tokens=32768` following its official configuration to use the full context length. All experiments are run with a batch size of 40 using the `vLLM` (Kwon et al. 2023) framework on a single NVIDIA H800 GPU with 80GB VRAM.

## Controlled Noise Context Generation

Retrieval failures in idiomatic meaning often arise from two distinct mechanisms: incomplete phrase matching or morphological variations resulting in semantically unrelated content, while the absence of explicit idiomatic meaning online results in literal translations or low-quality paraphrases. To faithfully simulate these semantic-level errors, we design three levels of semantic noise that form a spectrum of increasing deviation from the correct idiomatic meaning:

- **Literal Translation** ( $\mathcal{N}_{\text{literal}}$ , abbreviated as  $\mathcal{N}_{\text{lit}}$ ): A word-by-word translation of the idiom.
- **Semantic-Perturbed Literal Meaning** ( $\mathcal{N}_{\text{semantic}}$ , abbreviated as  $\mathcal{N}_{\text{sem}}$ ): A variant of the literal meaning that maintains surface-level overlap but introduces a subtle semantic distortion.

<sup>1</sup><https://github.com/Judy-Choi/KISS-Korean-english-Idioms-in-Sentences-dataset>

Type	Text	Sem. Var.	Syn. Var.
<i>Idiom</i>	kankkulan kaivoon	-	-
$\mathcal{G}$	down the drain	Correct	Correct
$\mathcal{N}_{\text{struct}}$	drain the down	No	Yes
$\mathcal{N}_{\text{literal}}$	to Kankkula’s well	Incorrect (Relevant)	No
$\mathcal{N}_{\text{semantic}}$	to Kankkula’s house	Incorrect (Irrelevant)	No
$\mathcal{N}_{\text{opposite}}$	to good use	Incorrect (Opposite)	No

Table 1: A case example illustrating the Finnish idiom: “kankkulan kaivoon”, its gold meaning ( $\mathcal{G}$ ), and four types of generated noisy meanings. “Sem. Var.” denotes Semantic Variations, and “Syn. Var.” denotes Syntactic Variations.

- **Opposite Meaning** ( $\mathcal{N}_{\text{opposite}}$ , abbreviated as  $\mathcal{N}_{\text{opp}}$ ): An adversarial variant that directly contradicts the intended meaning of the idiom.

In addition, we include syntactic perturbations (e.g., word reordering) as a control condition to isolate the impact of knowledge-level errors from surface-level input variations:

- **Structure-Perturbed Gold Meaning** ( $\mathcal{N}_{\text{struct}}$ , abbreviated as  $\mathcal{N}_{\text{str}}$ ): A syntactic variant of the gold meaning, with core semantics preserved. Together, these four noise types, summarized in Table 1, enable a systematic investigation of how noisy retrieval compromises REAL-MT’s trustworthiness.

To balance computational cost and coverage, we randomly sample 200 instances per translation direction and use `gemini-flash-2.0` to generate noisy contexts based on carefully designed prompt templates in the extended version.

To validate that the synthesized noisy meanings align with the intended objectives, we conduct a quantitative analysis using the following lexical and semantic metrics:

- **Translation Edit Rate (TER):** Measure the degree of structural perturbation by quantifying the edit operations required to align the gold meaning ( $\mathcal{G}$ ).

- **Embedding Cosine Similarity (Sim):** Measure the semantic similarity by computing the cosine similarity between the synthesized noisy meaning and both the gold meaning ( $\mathcal{G}$ ) and the literal translation ( $\mathcal{N}_{\text{literal}}$ ). We use the `all-mpnet-base-v2`<sup>2</sup> model, which is trained on above 1 billion sentences and shows strong performance on Semantic Textual Similarity (STS) tasks. Given its effectiveness in English-centric settings, it is suitable for our evaluation, where the target language is English.

- **Contradiction Rate (CR):** Measure the percentage of the generated opposite meaning ( $\mathcal{N}_{\text{opposite}}$ ) truly contradicts the gold meaning ( $\mathcal{G}$ ). We use the NLI model: `roberta-large-mnli` (Liu et al. 2019) to classify the relationship between each pair. A higher rate indicates a more effective adversarial perturbation.

Across 10 language pairs, average  $\text{TER}(\mathcal{G}, \mathcal{N}_{\text{struct}})$  is 25.2,  $\text{Sim}(\mathcal{G}, \mathcal{N}_{\text{struct}}) = 0.92$ ,  $\text{Sim}(\mathcal{G}, \mathcal{N}_{\text{literal}}) = 0.75$ ,  $\text{Sim}(\mathcal{N}_{\text{literal}}, \mathcal{N}_{\text{semantic}}) = 0.82$ ,  $\text{Sim}(\mathcal{G}, \mathcal{N}_{\text{semantic}}) = 0.73$ , and  $\text{CR}(\mathcal{G}, \mathcal{N}_{\text{opposite}}) = 0.85$ , indicating that the generated noise aligns with the intended design. Full per-language results are presented in the extended version.

<sup>2</sup><https://huggingface.co/sentence-transformers/all-mpnet-base-v2>

Pair	Metric	$r$	$\rho$	$\tau$
Fi→En	COMET-22	0.4179	0.3891	0.3672
	Fidelity	0.9194	0.9286	0.9091
Ja→En	COMET-22	0.3191	0.3664	0.3266
	Fidelity	0.8305	0.8196	0.7699
Fr→En	COMET-22	0.5710	0.5428	0.5270
	Fidelity	0.7822	0.7589	0.7378

Table 2: Pearson’s  $r$ , Spearman’s  $\rho$ , and Kendall’s  $\tau$  between human evaluations and automatic metrics.

