

RIVAL: Reinforcement Learning with Iterative and Adversarial Optimization for Machine Translation

Tianjiao Li[★], Mengran Yu[★], Chenyu Shi[★][◇],
Yanjun Zhao[◇][★]; Xiaojing Liu[★], Qiang Zhang[★][†]
Qi Zhang[◇], Xuanjing Huang[◇], Jiayin Wang[◇]

[★]Bilibili Inc., Shanghai, China

[◇]Xi'an Jiaotong University, China

[†]School of Computer Science, Fudan University, China

{litianjiao01, yumengran03, shichenyu}@bilibili.com

Abstract

Large language models (LLMs) possess strong multilingual capabilities, and combining Reinforcement Learning from Human Feedback (RLHF) with translation tasks has shown great potential. However, we observe that this paradigm performs unexpectedly poorly when applied to colloquial subtitle translation tasks. In this work, we investigate this issue and find that the offline reward model (RM) gradually diverges from the online LLM due to distributional shift, ultimately leading to undesirable training outcomes. To address this, we propose RIVAL, an adversarial training framework that formulates the process as a min-max game between the RM and the LLM. RIVAL iteratively updates the both models, with the RM trained to distinguish strong from weak translations (qualitative preference reward), and the LLM trained to enhance its translation for closing this gap. To stabilize training and improve generalizability, we also incorporate quantitative preference reward (e.g., BLEU) into the RM, enabling reference-free quality modeling aligned with human evaluation. Through extensive experiments, we demonstrate that the proposed training framework significantly improves upon translation baselines.¹

1 Introduction

Recent advances in pre-trained large language models (LLMs) have yielded state-of-the-art results across a wide range of benchmarks (Achiam et al., 2023; Grattafiori et al., 2024; Bai et al., 2023). In the field of neural machine translation (NMT), researchers have also been exploring ways to leverage the powerful capabilities of LLMs to improve translation quality. Although most existing methods (Wiseman and Rush, 2016; Ranzato et al.,

¹This work was done when Yanjun Zhao was an intern at Bilibili Inc.

[†]Corresponding Author

¹Our code and data are available at: <https://github.com/InvokerStark/RIVAL>.

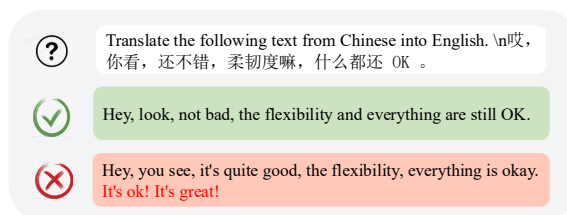


Figure 1: The illustration highlights the issue arising when directly applying vanilla RLHF to a colloquial subtitle dataset. The weak translator tends to generate content that is not present in the source text but receives a favorable score from the RM.

2015) rely on Maximum Likelihood Estimation (MLE)-based supervised fine-tuning, they suffer from exposure bias, leading to error accumulation and degraded translation quality. Moreover, such models often lack global coherence due to the limited modeling of sentence-level context (Kiegl and Kreutzer, 2021; Kreutzer et al., 2018). These issues have catalyzed growing interest in alternative approaches such as Reinforcement Learning from Human Feedback (RLHF) (He et al., 2025; Feng et al., 2025; Tan and Monz, 2025).

Currently, most NMT systems (Mohiuddin et al., 2022; Bansal et al., 2022) focus on formal and written language, with limited research addressing loosely structured and colloquial subtitles. To bridge this gap, we first construct a dedicated dataset and apply RLHF to optimize the translation quality for this task. This dataset is characterized by its diversity in both linguistic styles and domains. However, we find that the performance of vanilla RLHF is suboptimal. As shown in Figure 1, the model tends to generate content that does not exist in the source text for cheating the RM to obtain a high score, which is usually called reward hacking. Through analysis, we identify that this issue arises from the continuous distributional shift (Touvron et al., 2023; Luo et al., 2024) of the LLM during RL training, indicating that offline RM is

increasingly ineffective in providing reliable supervision signals, particularly in the context of online RL training.

To address the issue of distribution shift, we propose RIVAL, (**R**einforcement learning with **I**terative and **a**d**V**ersari**A**L optimization) an approach inspired by the adversarial training paradigm of generative adversarial networks (GANs) (Goodfellow et al., 2014). We reformulate the conventional RLHF training scheme as an explicit min-max game between the RM and the LLM. The RM is trained to maximize the score gap between the translation pairs produced by strong translators and weak translators. We name this signal the qualitative preference reward. The LLM training objective is to minimize the gap in translation quality between itself and strong translators. By iteratively optimizing both models, the RM is continuously updated to adapt to the distributional shifts introduced during LLM training.

In experiments, we further observe that as training progresses, the RM may learn a distribution divergent from the true reward signal. We attribute this issue to the large exploration space inherent in translation tasks, where purely qualitative preference reward may not consistently ensure beneficial optimization. To mitigate this, we introduce quantitative preference rewards (e.g., BLEU (Papineni et al., 2002)) into our adversarial framework to stabilize the iterative training process. However, directly using these signals suffers from poor generalizability (Sellam et al., 2020) and makes them unsuitable as dynamic rewards. We extend the RM to simultaneously predict this type of quantitative preference reward for incorporating them into our adversarial framework. This approach enables more generalized reference-free modeling of translation quality while maintaining alignment with human evaluation criteria through indirect reference guidance.

Our core contributions are as follows:

1. In this work, we propose an iterative and adversarial RL approach named RIVAL, which alternately optimizes the RM and the LLM through a competitive process, progressively improving performance from weak to strong by incorporating both qualitative and quantitative preference rewards.
2. Through extensive experiments, we demonstrate that our RIVAL effectively improves in-domain translation quality without compro-

promising the out-of-distribution performance of the model.

3. We provide a simple and effective pipeline for processing colloquial Chinese-English parallel subtitle translation datasets and release the processed data as an open-source resource.

2 Background

In the Bradley-Terry model (Bradley and Terry, 1952), a pairwise choice between items i and j depends on their latent utilities. The probability of selecting the item i increases monotonically with its utility relative to that of j . Formally, these probabilities are obtained by applying a softmax function to the log-utilities $r(\cdot)$ of the items.

$$Q(i \succ j) = \frac{\exp(r(i))}{\exp(r(i)) + \exp(r(j))} \quad (1)$$
$$= \text{softmax}(r(i), r(j)).$$

While this formulation elegantly captures static preference judgments, real-world applications—particularly in natural-language generation often require aligning model behavior with nuanced human values. RLHF addresses this gap by leveraging human evaluators to provide comparative or scalar feedback on model outputs. First, a RM is trained to predict these human judgments from pairs or ratings of model candidates (Christiano et al., 2017b; Stiennon et al., 2020). Next, the LLM is fine-tuned via policy optimization methods such as Proximal Policy Optimization (PPO) to maximize the learned reward signal (Schulman et al., 2017). This two-stage process has proven effective in producing more helpful, truthful, and harmless generations in large-scale language systems (Ouyang et al., 2022). In this paper, RLHF is not exclusively limited to human feedback, but also encompasses model-generated feedback.