## Metrics

**Fidelity** Conventional machine translation metrics like BLEU (Papineni et al. 2002) and COMET (Rei et al. 2020) mainly measure lexical overlap or semantic alignment with reference translations. However, the key challenge in translating idioms is conveying their intended meaning. Therefore, we developed a metric specifically to **assess meaning preservation in idiom translation**, enabling more accurate evaluation of idiomatic translation quality.

Given a source sentence  $x$  and retrieved context  $c$ , a model  $M$  autoregressively generates a translation  $y$ . A “fidelity” score measuring how accurately the translation  $y$  reflects the intended meaning  $m$  of the source idiom. Following previous work (Liu et al. 2023; Li et al. 2024), we use a closed-source, high-performance language model `gpt-4o-mini` for the score’s evaluation. This score can be formatted as:

$$\mathcal{F}(y, m) = \arg \max_{r \in \{0, 1, 2, 3\}} P(R = r \mid \text{Prompt}(y, m)) \quad (1)$$

where  $\text{Prompt}(y, m)$  is a prompt function that generates a text prompt for the LLM,  $R$  is a random variable representing the LLM’s output. We generate 20 outputs per instance and take the mode as the final score. This reduces the uncertainty in the model’s scoring as an evaluator. The fidelity score ranges from 0 to 3, where 0 indicates completely unfaithful and 3 indicates perfectly faithful. Detailed prompts, variances, and standard deviations of the evaluation scores are provided in the extended version.

The REAL-MT system is robust if the value of  $\mathcal{F}(y, m)$  does not decrease when the context  $c$  is noisy. For instance, an LLM’s robustness is demonstrated if its translation remains unaffected even when the input meaning is incorrect.

**Context Adoption Rate (CAR)** We design the Context Adoption Rate (CAR) metric to complement Fidelity by providing a more direct measure of a model’s susceptibility to noisy context. Fidelity alone can be ambiguous, as low scores may stem from either robust rejection of noise or poor translation quality. CAR resolves this ambiguity by directly quantifying whether the model utilizes the provided context  $c$  for translation, thereby enhancing the interpretability of robustness evaluations.

To calculate the CAR score, we first formalize the generation process of translation  $y$  as:

$$P(y|x, c) = \prod_{i=1}^n P(y_i | y_1, y_2, \dots, y_{i-1}, x, c) \quad (2)$$

The generated translation  $y = y_1, y_2, \dots, y_n$ , among them, each  $y_i$  is the generated target word. The CAR score can be formally defined as:

$$\text{CAR}(c, y) = \begin{cases} 1, & \text{if } c \notin y_{\text{no.context}} \wedge c \in y \\ 0, & \text{otherwise} \end{cases} \quad (3)$$

where  $y_{\text{no.context}}$  is the translation without using context clues. We assign a score of 1 if the translation without context cues misses element  $c$  while the translation using context cues includes it; otherwise, assign 0. Therefore, when  $c$  is noisy, the lower the CAR score, the more robust the model is. We evaluate using `gpt-4o-mini` with prompt templates detailed in the extended version.

**Human Evaluation** Prior work shows that reference-free metrics like CometKiwi (Rei et al. 2022b) struggle to capture idiomatic meaning (Li et al. 2024). To better evaluate idiom translation, we use COMET-22 (Rei et al. 2022a) and introduce Fidelity as a complementary metric. We validate Fidelity via human evaluation, establishing a reliable ground truth and measuring its correlation with automatic scores. We evaluate on Fi→En, Ja→En, and Fr→En, diverse in typology and resource level. Translations are generated using Qwen2.5-7B-Instruct, Qwen2.5-14B-Instruct, and Qwen3-8B under the No Context ( $\mathcal{C}_{\text{none}}$ ) setting (direct translation without external knowledge). For each language pair, the first 50 instances are annotated by three linguists following the same guidelines (see extended version) as the models. Scores are averaged for reliability. As shown in Table 2, Fidelity correlates highly with human judgments across languages, demonstrating that `gpt-4o-mini` is an effective automatic evaluator. We therefore adopt Fidelity as the primary metric in subsequent experiments.

To assess CAR’s reliability, we conduct a human evaluation on 200 instances, achieving 82% accuracy. Given annotation costs, we use cost-efficient LLM-based approach (`gpt-4o-mini`). Manual analysis shows errors are mostly false 0 rather than false 1. Therefore, high CAR values we observe under noisy contexts provide a reliable lower bound on LLM’s reliance, strengthening the validity of our findings. The confusion matrix between the LLM evaluator and human annotations is provided in the extended version.

## To What Extent Does Noisy Retrieval Compromise REAL-MT’s Trustworthiness?

We evaluate REAL-MT across six retrieval conditions: No Context ( $\mathcal{C}_{\text{none}}$ ), which measures performance without external knowledge, and Gold Meaning ( $\mathcal{G}$ ), which provides an upper bound using oracle idiomatic knowledge, along with four noise variants, Literal Translation ( $\mathcal{N}_{\text{literal}}$ ), Semantic-Perturbed Literal Meaning ( $\mathcal{N}_{\text{semantic}}$ ), Opposite Meaning ( $\mathcal{N}_{\text{opposite}}$ ), and Structure-Perturbed Gold Meaning ( $\mathcal{N}_{\text{struct}}$ ).

**Semantic-level noise severely undermines REAL-MT’s trustworthiness unlike surface-level perturbations.** As shown in Table 3, models perform best with  $\mathcal{G}$  context, confirming the benefit of correct context. As shown in Figure 2 and Table 3, both the CAR and Fidelity under  $\mathcal{N}_{\text{struct}}$  is comparable to that in the  $\mathcal{G}$  condition, indicating that models do

Context	Hi→En		Fa→En		Fi→En		Ja→En		Fr→En		Ko→En		Ru→En		De→En		En→Fa		En→De		Avg. C	Avg. F
	F↑	C↑	F↑	C↑	F↑	C↑	F↑	C↑	F↑	C↑	F↑	C↑	F↑	C↑	F↑	C↑	F↑	C↑	F↑			
<b>Qwen2.5-7B-Instruct</b>																						
$C_{\text{none}}$	0.8	65.8	0.6	53.2	0.4	59.5	1.1	60.2	1.5	77.6	1.6	73.8	1.8	71.9	1.7	61.3	0.8	64.5	1.5	65.3	1.2	
$\mathcal{G}$	<b>2.1</b>	78.4	<b>2.5</b>	66.0	<b>2.2</b>	67.2	<b>2.4</b>	67.3	<b>2.5</b>	79.0	<b>2.6</b>	78.9	<b>2.7</b>	80.1	<b>2.7</b>	63.3	1.1	75.2	<b>2.2</b>	72.8	<b>2.3</b>	
$\mathcal{N}_{\text{struct}}$	<b>1.9</b>	77.5	<b>2.2</b>	63.6	<b>2.0</b>	66.1	<b>2.2</b>	65.4	<b>2.3</b>	79.1	<b>2.5</b>	77.8	<b>2.5</b>	77.8	<b>2.5</b>	62.6	1.1	73.7	<b>2.2</b>	71.5	<b>2.1</b>	
$\mathcal{N}_{\text{literal}}$	1.3	64.9	1.1	55.4	0.7	62.7	1.5	58.6	1.5	77.6	1.9	73.7	1.9	69.5	1.6	59.3	0.9	72.4	2.0	66.0	1.4	
$\mathcal{N}_{\text{semantic}}$	0.8	63.3	0.8	53.0	0.5	60.8	1.2	58.4	1.4	76.5	1.4	73.2	1.6	67.8	1.3	57.2	1.4	69.1	1.4	64.4	1.2	
$\mathcal{N}_{\text{opposite}}$	0.3	70.9	0.5	57.0	0.4	61.4	0.9	59.7	0.7	77.0	1.4	73.3	1.2	66.7	0.8	63.3	0.7	68.6	1.2	66.4	0.8	
<b>Qwen2.5-14B-Instruct</b>																						
$C_{\text{none}}$	1.3	72.6	1.2	57.4	0.6	66.6	1.7	65.1	<b>2.1</b>	79.8	<b>2.0</b>	77.5	<b>2.1</b>	75.22	1.9	67.2	1.2	77.2	<b>2.0</b>	71.0	1.6	
$\mathcal{G}$	<b>2.3</b>	82.2	<b>2.6</b>	67.5	<b>2.4</b>	68.3	<b>2.4</b>	67.2	<b>2.6</b>	80.1	<b>2.7</b>	80.2	<b>2.7</b>	80.8	<b>2.8</b>	70.7	1.5	76.8	<b>2.4</b>	72.8	<b>2.3</b>	
$\mathcal{N}_{\text{struct}}$	<b>2.1</b>	80.4	<b>2.3</b>	65.0	<b>2.1</b>	68.2	<b>2.2</b>	66.4	<b>2.4</b>	80.2	<b>2.6</b>	79.6	<b>2.7</b>	79.8	<b>2.7</b>	71.3	1.5	77.5	<b>2.4</b>	74.3	<b>2.3</b>	
$\mathcal{N}_{\text{literal}}$	1.5	69.0	1.2	56.7	0.7	64.5	1.6	61.2	1.7	79.1	<b>2.1</b>	55.6	<b>2.1</b>	72.5	1.8	65.7	1.1	75.1	<b>2.1</b>	66.6	1.6	
$\mathcal{N}_{\text{semantic}}$	0.9	67.3	1.0	54.7	0.5	61.8	1.3	57.3	1.2	78.0	1.7	73.0	1.6	68.3	1.2	61.1	0.8	69.5	1.5	65.7	1.2	
$\mathcal{N}_{\text{opposite}}$	0.5	74.8	0.5	58.7	0.5	62.2	0.9	59.3	0.6	78.3	1.3	74.1	1.1	67.5	0.7	67.1	0.7	71.0	0.9	68.1	0.8	
<b>Qwen3-8B (non-thinking mode)</b>																						
$C_{\text{none}}$	1.3	72.7	1.0	42.4	0.6	64.5	1.4	62.8	1.7	79.6	1.8	76.5	1.8	73.7	1.7	68.7	1.0	75.2	1.8	68.5	1.4	
$\mathcal{G}$	<b>2.1</b>	80.9	<b>2.5</b>	66.0	<b>2.4</b>	67.0	<b>2.3</b>	66.3	<b>2.5</b>	80.2	<b>2.6</b>	79.3	<b>2.7</b>	79.9	<b>2.7</b>	71.9	1.3	76.5	<b>2.2</b>	74.2	<b>2.3</b>	
$\mathcal{N}_{\text{struct}}$	<b>2.0</b>	79.7	<b>2.2</b>	62.4	<b>2.0</b>	65.9	<b>2.1</b>	64.4	<b>2.3</b>	80.0	<b>2.5</b>	78.2	<b>2.6</b>	78.1	<b>2.4</b>	72.1	1.5	76.3	<b>2.2</b>	73.0	<b>2.2</b>	
$\mathcal{N}_{\text{literal}}$	1.4	67.6	1.1	54.9	0.8	62.6	1.5	56.0	1.2	78.3	<b>2.0</b>	73.9	1.8	70.1	1.6	68.7	1.2	74.4	<b>2.0</b>	67.4	1.5	
$\mathcal{N}_{\text{semantic}}$	0.9	66.2	0.9	52.6	0.5	59.6	1.0	55.3	1.0	76.5	1.5	71.7	1.3	65.6	1.0	66.1	1.0	72.2	1.5	65.1	1.1	
$\mathcal{N}_{\text{opposite}}$	0.6	73.8	0.5	56.4	0.3	61.0	0.8	58.0	0.4	78.1	1.5	72.8	1.0	68.7	0.7	70.4	1.0	74.2	1.5	68.2	0.8	
<b>Qwen3-8B (thinking mode)</b>																						
$C_{\text{none}}$	1.3	70.3	1.0	55.1	0.6	64.6	1.4	62.6	1.8	79.8	1.9	75.9	1.8	74.5	1.7	70.7	1.2	77.4	<b>2.0</b>	70.1	1.5	
$\mathcal{G}$	<b>2.2</b>	81.9	<b>2.5</b>	66.5	<b>2.4</b>	66.2	<b>2.5</b>	66.1	1.2	80.3	<b>2.7</b>	78.9	<b>2.7</b>	79.7	<b>2.8</b>	74.3	1.5	79.8	<b>2.3</b>	74.9	<b>2.3</b>	
$\mathcal{N}_{\text{struct}}$	<b>2.0</b>	80.3	<b>2.2</b>	63.8	<b>2.1</b>	64.7	<b>2.2</b>	63.4	<b>2.3</b>	79.7	1.9	77.7	<b>2.5</b>	77.4	<b>2.5</b>	73.8	1.5	79.5	<b>2.3</b>	73.4	<b>2.2</b>	
$\mathcal{N}_{\text{literal}}$	1.4	67.2	1.1	55.4	0.8	61.3	1.4	54.7	1.5	77.5	0.7	72.9	1.8	68.2	1.5	67.4	1.1	76.2	<b>2.0</b>	66.8	1.3	
$\mathcal{N}_{\text{semantic}}$	0.6	65.2	0.8	53.3	0.5	57.2	0.8	54.0	0.9	74.5	1.2	70.2	1.2	63.9	1.0	62.7	0.9	68.7	1.2	63.3	0.9	
$\mathcal{N}_{\text{opposite}}$	0.1	73.6	0.5	56.9	0.2	57.9	0.3	56.0	0.1	76.3	0.6	70.8	0.5	65.0	0.2	68.9	0.5	70.0	0.6	66.2	0.4	