3 Pilot Experiment and Analysis

In this section, we first explore the application of vanilla RLHF to NMT and conduct a preliminary analysis of the experimental results.

3.1 How to Apply RLHF to NMT?

Inspired by RLHF, we model the translation quality problem as a rank-wise comparison between a weak translator and a strong translator. Here, we refer to the translation model being trained as a

weak translator, while grouping together gold references and outputs from stronger models under the term strong translator. Accordingly, we adopt the RM framework to train an evaluator for translation quality r_ϕ , as formulated below (Ouyang et al., 2022):

$$\mathcal{L}_{\text{rank}}(r_\phi; \mathcal{D}_{\text{RM}}) = -\mathbb{E}_{\mathcal{D}_{\text{RM}}} [\log Q(y^s \succ y^w | x)], \quad (2)$$

where $\mathcal{D}_{\text{RM}} = (x, y^s, y^w)$ is parallel translation dataset, y^s and y^w are the strong translation and the weak translation of the text x , respectively.

With a learned RM, we first sample a set of candidate responses $\{y_1, y_2, \dots, y_G\}$ from the previous policy $\pi_{\theta_{\text{old}}}$ when presented with an original text x . Then the advantage function A_i is calculated by normalizing each individual reward relative to the distribution of all rewards in the group. Specifically: $A_i = \frac{r_i - \mu_r}{\sigma_r}$ where $\mu_r = \text{mean}(\{r_1, r_2, \dots, r_G\})$ represents the average of rewards, $\sigma_r = \text{std}(\{r_1, r_2, \dots, r_G\})$ denotes the standard deviation, and $\{r_1, r_2, \dots, r_G\}$ are the rewards scored by the learned RM. Finally, GRPO(Shao et al., 2024) aims to maximize the following objective function to optimize π_θ :

$$J_{\text{GRPO}}(\theta) = \mathbb{E}_{x \sim P(X), \{y_i\}_{i=1}^G \sim \pi_{\theta_{\text{old}}}(Y|x)} \left[\frac{1}{G} \sum_{i=1}^G \min \left(\frac{\pi_\theta(y_i | x)}{\pi_{\theta_{\text{old}}}(y_i | x)} A_i, \text{clip} \left(\frac{\pi_\theta(y_i | x)}{\pi_{\theta_{\text{old}}}(y_i | x)}, 1 - \varepsilon, 1 + \varepsilon \right) A_i \right) - \beta D_{\text{KL}}(\pi_\theta \| \pi_{\text{ref}}) \right], \quad (3)$$

where the hyperparameter ε controls the PPO clipping threshold and β penalizes the Kullback–Leibler (KL) divergence between the optimized policy π_θ and the initial policy π_{ref} .

According to the approach mentioned above, we conduct pilot experiments applying RLHF to NMT in the task of colloquial subtitle translation. We detail the construction of a diverse and colloquial subtitle translation dataset drawn from real-world videos. The data processing pipeline follows a systematic approach:

1. Subtitle Extraction. We collect real-world videos and employ Automatic Speech Recognition technology to transcribe spoken content. The transcriptions are then segmented into discrete sen-

tences based on natural pauses and semantic completeness.

2. Subtitle Processing. We remove background music and semantically insignificant utterances. To enhance contextual coherence and conversational continuity, we structure the data into groups of ten sentences. For clarity and standardization, our prompt instructs models to follow in JSON format.

3. Translation Generation. Given the absence of golden translations for real-world video subtitles, we implement a dual-model approach to generate translation pairs of varying quality. We utilize GPT-4o (Achiam et al., 2023) to produce strong translations that served as targets and employ Qwen2.5-7B-Chat (Team, 2024) as our baseline model to generate weak translations, representing the initial performance that we aim to improve.

4. Dataset Construction. To construct clean data, we conduct simple low-quality filtering strategies on the model-generated responses, including: repetitive generation, failure to follow JSON format, empty generation, target language mismatch, etc., to ensure that the remain translation data is valid. Subsequently, we calculated the 2-gram similarity of the corresponding strong-weak translation pairs and heuristically set different filtering thresholds to filter similar responses. The filtered prompts and paired translations are incorporated into the RL corpus, while the combinations of prompts and strong translations form the supervised fine-tuning (SFT) dataset.

We provide several cases in Appendix A. Based on the datasets above, we conduct experiments on optimizing NMT using RLHF. Since colloquial translation prioritize free translation and semantic alignment rather than requiring word-for-word correspondence with the target text, we employ COMETkiwi(Rei et al., 2022) and model judgment, such as GPT-4o, as evaluation metrics.

3.2 Why Vanilla RLHF Fails in Colloquial Subtitle Translation?

To investigate the causes of the aforementioned phenomena, Figure 2 shows the scoring differential between strong translations and the translations generated by the optimized model, which are evaluated by the RM and GPT-4o. As training progresses, the scoring differential from the RM gradually decreases, suggesting that the RM perceives a progressive improvement in the model’s generated translations. However, the scoring differential from

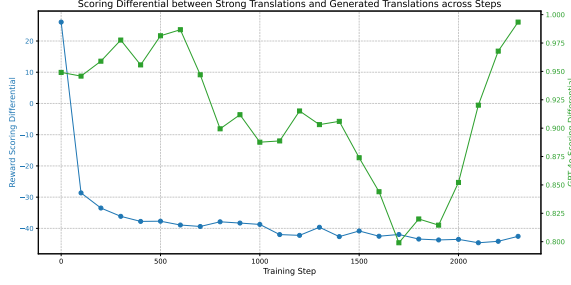


Figure 2: Scoring Differential of the RM and GPT-4o between strong translations and the translations generated by the optimized model.

GPT-4o exhibits an initial decrease followed by a subsequent increase. The significant discrepancy between these two evaluation metrics indicates that the RM trained solely on weak translations from the initial model struggles to adapt to the distributional shift occurring in the training process.

4 Methods

Based on the experimental findings in the previous section, we observe that vanilla RLHF is not well-suited for colloquial subtitle translation tasks. Therefore, inspired by the principles of GANs (Goodfellow et al., 2014), we propose an adversarial RL approach, RIVAL, that enables both the RM and the LLM to progressively improve from weak to strong.

Specifically, we formulate the two-stage training process of RLHF as a min-max adversarial game between LLM π_θ (Generator) and RM r_ϕ (Discriminator), as expressed by the following objective:

$$\min_{r_\phi} \max_{\pi_\theta} \mathbb{E}[r_\phi(x, \pi_\theta(y|x))] - \mathbb{E}_{y \in P_{\text{strong}}}[r_\phi(x, y)] \quad \text{s.t.} \quad \text{KL}[\pi_\theta(y|x) \parallel \pi_{\text{ref}}(y|x)] < \eta, \quad (4)$$

where P_{strong} denotes the output distribution of a strong translator. The LLM, as a weak translator, is trained to minimize the gap between its output distribution and that of the strong translator. In contrast, the RM is trained to maximize this distributional gap by distinguishing between the outputs of the weak and strong translators.