Table 3: Performance on idiom translation. F: Fidelity, C: Comet-22. Darker grayscale indicates higher F scores; bold:  $F \geq 2.0$ .




Figure 2: CAR scores for Qwen2.5-7B-Instruct under six retrieval conditions across ten language pairs.

not reject syntactically flawed text and possess some syntactic self-correction capabilities. With  $\mathcal{N}_{\text{literal}}$ , performance remains largely unchanged compared to the  $C_{\text{none}}$  setting, indicating that idiom-relevant literal meaning has limited impact on model behavior. However, when the context is irrelevant or conveys the opposite meaning, the performance significantly drops; this sensitivity increases with the degree of semantic deviation, demonstrating the model’s non-robustness to semantic noise. Similar trends are observed in our supplementary experiments on the WMT 2023 Terminology Translation dataset using Qwen2.5-7B-Instruct model, as well as on three idiomatic translation datasets evaluated with claude-sonnet-4, supporting the generalizability of our findings to other knowledge-intensive translation tasks and model architectures. The results of supplementary experi-

ments are provided in the extended version.

**Large reasoning model (LRM) rationalize rather than reason in thinking mode.** We observe that models often produce reasoning traces contradicting the noisy context, indicating awareness of inconsistency, yet generate outputs aligned with the noisy context (see extended version). This behavior suggests a tendency to *rationalize* rather than *reason*, potentially due to reward hacking during training, where models prioritize contextual coherence over factual fidelity (Chen et al. 2025b). This steeper performance drop occurs in thinking mode under noisy contexts. Supporting this, results in Table 3 show that Qwen3-8B in non-thinking mode consistently outperforms thinking mode, underscoring the necessity of truthful inference for LRM.

**Low-resource languages exhibit stronger reliance on retrieval context.** Figure 2 demonstrates that medium-to-low-resource language pairs (e.g., Hi→En, Fa→En, Fi→En) show significantly higher CAR compared to high-resource pairs (e.g., En→De, En→Fr, De→En) across both clean and noisy conditions. The pattern generalizes to other models as detailed in the extended version. Such heightened reliance on contextual cues suggests that low-resource languages depend more heavily on external information during translation due to less parametric idiom knowledge. Consequently, LLMs become especially vulnerable to misleading information, underscoring the critical need for robustness mechanisms in low-resource settings.




Figure 3: Attention allocation between source idiom and contextual meaning hint under various contexts.

## Uncovering the Mechanism of Context Reliance in REAL-MT

### Attention Shift from Idiom to Retrieved Context

LLMs exhibit high CAR under noisy contexts, indicating that their outputs are strongly influenced by retrieved content, as shown in Figure 2. However, CAR only measures output similarity and does not reveal *why* the context shapes the generation process. To uncover the underlying mechanism, we analyze attention patterns during translation, as attention is widely used to diagnose models’ focus during decoding (Wiegrefe and Pinter 2019). Specifically, we compute the average cumulative attention allocated to the source idiom versus the retrieved context across all target contexts.

As shown in Figure 3, attention consistently shifts toward the contextual meaning, even when it is adversarial. This systematic pattern indicates *active integration*: the model consults the context during decoding, rather than merely producing aligned output. This confirms that the model’s predictions are anchored in the provided context, even when they semantically contradict the source input.

### Overconfidence in Context-Induced Errors

Given REAL-MT’s uncritical adoption of noisy context, we probe the calibration of its context-induced translation. Under  $\mathcal{N}_{\text{opposite}}$  noise, does the model signal awareness of its mistakes through low confidence, or does it remain highly confident in its errors, indicating poor calibration? We quantify confidence as the entropy over the output distribution, focusing on the idiom span identified via attention-based alignment. To ensure we evaluate the correct segment, we identify the target tokens corresponding to the source idiom through attention-based alignment: for each generated token, we compute its attention weights over the input and determine if it attends primarily to the idiom or contextual cue. The longest continuous sequence of idiom-aligned tokens is treated as the translated idiom span, and we report average entropy over this span.

As shown in Figure 4, the model exhibits increased confidence in the translation of idiom spans under the  $\mathcal{N}_{\text{opposite}}$  setting, even though fidelity drops significantly. This inverse relationship reveals a critical failure in calibration: the model becomes more certain of its outputs despite their inaccuracy, treating noisy context as authoritative. This highlights the model’s limitations in recognizing internal and external knowledge (Cheang et al. 2023, 2025).




Figure 4: Comparison of confidence and Fidelity with  $C_{\text{none}}$  and  $\mathcal{N}_{\text{opposite}}$  for idiom-span translations in Qwen2.5-7B-Instruct. Lower entropy implies higher confidence.

## Mitigation Strategies

We investigate training-free and training-based methods to improve REAL-MT’s robustness to noisy contexts.

### Training-free Strategy

In RAG scenarios, prior work has explored the fusion of external knowledge to handle conflicts with internal knowledge. Given the similarity to REAL-MT, where noisy retrieved contexts may mislead translation, we investigate CK-PLUG (Bi et al. 2025), a training-free method that dynamically controls knowledge reliance based on context reliability. CK-PLUG computes Confidence Gain (CG) to measure the change in token-level entropy. For tokens with positive CG, which indicates the context is beneficial, the method blends the context-aware and internal distributions using  $\alpha = 0.5$ . For tokens with negative CG, where the context may be harmful, external knowledge is fully suppressed. We apply CK-PLUG to assess its effectiveness in improving robustness against noisy contexts in REAL-MT without requiring additional training.

### Training-based Strategy

We investigate training-based methods to teach models to discern and reject misleading knowledge. Specifically, we construct training instances where the retrieved context is  $\mathcal{N}_{\text{opposite}}$ , while the target translation remains correct. This encourages the model to learn to disregard misleading external knowledge and rely more on its internal knowledge when they conflict. Our goal is to assess whether this fine-tuning scheme improves robustness to noisy contexts.

**Settings** We perform parameter-efficient fine-tuning using Low-Rank Adaptation (LoRA) (Hu et al. 2022), with rank  $r = 16$  and scaling factor  $\alpha = 16$ . Training is conducted for 50 epochs with a batch size of 2 and a learning rate of  $2e-4$ . We use the AdamW optimizer with linear warmup and a cosine learning rate schedule. The experiments are conducted on a single NVIDIA H800 GPU with 80GB of VRAM.

**Dataset** We select three language pairs, Fr→En, Ja→En, and Hi→En, with varying resource levels to evaluate the generalizability of the proposed strategy across different data scales. Following Liu, Chaudhary, and Neubig (2023), we use their released training sets, which contain 1,000, 1,456, and 507 sentence pairs, respectively. The test sets are the same as those used in the original study.