By iteratively optimizing both models and using the current LLM to reconstruct new training data for the RM, the RM can effectively learn to serve as a high-quality evaluator for translation quality. Meanwhile, a well-trained RM can in turn provide effective guidance for the LLM, enabling

it to explore the open-ended output space and progressively learn to become a strong translator.

4.1 RM and LLM Optimization Step

In this subsection, we provide a detailed explanation of the RM and LLM optimization steps.

4.1.1 RM Step

The goal of the RM step is to train a high-quality RM to accurately evaluate the quality of the translation. In this step, we keep π_θ fixed and optimize only r_ϕ . Therefore, the objective 4 can be simplified to the following form:

$$\min_{r_\phi} \mathbb{E}[r_\phi(x, \pi_\theta(y|x))] - \mathbb{E}_{y \in P_{\text{strong}}}[r_\phi(x, y)]. \quad (5)$$

This objective means that the model aims to maximize the gap between the strong translator and the current-round translator. Our translation task uses source texts x and their translations y_{strong} from the dataset \mathcal{D}_{RM} , with the translations originating from a strong translator—either a more powerful model or a human expert. We then use the current-round LLM π_θ to translate these source texts, obtaining the weak translator’s outputs $\pi_\theta(y|x)$. At this stage, we apply a simple filtering process to exclude samples where the translations from the weak and strong translators are already highly similar, as they offer limited learning signal. Therefore, we can derive Equation 5 as follows:

$$\begin{aligned} & \min_{r_\phi} \mathbb{E}[r_\phi(x, \pi_\theta(y|x))] - \mathbb{E}_{y \in P_{\text{strong}}}[r_\phi(x, y)] \\ &= \min_{r_\phi} \mathbb{E}[r_\phi(x, y_{\text{weak}})] - \mathbb{E}[r_\phi(x, y_{\text{strong}})] \\ &= \max_{r_\phi} \mathbb{E}[r_\phi(x, y_{\text{strong}}) - r_\phi(x, y_{\text{weak}})]. \quad (6) \end{aligned}$$

The objective in Equation 6 aligns with the structure of the RM, and thus we directly adopt the traditional rank loss as the training loss for the RM step. In this work, we refer to the rank loss as a qualitative preference loss. Furthermore, during subsequent iterations, we not only utilize translations generated by the current-round LLM but also replay a subset of outputs from previous rounds. This approach helps prevent excessive distributional shifts, enhances data diversity, and effectively improves the model’s robustness.

4.1.2 LLM Step

In the LLM update step, we similarly keep the reward model r_ϕ fixed and update only the param-

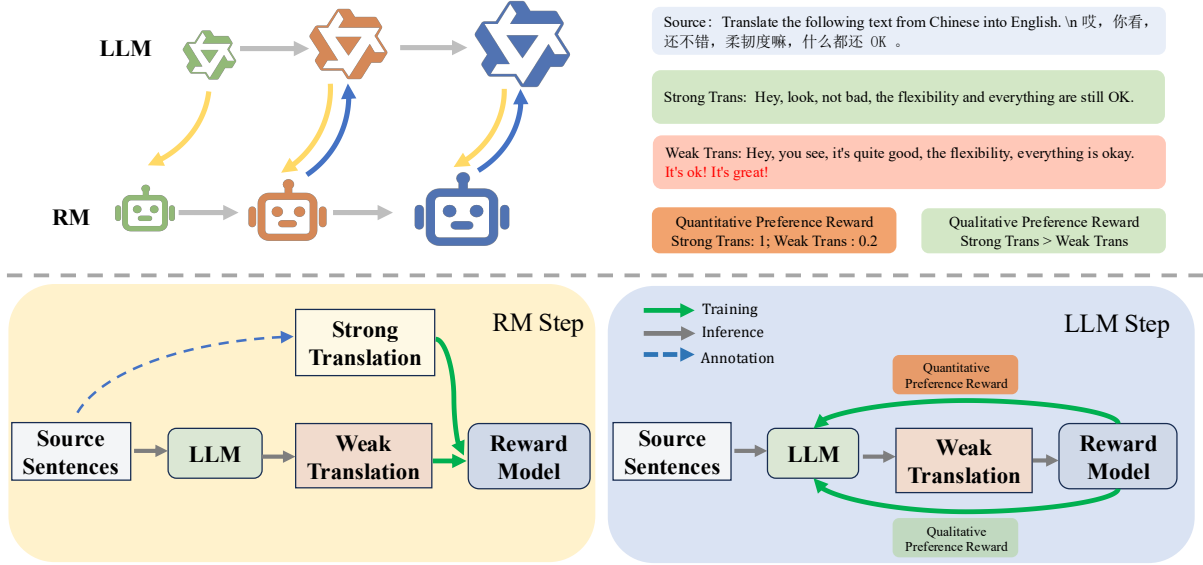


Figure 3: The RIVAL framework. In the RM updating step, the RM is trained to maximize the score gap between translation pairs produced by strong and weak translators. In the LLM updating step, the LLM is to alleviate the quality gap between its own outputs and those of the strong translators.

ters of the LLM π_θ :

$$\begin{aligned} \max_{\pi_\theta} \mathbb{E}_{x \in D_{LLM}} [r_\phi(x, \pi_\theta(y|x))] \\ \text{s.t. } \text{KL}[\pi_\theta(y|x) || \pi_{\text{ref}}(y|x)] < \eta. \end{aligned} \quad (7)$$

The objective of this training step is to optimize the LLM using the signals provided by the RM. This procedure is aligned with the standard preference learning paradigm, where the LLM is guided to produce the outputs preferred by the RM. In line with the earlier stages of our pipeline, we continue to adopt the GRPO algorithm as our RL method in this phase.

4.2 Incorporate Quantitative Preference Reward

Moreover, our experiments reveal that using only the qualitative preference reward leads to instability across training iterations. This can be attributed to the translation model’s large exploration space, which will be discussed in detail later in the paper. To address this, we introduce a quantitative preference reward to stabilize and align the optimization trajectory of the π_θ throughout the training iterations. Given that the qualitative preference reward emphasizes semantic alignment, we opt for BLEU as a quantitative preference reward to facilitate lexical alignment. Specifically, rather than directly employing BLEU as a reward signal, we train a RM to approximate BLEU scores.

The key insight is that reference translations may contain errors (Xu et al., 2024b) or lack diversity,

resulting in biased BLEU calculations that lack robustness. Prior work (Benedetti and Ventura, 2024; de Jorge Aranda et al., 2022; Bishop, 1995) show that noisy-based training can, to some extent, be equivalent to Tikhonov regularization, effectively mitigating overfitting and improving model robustness. Therefore, this work employs noisy data to train a RM that approximates BLEU scores, resulting in reduced sensitivity to noise and enhanced robustness.

Moreover, instead of training a separate model for quantitative preference reward, we leverage the same model backbone used for qualitative preference reward and introduce an additional output head dedicated to quantitative preference reward prediction. This design offers two key benefits: it reduces the computational cost of model training and enables mutual learning between the two tasks with different supervision forms, thereby mitigating the risk of over-optimization (Ahmed et al., 2024). Finally, our multi-head RM loss function incorporating BLEU is defined as follows:

$$\begin{aligned} \mathcal{L}_{RM} &= \mathcal{L}_{\text{qualitative}} + \alpha \mathcal{L}_{\text{quantitative}} \\ &= \mathcal{L}_{\text{rank}} + \alpha \mathcal{L}_{\text{MAE}}(r_\phi; \mathcal{D}_{RM}) \\ &= \mathcal{L}_{\text{rank}} + \alpha \mathbb{E}_{\mathcal{D}_{RM}} [|y_{\text{BLEU}}^s - y_{\text{BLEU}}^w|], \end{aligned} \quad (8)$$

where α is a coefficient used to balance the weights of the two components. In this paper, we do not assign it a specific value and thus set it to 1. Additionally, we recommend using Mean Absolute Error (MAE) as the loss function instead of Mean

Algorithm 1: RIVAL

Require: $\mathcal{D}_{\mathcal{RM}}, \mathcal{D}_{\mathcal{RM}}^*, \mathcal{D}_{\mathcal{LLM}}, \pi_\theta, r_\phi$, iterations N , Training steps T_{RM}, T_{LLM} for RM and LLM

- 1: **for** $k = 1, \dots, N$ **do**
- 2: //Process Data and Get BLEU.
- 3: **for** $i = 1, \dots, |\mathcal{D}_{\mathcal{RM}}|$ **do**
- 4: Sample piece $(x, y_i^s, y_i^w) \in \mathcal{D}_{\mathcal{RM}}$
- 5: **if** $\text{sim}(y_i^s, y_i^w) < \tau$ **then**
- 6: //Get BLEU for translator.
- 7: Update $\mathcal{D}_{\mathcal{RM}}^* \leftarrow (x, \text{BLEU}(y_i^s, y_i^s), \text{BLEU}(y_i^s, y_i^w))$
- 8: **end if**
- 9: **end for**
- 10: //RM Step.
- 11: **for** $t = 1, \dots, T_{RM}$ **do**
- 12: Sample batch $\mathcal{B}_{\mathcal{RM}} \subset \mathcal{D}_{\mathcal{RM}}^*$
- 13: //Update RM using Equation 8
- 14: $\phi_t = \phi_{t-1} - \eta_t^{RM} * \nabla_\phi \mathcal{L}_{\mathcal{RM}}(\phi_{t-1})$
- 15: **end for**
- 16: //LLM Step.
- 17: **for** $t = 1, \dots, T_{LLM}$ **do**
- 18: Sample $\mathcal{B}_{\mathcal{LLM}} \subset \mathcal{D}_{\mathcal{LLM}}$
- 19: Generate $\hat{y} = \text{LLM}_{\theta_t}(x), x \in \mathcal{B}_{\mathcal{LLM}}$
- 20: Compute rewards $r = \text{RM}_{\phi_k}(x, \hat{y})$
- 21: //Update LLM using Equation 3
- 22: $\theta_t = \theta_{t-1} - \eta_t^{LLM} * \nabla_\theta J(\theta_{t-1})$
- 23: **end for**
- 24: //Process new RM Training Data.
- 25: Update $\mathcal{D}_{\mathcal{RM}} \leftarrow (x, y^s, \text{LLM}_{\theta_k}(x))$
- 26: **end for**
- 27: **return** θ_{LLM}, ϕ_{RM}

Squared Error (MSE); this choice will be further discussed in experiments. The detailed algorithmic procedure can be found in Algorithm 1.

5 Experiment

We will introduce experimental setup, main results and analysis in this section.

5.1 Experimental Setup

In this subsection, we present the experimental setup and implementation details.

Datasets. Our main experiments are conducted on the WMT dataset and our proposed subtitle dataset. For the WMT dataset, we collect Chinese(ZH)-English(EN) parallel corpora from WMT 2017-2020² and perform basic preprocessing, removing

²<https://www.statmt.org/wmt20/>

sentence pairs with fewer than 30 characters(Xu et al., 2024a; Feng et al., 2024). Additionally, we use English(EN)-German(DE) and Chinese(ZH)-German(DE) medical translation tasks from the WMT dataset as out-of-distribution (OOD) language settings for evaluation.

Evaluation Metrics. On the subtitle dataset, we use COMETKiwi and GPT-4o for evaluation because (1) subtitle tasks prioritize semantic content, making BLEU less appropriate; and (2) the references are GPT-4o-generated, making it infeasible to compute BLEU scores against genuine reference translations. For the WMT dataset, both BLEU and COMETKiwi are used to capture lexical fidelity and semantic adequacy.

GPT-4o-based scoring follows the LLM-as-a-Judge(Lee et al., 2024; Weng et al., 2022) paradigm, evaluating translations across four dimensions: accuracy, completeness, coherence, and stylistic consistency. The specific prompt used for evaluation can be found in the Appendix B. The consistency between human evaluation and GPT-4o scoring can be found in the Appendix C.

Baselines. Our baselines include the strong general-purpose model GPT-4o, the translation-specific model Tower-7B-v0.2(Alves et al., 2024), and the models obtained by applying SFT.

Implementation Details. We train our RM using Megatron framework (Shoeybi et al., 2019) with the Qwen2.5-72B-Chat. For LLM training, we employ the Verl framework(Sheng et al., 2024) with the Qwen2.5-7B-Chat. For more training details, please refer to the Appendix D.

5.2 Main Results

Table 1 displays the performance on the subtitle task with only the qualitative preference reward, and Table 2 presents results on WMT task with both qualitative and quantitative preference rewards. We also provide the result of significance test in Appendix E

Colloquial subtitle translation is a particularly challenging task. As shown in Table 1, even TowerInstruct-7B-v0.2—a translation-specific LLM trained with approximately 20B tokens of continued pretraining—fails to perform well on this task. According to GPT-4o evaluations, the primary issue lies in insufficient accuracy. This indicates that our ASR-based dataset contains highly diverse and substantial knowledge-intensive content, making it a relatively challenging task that

	In-Domain Subtitle ZH->EN						
	Accuracy	Completeness	Coherence	Stylistic Consistency	Avg	COMETkiwi	
GPT-4o	3.99	4.00	4.00	4.00	4.00	66.82	
TowerInstruct-7B-v0.2	2.69	3.69	3.56	3.71	3.37	64.52	
Qwen2.5-7B-Chat	2.52	3.4	3.32	3.48	3.13	64.03	
Qwen2.5-7B-SFT	2.97	3.74	3.46	3.75	3.46	65.89	
RIVAL-Iter0-Qual	2.69	3.4	3.32	3.48	3.26	65.22	
RIVAL-Iter1-Qual	3.28	3.86	3.84	3.89	3.68	66.27	
RIVAL-Iter2-Qual	3.06	3.82	3.57	3.79	3.53	66.49	
	OOD Medical ZH->DE						
	Accuracy	Completeness	Coherence	Stylistic Consistency	Avg	COMETkiwi	
Qwen2.5-7B-Chat	2.33	3.19	3.14	3.28	2.94	52.58	
Qwen2.5-7B-SFT	2.34	2.99	2.98	3.08	2.81	49.15	
RIVAL-Iter0-Qual	2.48	3.15	3.10	3.26	2.96	52.41	
RIVAL-Iter1-Qual	2.46	3.23	3.12	3.29	2.99	53.42	
RIVAL-Iter2-Qual	2.33	3.16	3.06	3.21	2.93	51.76	

Table 1: Performance comparison on in-domain translation direction (Subtitle ZH-EN) and OOD translation task (Medical ZH-DE) using GPT-4o scores and COMETKiwi. Only the qualitative preference reward is used. All results are averaged by 3 times.