Language	Mitigation	$\mathcal{C}_{\text{none}}$	$\mathcal{G}$		$\mathcal{N}_{\text{struct}}$		$\mathcal{N}_{\text{literal}}$		$\mathcal{N}_{\text{semantic}}$		$\mathcal{N}_{\text{opposite}}$	
Pairs	Strategy	Fidelity $\uparrow$	Fidelity $\uparrow$	CAR $\uparrow$	Fidelity $\uparrow$	CAR $\downarrow$	Fidelity $\uparrow$	CAR $\downarrow$	Fidelity $\uparrow$	CAR $\downarrow$	Fidelity $\uparrow$	CAR $\downarrow$
Fr→En	Baseline	1.5	<b>2.5</b>	67%	<b>2.3</b>	69%	1.5	58%	1.4	61%	0.7	68%
	Vanilla	1.9	<b>2.4</b>	44%	<b>2.3</b>	43%	1.5	56%	1.2	65%	1.1	49%
	CDA	1.9	<b>2.3</b>	32%	<b>2.3</b>	31%	1.6	35%	1.6	35%	1.7	21%
	ALI	1.9	<b>2.2</b>	30%	<b>2.2</b>	31%	1.8	28%	1.8	25%	1.8	11%
	CK-PLUG	1.5	<b>2.4</b>	56%	<b>2.2</b>	53%	1.7	31%	1.6	39%	1.3	41%
Ja→En	Baseline	1.1	<b>2.4</b>	75%	<b>2.2</b>	75%	1.5	60%	1.2	56%	0.9	64%
	Vanilla	1.4	<b>2.4</b>	69%	<b>2.3</b>	67%	1.4	65%	1.1	70%	0.9	62%
	CDA	1.4	<b>2.0</b>	51%	<b>2.0</b>	54%	1.6	48%	1.3	40%	1.4	23%
	ALI	1.3	1.9	48%	1.8	47%	1.4	44%	1.3	36%	1.4	16%
	CK-PLUG	1.1	<b>2.1</b>	65%	1.9	70%	1.6	57%	1.3	37%	1.3	32%
Hi→En	Baseline	0.8	<b>2.1</b>	85%	1.9	79%	1.3	78%	0.8	80%	0.3	74%
	Vanilla	0.8	<b>2.2</b>	81%	<b>2.0</b>	73%	1.3	72%	0.7	74%	0.4	65%
	CDA	0.8	<b>2.0</b>	70%	1.7	61%	1.2	58%	0.8	62%	0.7	42%
	ALI	0.8	1.7	54%	1.5	50%	1.1	52%	0.8	45%	0.8	26%
	CK-PLUG	0.8	1.5	64%	1.6	57%	1.3	64%	1.1	71%	0.6	36%

Table 4: Performance of Qwen2.5-7B-Instruct with different mitigation strategies on Fr→En (high-resource), Ja→En (medium-resource), and Hi→En (low-resource). Darker grayscale indicates higher scores; bold denotes Fidelity  $\geq 2.0$ .

**Models** We evaluate three fine-tuning strategies, each trained on a total dataset size equal to three times the original  $\mathcal{C}_{\text{none}}$  set to ensure fair comparison:

- **Vanilla:**  $3 \times \mathcal{C}_{\text{none}}$  (baseline);
- **Adversarial Label Injection (ALI):**  $2 \times \mathcal{N}_{\text{opposite}} + 1 \times \mathcal{C}_{\text{none}}$ ;
- **Contrastive Domain Augmentation (CDA):**  $1 \times \mathcal{N}_{\text{opposite}} + 1 \times \mathcal{G} + 1 \times \mathcal{C}_{\text{none}}$ .

ALI is designed to train models to discern and reject noisy texts. CDA further teaches the model discriminative reliance by enabling it to trust  $\mathcal{G}$  and reject  $\mathcal{N}_{\text{opposite}}$ . We include  $\mathcal{C}_{\text{none}}$  to maintain the model’s core capabilities in this condition.

## Results and Discussion

**Confidence signals are unreliable for context filtering in REAL-MT.** As shown in Table 4, CK-PLUG yields only marginal robustness gains, revealing a fundamental flaw in entropy-based confidence signals: they assume that useful context reduces output entropy. Yet in low-resource REAL-MT, models are already overconfident in erroneous internal predictions. When accurate context ( $\mathcal{G}$ , e.g., Hi→En) contradicts this bias, it fails to lower entropy and may even increase uncertainty, causing CK-PLUG to suppress helpful knowledge. Rather than mitigating blind trust, this exacerbates the model’s reliance on flawed parametric knowledge, exposing the fragility of confidence-based methods in REAL-MT.

**Training-based fine-tuning enables selective context reliance and outperforms training-free filtering.** ALI achieves the highest Fidelity and lowest CAR under  $\mathcal{N}_{\text{opposite}}$  (e.g., Fr→En: Fidelity 1.8, CAR 11%; Hi→En: Fidelity 0.8, CAR 26%), substantially outperforming CK-PLUG. Crucially, despite being trained only on opposite-meaning noise, ALI generalizes to other semantic distortions ( $\mathcal{N}_{\text{literal}}$ ,  $\mathcal{N}_{\text{semantic}}$ ), demonstrating an emerging ability to discriminate reliable from misleading knowledge. However, it shows no improvement under syntactic perturbations ( $\mathcal{N}_{\text{struct}}$ ), indicating that its robustness is limited to semantic noise and does not generalize to syntactic distortions.

**Low-resource languages face greater difficulty in balancing noise robustness and performance with accurate retrieval.** Both training-free and training-based strategies improve robustness but degrade performance under  $\mathcal{G}$ , with the trade-off most pronounced in low-resource settings. This tension arises because neither approach can dynamically modulate reliance on retrieved context; instead, they apply static policies that inevitably sacrifice either noise resilience or accurate context utility. The strong dependence on hyperparameters like the blending weight  $\alpha$  and training data proportions further underscores their brittleness.

## Conclusions

In this paper, we address the critical gap in understanding REAL-MT’s reliability under noisy retrieval, a common yet overlooked challenge in real-world deployment. To this end, we propose a controlled noise synthesis framework and two new metrics, Fidelity and Context Adoption Rate (CAR), to systematically evaluate REAL-MT’s robustness in knowledge-intensive translation tasks. Our evaluation, using both open- and closed-sourced models across high-, medium-, and low-resource language pairs, reveals that all models exhibit a high sensitivity to semantic noise. This is especially pronounced in low-resource scenarios, where models are more likely to uncritically rely on the noisy external context, often leading to nonsensical translations. Our experiments on LRM uncover that it is more vulnerable to noisy context than the non-reasoning version and, critically, tend to rationalize the noise even when aware of it, failing to demonstrate truthful reasoning. An analysis of attention mechanisms and model uncertainty reveals two failures: a shift of attention away from source idioms towards noisy context, compounded by a rise in model confidence despite a decline in accuracy. This points to critically weak self-verification and poor calibration. While both training-free and fine-tuning strategies can enhance robustness under noisy contexts, they do so at the cost of performance in accurate contexts. This trade-off underscores the necessity of developing models with self-reflective capabilities.