In-Domain	WMT EN->ZH		WMT ZH->EN	
	BLEU	COMETkiwi	BLEU	COMETkiwi
GPT-4o	38.98	75.64	32.4	73.59
TowerInstruct-7B-v0.2	38.69	71.45	31.61	71.67
Qwen2.5-7B-Chat	31.50	65.41	25.60	68.61
Qwen2.5-7B-SFT	38.77	71.39	32.22	71.88
RIVAL-Iter0-Qual	33.34	69.64	28.32	69.91
RIVAL-Iter1-Qual	31.74	70.83	27.65	71.19
RIVAL-Iter2-Qual	30.14	71.91	26.73	73.28
RIVAL-Iter0-Qual+Quant	34.76	69.12	29.52	70.16
RIVAL-Iter1-Qual+Quant	38.62	69.72	32.90	72.37
RIVAL-Iter2-Qual+Quant	39.39	72.60	33.42	73.61
OOD	WMT EN->DE		WMT DE->EN	
	BLEU	COMETkiwi	BLEU	COMETkiwi
Qwen2.5-7B-Chat	27.23	71.53	36.85	74.18
Qwen2.5-7B-SFT	22.16	67.88	34.69	71.19
RIVAL-Iter0-Qual	25.15	71.48	35.17	73.69
RIVAL-Iter1-Qual	24.71	70.93	34.97	73.06
RIVAL-Iter2-Qual	23.96	69.81	34.82	72.76
RIVAL-Iter0-Qual+Quant	26.44	71.50	35.35	73.49
RIVAL-Iter1-Qual+Quant	25.70	69.93	35.32	72.66
RIVAL-Iter2-Qual+Quant	25.25	68.74	35.25	72.23

Table 2: Performance comparison on in-domain WMT(EN-ZH) and OOD WMT(EN-DE) using BLEU and COMETKiwi. Both preference rewards are used. All results are averaged by 3 times.

has not yet been well addressed.

As the iterations progress, the model is able to gradually discover improved translation. As shown in Tables 1 and Table 2, our RIVAL method is able to progressively discover higher-quality translations through iterative optimization. The model outperforms current open-source translation-specific LLMs and, on certain tasks, the performance of RIVAL even exceeds that of strong general-purpose models. These results not only demonstrate the effectiveness of our approach but also highlight the potential of iterative optimization to enable autonomous exploration and continuous

adjustment toward improved translation quality.

Different reward signals serve different purposes. As shown in Table 2, our RIVAL method achieves comparable performance under both qualitative and quantitative rewards. However, when the quantitative reward is incorporated, the corresponding BLEU score improves significantly. In fact, omitting the quantitative reward leads to a drop in BLEU but an increase in COMET, suggesting that the model is able to explore effectively at the semantic level without being constrained by surface-level lexical similarity.

In Language OOD scenarios, our approach better preserves the model’s capabilities. As shown in Tables 1 and Table 2, SFT severely degrades performance on OOD languages. According to GPT-based evaluation, this performance drop is primarily observed in completeness, coherence, and stylistic consistency, while accuracy remains largely unaffected. This suggests that different languages share parameters related to knowledge within the model, whereas language-specific capabilities are more tightly coupled with the language itself (Xu et al., 2023). In contrast, our RIVAL method results in significantly less degradation in OOD language performance. Moreover, we observe an interesting phenomenon: on the WMT-Medical test dataset, the performance at RIVAL-Iter1 even surpasses that of the original model. This indirectly indicates that our method encourages the exploration of more effective translation strategies, rather than relying on rote memorization of specific patterns.

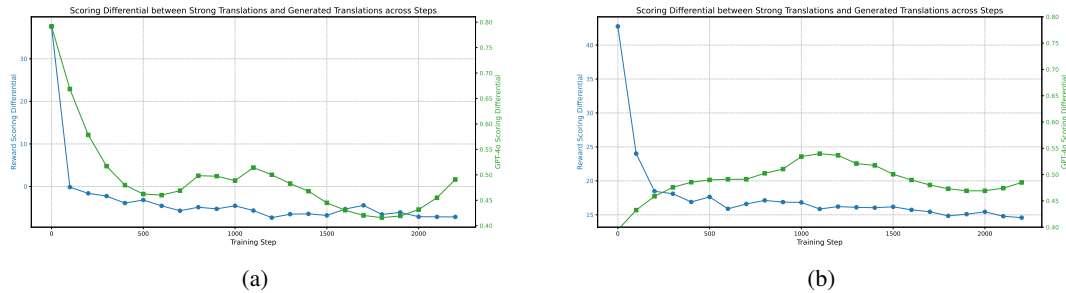


Figure 4: Scoring Differential of the RM and GPT-4o between strong and weak translations. (a) for RIVAL-Iter1 (b) for RIVAL-Iter2.

Additionally, we observed that models trained solely with qualitative rewards outperform the SFT method in OOD scenarios. Consistent with our in-domain findings, evaluation on both BLEU and COMETkiwi metrics revealed that training with qualitative rewards resulted in a significant reduction in character-level similarity, while semantic-level performance remained largely unaffected.

Purely the qualitative preference reward cannot guarantee a reliable optimization direction. As observed in Table 1, the performance of RIVAL-Iter2 is lower than that of RIVAL-Iter1. We hypothesize that this is due to the vast exploration space inherent in open-ended generation tasks, where purely qualitative preference rewards fail to effectively constrain the model’s exploration, leading to behaviors misaligned with the true reward signal. To address this, we introduce the quantitative preference reward, and as shown in Table 2, the model’s exploration results exhibit a high degree of consistency, indicating more stable and reliable.

5.3 Analysis

MAE is better than MSE. In our experiments, we find that using MAE as the loss function for training the RM yields significantly better results than using MSE. Detailed case studies and exact figures are provided in the Appendix F. MAE reduces the error by 80% compared to MSE and improves accuracy by 2.5%. We attribute this phenomenon to the squared term in MSE, which greatly diminishes the loss when the target values are less than 1, making it difficult for the model to effectively learn how to fit the quantitative preference rewards. Therefore, we recommend using MAE as the loss function.

RIVAL effectively mitigates distributional shift. We also plot the score differences between strong and weak translations in the last two iterations. As shown in the Figure 4, we can observe that the score differences remain within a relatively small range throughout training and do not exhibit sud-

den increases in the later stages. This suggests that during effective training steps, the issue of distributional shift is significantly alleviated and iterative adversarial approach enhances the generalization capabilities of both the RM and the LLM.

More iterations lead to better performance. Experimental results indicate that the model often reaches or exceeds the SFT baseline within two iterations. Continued gains observed in WMT dataset at the third iteration imply that further self-exploration may help the model surpass strong translators, which we consider a potential avenue for future work.

6 Related works

Here we present a concise overview on RL-based method for the MT(details in Appendix G).

Reinforcement Learning for Machine Translation Early RL studies alleviated the training–evaluation gap by directly optimizing corpus-level BLEU and alleviating exposure bias (Ranzato et al., 2016; Edunov et al., 2018; Wang and Sennrich, 2020). Later work used human-in-the-loop feedback for domain or user adaptation (Sokolov et al., 2016; Nguyen et al., 2017). Reward-shaping variants inject intermediate signals, yet even token-wise BLEU remains too coarse for fine-grained rewards (Wu et al., 2018; Goyal et al., 2019; Devidze et al., 2022).

Modern RL-enhanced LLMs such as OpenAI O1 and DeepSeek-R1 achieve state-of-the-art performance accross diverse benchmarks (OpenAI et al., 2024; DeepSeek-AI et al., 2025). Motivated by these, R1-Translator introduces chain-of-thought alignment for MT (He et al., 2025); MT-R1-Zero adds zero-shot hybrid rewards (Feng et al., 2025). ReMedy frames MT evaluation as pairwise reward modeling (Tan and Monz, 2025), while severity-weighted token rewards from XCOMET provide richer signals (Ramos et al., 2025).

Reward Hacking in Reinforcement Learn-

ing Reward hacking happens when an RL agent finds loopholes in the reward and chases the proxy instead of the real goal (Weng, 2024; Everitt and Hutter, 2019). It has appeared in robotics, Atari-style games, and RLHF language models, where agents loop waypoints, rewrite tests, or game metrics to boost scores without real gains (Lehman et al., 2020; Gleave et al., 2019; Denison et al., 2024). To prevent the above issue, recent work suggests the following methods of defence: (1) fix the reward with ensembles (Wortsman et al., 2022), adversarial training, caps, or trip-wires (Amodei et al., 2016); (2) harden the policy through look-ahead planning, adversarial blinding, or indifference methods (Uesato et al., 2020); (3) add human checks, anomaly detection, and data audits after training (Pan et al., 2022; Revel et al., 2025).

7 Conclusion

Using a real-world colloquial subtitle dataset, we find that vanilla RLHF struggles to improve translation quality due to distributional shifts from offline RM. To address this, we propose RIVAL—an adversarial framework enabling iterative optimization between the RM and the LLM, incorporating both qualitative and quantitative preference rewards. Through extensive experiments and analysis, we demonstrate the effectiveness and generalizability of this framework. Our work offers valuable insights for applying RL to NMT, and more broadly, serves as a reference for RL-based approaches in general post-training of LLMs.

Limitations

In this paper, we focus on how to optimize translation performance using the RLHF paradigm, with particular emphasis on the previously underexplored task of colloquial subtitle translation. Although our RIVAL method demonstrates superior performance over the baselines in experiments, we have not yet explored its full potential. We believe that with additional rounds of iteration, the model could achieve performance comparable to, or even surpass, that of strong translators. Moreover, we believe that this iterative adversarial optimization approach can be extended to more general post-training settings. Given its generalizability, we hope it can serve as an effective alternative to supervised fine-tuning.

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Ahmed M Ahmed, Rafael Rafailov, Stepan Sharkov, Xuechen Li, and Sanmi Koyejo. 2024. Scalable ensembling for mitigating reward overoptimisation. *arXiv preprint arXiv:2406.01013*.
- Duarte M Alves, José Pombal, Nuno M Guerreiro, Pedro H Martins, João Alves, Amin Farajian, Ben Peters, Ricardo Rei, Patrick Fernandes, Sweta Agrawal, et al. 2024. Tower: An open multilingual large language model for translation-related tasks. *arXiv preprint arXiv:2402.17733*.
- Dario Amodei, Chris Olah, Jacob Steinhardt, Paul Christiano, John Schulman, and Dan Mané. 2016. *Concrete problems in ai safety*. *arXiv preprint arXiv:1606.06565*.
- Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. 2023. Qwen technical report. *arXiv preprint arXiv:2309.16609*.
- Trapit Bansal, Jakub W. Pachocki, Szymon Sidor, Ilya Sutskever, and Igor Mordatch. 2017. *Emergent complexity via multi-agent competition*. *ArXiv*, abs/1710.03748.
- Yamini Bansal, Behrooz Ghorbani, Ankush Garg, Biao Zhang, Colin Cherry, Behnam Neyshabur, and Orhan Firat. 2022. Data scaling laws in nmt: The effect of noise and architecture. In *International Conference on Machine Learning*, pages 1466–1482. PMLR.
- Marco Benedetti and Enrico Ventura. 2024. Training neural networks with structured noise improves classification and generalization. *Journal of Physics A: Mathematical and Theoretical*, 57(41):415001.
- Chris M Bishop. 1995. Training with noise is equivalent to tikhonov regularization. *Neural computation*, 7(1):108–116.
- Ralph Allan Bradley and Milton E. Terry. 1952. *Rank analysis of incomplete block designs: I. the method of paired comparisons*. *Biometrika*, 39(3/4):324–345.
- Paul F Christiano, Jan Leike, Tom Brown, Miljan Martić, Shane Legg, and Dario Amodei. 2017a. Deep reinforcement learning from human preferences. *Advances in neural information processing systems*, 30.
- Paul F. Christiano, Jan Leike, Tom B. Brown, Miljan Martić, Shane Legg, and Dario Amodei. 2017b. Deep reinforcement learning from human preferences. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, NIPS’17, page 4302–4310, Red Hook, NY, USA. Curran Associates Inc.