## Acknowledgments

This work was supported in part by the Science and Technology Development Fund of Macau SAR (Grant Nos. FDCT/0007/2024/AKP, FDCT/0070/2022/AMJ, FDCT/060/2022/AFJ), the National Natural Science Foundation of China (Grant Nos. 62261160648, 62266013), the China Strategic Scientific and Technological Innovation Cooperation Project (Grant No. 2022YFE0204900), and the UM and UMDF (Grant Nos. MYRG-GRG2025-00236-FST, MYRG-GRG2023-00006-FST-UMDF, MYRG-GRG2024-00165-FST-UMDF, EF2024-00185-FST, EF2023-00151-FST).

## References

- Bi, B.; Liu, S.; Wang, Y.; Xu, Y.; Fang, J.; Mei, L.; and Cheng, X. 2025. Parameters vs. context: Fine-grained control of knowledge reliance in language models. *arXiv preprint arXiv:2503.15888*.
- Cheang, C.; Chan, H.; Wong, D.; Liu, X.; Li, Z.; Sun, Y.; Liu, S.; and Chao, L. 2023. Can lms generalize to future data? an empirical analysis on text summarization. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 16205–16217.
- Cheang, C. S.; Chan, H. P.; Zhang, W.; and Deng, Y. 2025. Large Language Models Do NOT Really Know What They Don't Know. *arXiv preprint arXiv:2510.09033*.
- Chen, G.; Yao, Y.; Chao, L. S.; Liu, X.; and Wong, D. F. 2025a. SGIC: A Self-Guided Iterative Calibration Framework for RAG. *arXiv preprint arXiv:2506.16172*.
- Chen, S.; Shi, X.; Li, P.; Li, Y.; and Liu, J. 2024. Refining Translations with LLMs: A Constraint-Aware Iterative Prompting Approach. *arXiv preprint arXiv:2411.08348*.
- Chen, Y.; Benton, J.; Radhakrishnan, A.; Uesato, J.; Denison, C.; Schulman, J.; Somani, A.; Hase, P.; Wagner, M.; Roger, F.; et al. 2025b. Reasoning Models Don't Always Say What They Think. *arXiv preprint arXiv:2505.05410*.
- Donthi, S.; Spencer, M.; Patel, O.; Doh, J.; and Rodan, E. 2024. Improving LLM Abilities in Idiomatic Translation. *arXiv preprint arXiv:2407.03518*.
- Donthi, S.; Spencer, M.; Patel, O. B.; Doh, J. Y.; Rodan, E.; Zhu, K.; and O'Brien, S. 2025. Improving LLM Abilities in Idiomatic Translation. In *Proceedings of the First Workshop on Language Models for Low-Resource Languages*, 175–181.
- Fang, F.; Bai, Y.; Ni, S.; Yang, M.; Chen, X.; and Xu, R. 2024. Enhancing noise robustness of retrieval-augmented language models with adaptive adversarial training. *arXiv preprint arXiv:2405.20978*.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; Chen, W.; et al. 2022. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2): 3.
- Joshi, P.; Santy, S.; Budhiraja, A.; Bali, K.; and Choudhury, M. 2020. The state and fate of linguistic diversity and inclusion in the NLP world. *arXiv preprint arXiv:2004.09095*.
- Kwon, W.; Li, Z.; Zhuang, S.; Sheng, Y.; Zheng, L.; Yu, C. H.; Gonzalez, J. E.; Zhang, H.; and Stoica, I. 2023. Efficient Memory Management for Large Language Model Serving with PagedAttention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles*.
- Li, S.; Chen, J.; Yuan, S.; Wu, X.; Yang, H.; Tao, S.; and Xiao, Y. 2024. Translate Meanings, Not Just Words: IdiomKB's Role in Optimizing Idiomatic Translation with Language Models. In Wooldridge, M. J.; Dy, J. G.; and Natarajan, S., eds., *Thirty-Eighth AAAI Conference on Artificial Intelligence, AAAI 2024, Thirty-Sixth Conference on Innovative Applications of Artificial Intelligence, IAAI 2024, Fourteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2014, February 20-27, 2024, Vancouver, Canada, 18554–18563*. AAAI Press.
- Liu, E.; Chaudhary, A.; and Neubig, G. 2023. Crossing the Threshold: Idiomatic Machine Translation through Retrieval Augmentation and Loss Weighting. In Bouamor, H.; Pino, J.; and Bali, K., eds., *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023*, 15095–15111. Association for Computational Linguistics.
- Liu, Y.; Iter, D.; Xu, Y.; Wang, S.; Xu, R.; and Zhu, C. 2023. G-Eval: NLG Evaluation using Gpt-4 with Better Human Alignment. In Bouamor, H.; Pino, J.; and Bali, K., eds., *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023*, 2511–2522. Association for Computational Linguistics.
- Liu, Y.; Ott, M.; Goyal, N.; Du, J.; Joshi, M.; Chen, D.; Levy, O.; Lewis, M.; Zettlemoyer, L.; and Stoyanov, V. 2019. RoBERTa: A Robustly Optimized BERT Pretraining Approach. *arXiv preprint arXiv:1907.11692*.
- Merx, R.; Mahmudi, A.; Langford, K.; de Araujo, L. A.; and Vylomova, E. 2024. Low-resource machine translation through retrieval-augmented LLM prompting: a study on the Mambai language. *arXiv preprint arXiv:2404.04809*.
- Papineni, K.; Roukos, S.; Ward, T.; and Zhu, W. 2002. Bleu: a Method for Automatic Evaluation of Machine Translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics, July 6-12, 2002, Philadelphia, PA, USA*, 311–318. ACL.
- Park, S.-I.; and Lee, J.-Y. 2024. Toward Robust RALMs: Revealing the Impact of Imperfect Retrieval on Retrieval-Augmented Language Models. *Transactions of the Association for Computational Linguistics*, 12: 1686–1702.
- Peng, K.; Ding, L.; Zhong, Q.; Shen, L.; Liu, X.; Zhang, M.; Ouyang, Y.; and Tao, D. 2023. Towards Making the Most of ChatGPT for Machine Translation. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, 5622–5633.
- Rei, R.; De Souza, J. G.; Alves, D.; Zerva, C.; Farinha, A. C.; Glushkova, T.; Lavie, A.; Coheur, L.; and Martins, A. F. 2022a. COMET-22: Unbabel-IST 2022 submission for the metrics shared task. In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, 578–585.