- Pau de Jorge Aranda, Adel Bibi, Riccardo Volpi, Amartya Sanyal, Philip Torr, Grégory Rogez, and Puneet Dokania. 2022. Make some noise: Reliable and efficient single-step adversarial training. *Advances in Neural Information Processing Systems*, 35:12881–12893.
- DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, and Ruoyu Zhang. 2025. [Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning](#).
- Carson Denison, Monte MacDiarmid, Fazl Barez, David Duvenaud, Shauna Kravec, Samuel Marks, Nicholas Schiefer, Ryan Soklaski, Alex Tamkin, Jared Kaplan, et al. 2024. Sycophancy to subterfuge: Investigating reward-tampering in large language models. *arXiv preprint arXiv:2406.10162*.
- Rati Devidze, Parameswaran Kamalaruban, and Adish Singla. 2022. [Exploration-guided reward shaping for reinforcement learning under sparse rewards](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 5829–5842. Curran Associates, Inc.
- Lauro Langosco di Langosco, Jack Koch, Lee D. Sharkey, Jacob Pfau, and David Krueger. 2021. [Goal misgeneralization in deep reinforcement learning](#). In *International Conference on Machine Learning*.
- Sergey Edunov, Myle Ott, Michael Auli, David Grangier, and Marc’Aurelio Ranzato. 2018. [Classical structured prediction losses for sequence to sequence learning](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 355–364, New Orleans, Louisiana. Association for Computational Linguistics.
- Tom Everitt and Marcus Hutter. 2019. [Reward tampering problems and solutions in reinforcement learning: A causal influence diagram perspective](#). *ArXiv*, abs/1908.04734.
- Tom Everitt, Victoria Krakovna, Laurent Orseau, and Shane Legg. 2017. [Reinforcement learning with a corrupted reward channel](#). In *International Joint Conference on Artificial Intelligence*.
- Zhaopeng Feng, Shaosheng Cao, Jiahao Ren, Jiayuan Su, Ruizhe Chen, Yan Zhang, Zhe Xu, Yao Hu, Jian Wu, and Zuozhu Liu. 2025. [Mt-r1-zero: Advancing llm-based machine translation via r1-zero-like reinforcement learning](#).
- Zhaopeng Feng, Ruizhe Chen, Yan Zhang, Zijie Meng, and Zuozhu Liu. 2024. [Ladder: A model-agnostic framework boosting LLM-based machine translation to the next level](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 15377–15393, Miami, Florida, USA. Association for Computational Linguistics.
- Adam Gleave, Michael Dennis, Cody Wild, Neel Kant, Sergey Levine, and Stuart Russell. 2019. Adversarial policies: Attacking deep reinforcement learning. *arXiv preprint arXiv:1905.10615*.
- Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. *Advances in neural information processing systems*, 27.
- Prasoon Goyal, Scott Niekum, and Raymond J. Mooney. 2019. [Using natural language for reward shaping in reinforcement learning](#). In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pages 2385–2391. International Joint Conferences on Artificial Intelligence Organization.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Mingui He, Yilun Liu, Shimin Tao, Yuanchang Luo, Hongyong Zeng, Chang Su, Li Zhang, Hongxia Ma, Daimeng Wei, Weibin Meng, Hao Yang, Boxing Chen, and Osamu Yoshie. 2025. [R1-t1: Fully incentivizing translation capability in llms via reasoning learning](#).
- Samuel Kiegl and Julia Kreutzer. 2021. [Revisiting the weaknesses of reinforcement learning for neural machine translation](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1673–1681, Online. Association for Computational Linguistics.
- Julia Kreutzer, Shahram Khadivi, Evgeny Matusov, and Stefan Riezler. 2018. [Can neural machine translation be improved with user feedback?](#) In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 3 (Industry Papers)*, pages 92–105, New Orleans - Louisiana. Association for Computational Linguistics.
- Yebin Lee, Imseong Park, and Myungjoo Kang. 2024. [Fleur: An explainable reference-free evaluation metric for image captioning using a large multimodal model](#). *arXiv preprint arXiv:2406.06004*.
- Joel Lehman, Jeff Clune, Dusan Misevic, Christoph Adami, Lee Altenberg, Julie Beaulieu, Peter J Bentley, Samuel Bernard, Guillaume Beslon, David M Bryson, et al. 2020. The surprising creativity of digital evolution: A collection of anecdotes from the evolutionary computation and artificial life research communities. *Artificial life*, 26(2):274–306.
- Wang Luo, Haoran Li, Zicheng Zhang, Congying Han, Jiayu Lv, and Tiande Guo. 2024. [Sambo-r1: Shifts-aware model-based offline reinforcement learning](#). *arXiv preprint arXiv:2408.12830*.

- Tasnim Mohiuddin, Philipp Koehn, Vishrav Chaudhary, James Cross, Shruti Bhosale, and Shafiq Joty. 2022. Data selection curriculum for neural machine translation. *arXiv preprint arXiv:2203.13867*.
- Khanh Nguyen, Hal Daumé III, and Jordan Boyd-Graber. 2017. Reinforcement learning for bandit neural machine translation with simulated human feedback. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1464–1474, Copenhagen, Denmark. Association for Computational Linguistics.
- OpenAI, :, Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec Helyar, Aleksander Madry, and Alex Beutel. 2024. *Openai o1 system card*.
- Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. 2022. Training language models to follow instructions with human feedback. In *Proceedings of the 36th International Conference on Neural Information Processing Systems, NIPS '22*, Red Hook, NY, USA. Curran Associates Inc.
- Alexander Pan, Kush Bhatia, and Jacob Steinhardt. 2022. The effects of reward misspecification: Mapping and mitigating misaligned models. *ArXiv*, abs/2201.03544.
- 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.
- Romain Paulus, Caiming Xiong, and Richard Socher. 2017. A deep reinforced model for abstractive summarization. *ArXiv*, abs/1705.04304.
- Miguel Moura Ramos, Tomás Almeida, Daniel Vareta, Filipe Azevedo, Sweta Agrawal, Patrick Fernandes, and André F. T. Martins. 2025. Fine-grained reward optimization for machine translation using error severity mappings.
- Marc’Aurelio Ranzato, Sumit Chopra, Michael Auli, and Wojciech Zaremba. 2015. Sequence level training with recurrent neural networks. *arXiv preprint arXiv:1511.06732*.
- Marc’Aurelio Ranzato, Sumit Chopra, Michael Auli, and Wojciech Zaremba. 2016. Sequence level training with recurrent neural networks.
- Ricardo Rei, Marcos Treviso, Nuno M Guerreiro, Chrysoula Zerva, Ana C Farinha, Christine Maroti, José GC De Souza, Taisiya Glushkova, Duarte M Alves, Alon Lavie, et al. 2022. Cometkiwi: Ist-unbabel 2022 submission for the quality estimation shared task. *arXiv preprint arXiv:2209.06243*.
- Manon Revel, Matteo Cargnelutti, Tyna Eloundou, and Greg Leppert. 2025. Seal: Systematic error analysis for value alignment. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 27599–27607.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Thibault Sellam, Dipanjan Das, and Ankur P Parikh. 2020. Bleurt: Learning robust metrics for text generation. *arXiv preprint arXiv:2004.04696*.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Y Wu, et al. 2024. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*.
- Guangming Sheng, Chi Zhang, Zilingfeng Ye, Xibin Wu, Wang Zhang, Ru Zhang, Yanghua Peng, Haibin Lin, and Chuan Wu. 2024. Hybridflow: A flexible and efficient rlhf framework. *arXiv preprint arXiv:2409.19256*.
- Mohammad Shoeybi, Mostofa Patwary, Raul Puri, Patrick LeGresley, Jared Casper, and Bryan Catanzaro. 2019. Megatron-lm: Training multi-billion parameter language models using model parallelism. *arXiv preprint arXiv:1909.08053*.
- Artem Sokolov, Julia Kreutzer, Stefan Riezler, and Christopher Lo. 2016. Stochastic structured prediction under bandit feedback. In *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc.
- Nisan Stiennon, Long Ouyang, Jeffrey Wu, Daniel Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul F Christiano. 2020. Learning to summarize with human feedback. *Advances in neural information processing systems*, 33:3008–3021.
- Shaomu Tan and Christof Monz. 2025. Remedy: Learning machine translation evaluation from human preferences with reward modeling.
- Qwen Team. 2024. *Qwen2.5: A party of foundation models*.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Jonathan Uesato, Ramana Kumar, Victoria Krakovna, Tom Everitt, Richard Ngo, and Shane Legg. 2020. Avoiding tampering incentives in deep rl via decoupled approval. *arXiv preprint arXiv:2011.08827*.