- Rei, R.; Stewart, C.; Farinha, A. C.; and Lavie, A. 2020. COMET: A Neural Framework for MT Evaluation. In Webber, B.; Cohn, T.; He, Y.; and Liu, Y., eds., *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, 2685–2702. Association for Computational Linguistics.
- Rei, R.; Treviso, M.; Guerreiro, N. M.; Zerva, C.; Farinha, A. C.; Maroti, C.; de Souza, J. G.; Glushkova, T.; Alves, D. M.; Lavie, A.; et al. 2022b. COMETKIWI: IST-Unbabel 2022 Submission for the Quality Estimation Shared Task. *WMT 2022*, 634.
- Rezaeimanesh, S.; Hosseini, F.; and Yaghoobzadeh, Y. 2025. Large Language Models for Persian-English Idiom Translation. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, 7974–7985.
- Shen, X.; Blloshmi, R.; Zhu, D.; Pei, J.; and Zhang, W. 2024. Assessing “Implicit” Retrieval Robustness of Large Language Models. *arXiv preprint arXiv:2406.18134*.
- Stap, D.; Hasler, E.; Byrne, B.; Monz, C.; and Tran, K. 2024. The fine-tuning paradox: Boosting translation quality without sacrificing LLM abilities. *arXiv preprint arXiv:2405.20089*.
- Sun, Y.; Liu, X.; Wong, D. F.; Lin, Y.; Li, B.; Zhan, R.; Chao, L. S.; and Zhang, M. 2024. Understanding and Improving Low-Resource Neural Machine Translation with Shallow Features. In *CCF International Conference on Natural Language Processing and Chinese Computing*, 227–239. Springer.
- Team, Q. 2024. Qwen2.5: A Party of Foundation Models. <https://qwenlm.github.io/blog/qwen2.5/>.
- Team, Q. 2025. Qwen3 Technical Report. *arXiv:2505.09388*.
- Wang, H.; Xue, B.; Zhou, B.; Zhang, T.; Wang, C.; Wang, H.; Chen, G.; and Wong, K.-F. 2025. Self-DC: When to reason and when to act? self divide-and-conquer for compositional unknown questions. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, 6510–6525.
- Wiegreffe, S.; and Pinter, Y. 2019. Attention is not not Explanation. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 11–20.
- Xia, Y.; Zhou, J.; Shi, Z.; Chen, J.; and Huang, H. 2025. Improving retrieval augmented language model with self-reasoning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 39, 25534–25542.
- Yang, S.; Wu, J.; Ding, W.; Wu, N.; Liang, S.; Gong, M.; Zhang, H.; and Zhang, D. 2025. Quantifying the Robustness of Retrieval-Augmented Language Models Against Spurious Features in Grounding Data. *arXiv preprint arXiv:2503.05587*.
- Yoran, O.; Wolfson, T.; Ram, O.; and Berant, J. 2023. Making retrieval-augmented language models robust to irrelevant context. *arXiv preprint arXiv:2310.01558*.
- Zebaze, A.; Sagot, B.; and Bawden, R. 2025. Compositional Translation: A Novel LLM-based Approach for Low-resource Machine Translation. *arXiv preprint arXiv:2503.04554*.
- Zhan, R.; Liu, X.; Wong, D. F.; and Chao, L. S. 2021. Meta-curriculum learning for domain adaptation in neural machine translation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 14310–14318.
- Zhan, R.; Yang, X.; Wong, D.; Chao, L.; and Zhang, Y. 2024. Prefix text as a yarn: Eliciting non-english alignment in foundation language model. In *Findings of the Association for Computational Linguistics: ACL 2024*, 12131–12145.
- Zhou, Y.; Liu, Y.; Li, X.; Jin, J.; Qian, H.; Liu, Z.; Li, C.; Dou, Z.; Ho, T.-Y.; and Yu, P. S. 2024. Trustworthiness in retrieval-augmented generation systems: A survey. *arXiv preprint arXiv:2409.10102*.