Chaojun Wang and Rico Sennrich. 2020. [On exposure bias, hallucination and domain shift in neural machine translation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3544–3552, Online. Association for Computational Linguistics.

Lilian Weng. 2024. [Reward hacking in reinforcement learning](#). *lilianweng.github.io*.

Yixuan Weng, Minjun Zhu, Fei Xia, Bin Li, Shizhu He, Shengping Liu, Bin Sun, Kang Liu, and Jun Zhao. 2022. Large language models are better reasoners with self-verification. *arXiv preprint arXiv:2212.09561*.

Sam Wiseman and Alexander M. Rush. 2016. [Sequence-to-sequence learning as beam-search optimization](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1296–1306, Austin, Texas. Association for Computational Linguistics.

Mitchell Wortsman, Gabriel Ilharco, Samir Ya Gadre, Rebecca Roelofs, Raphael Gontijo-Lopes, Ari S Morcos, Hongseok Namkoong, Ali Farhadi, Yair Carmon, Simon Kornblith, et al. 2022. Model soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference time. In *International conference on machine learning*, pages 23965–23998. PMLR.

Lijun Wu, Fei Tian, Tao Qin, Jianhuang Lai, and Tie-Yan Liu. 2018. [A study of reinforcement learning for neural machine translation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3612–3621, Brussels, Belgium. Association for Computational Linguistics.

Haoran Xu, Young Jin Kim, Amr Sharaf, and Hany Hassan Awadalla. 2023. A paradigm shift in machine translation: Boosting translation performance of large language models. *arXiv preprint arXiv:2309.11674*.

Haoran Xu, Young Jin Kim, Amr Sharaf, and Hany Hassan Awadalla. 2024a. [A paradigm shift in machine translation: Boosting translation performance of large language models](#). In *The Twelfth International Conference on Learning Representations*.

Haoran Xu, Amr Sharaf, Yunmo Chen, Weiting Tan, Lingfeng Shen, Benjamin Van Durme, Kenton Murray, and Young Jin Kim. 2024b. Contrastive preference optimization: Pushing the boundaries of llm performance in machine translation. *arXiv preprint arXiv:2401.08417*.

A Subtitle Case

Source: 你将一个视频中的多个语音识别文本逐条翻译成英文。输入为一个json格式，key为序号，value为待翻译的语音文本，一共有10个文本。示例如下：
输入：{"1": "xxx", "2": "xxx"}

输出：{"1": "xxx", "2": "xxx"}

现在请根据上述要求完成如下片段的翻译，输出一共10个翻译后的结果，需要和输入一一对应。直接输出翻译后的英文，不要进行任何解释。输入：{"1": "牛老师惊恐的说：你大的要来了？我点点头：要来了，我瞒不住了。", "2": "其实我是一名鬼杀队队员。牛老师大怒：你是不是有病？我都这样了，你还不放过我？", "3": "而且还要 COS 鬼灭之刃。我叹了口气：You KNOW NOTHING，牛牛桑，鬼灭之刃其实不是虚构的，而是一部纪录片。", "4": "世界上是真的有鬼，也是真的有鬼杀队员。", "5": "如果今天不是我及时出现，大爷们就要白发人送黑发人了。", "6": "刘老师说：那问题是我刚染了黄毛呀，不是黑发人呀。", "7": "我说：我说的是你的脸，印堂发黑去非洲怎么着也得当一个王妃。", "8": "之所以你认为这个世界上没有鬼存在，都是我们这些鬼杀队队员的功劳，把你这些普通人类保护的太好。", "9": "当然，也正因为跟我们跟鬼残酷的战斗了几百年，以至于人才凋落，很多鬼都管不过来，才让你被鬼白嫖。", "10": "牛老师说：你快拉倒吧，你这人全身上下只有嘴硬，就你还杀鬼，你只会鬼吹灯。"}
输出：

Strong Translation: {"1": "Teacher Niu said in horror: Your big one is coming? I nodded: It's coming, I can't hide it anymore.", "2": "Actually, I am a member of the Demon Slayer Corps. Teacher Niu was furious: Are you sick? I'm already like this, and you still won't let me go?", "3": "And I have to COS Demon Slayer. I sighed: You KNOW NOTHING, Niu Niu-san, Demon Slayer is not fictional, it's actually a documentary.", "4": "There really are demons in the world, and there really are demon slayers.", "5": "If I hadn't appeared in time today, the old men would have had to send off the young.", "

"6": "Teacher Liu said: But the problem is I just dyed my hair yellow, I'm not a young person with black hair.", "7": "I said: I'm talking about your face, with a darkened forehead, you could at least be a princess in Africa.", "8": "The reason you think there are no demons in this world is all thanks to us demon slayers, protecting you ordinary humans too well.", "9": "Of course, it's also because we've been fighting brutally with demons for hundreds of years, leading to a decline in talent, and many demons are left unchecked, which is why you got taken advantage of by demons.", "10": "Teacher Niu said: Stop it, you're all talk and no action, you killing demons? You only know how to blow hot air."}]"

Source: "请你将一个视频中的多个语音识别文本逐条翻译成英文。输入为一个json格式，key为序号，value为待翻译的语音文本，一共有10个文本。示例如下：

输入: {"1": "xxx", "2": "xxx"}

输出: {"1": "xxx", "2": "xxx"}

现在请根据上述要求完成如下片段的翻译，输出一共10个翻译后的结果，需要和输入一一对应。直接输出翻译后的英文，不要进行任何解释。

输入: {"1": "啊哈。但是你听到后面跟的是个降调的感觉。听啊，", "2": "啊，再来一下。", "3": "当然了，倒水最主要声音最响就是这个叭叭这种声音。", "4": "贯穿始终有一种大珠小珠落玉盘的感觉，它呢，其实就是这种入水的声音。就是这，", "5": "是吧？", "6": "我应该找 B.Boxer 酋长来录这一期的。", "7": "总之呢，倒一杯水有三种声音诶，吧吧。", "8": "但是在科学上，你不能用得儿劲儿蹦来描述谁的声音。听着不太聪明的样子。", "9": "声音的本质是物体的振动。那这三种声音都是啥在振动呢？那个声调的怎？", "10": "N 是杯子里面水面上方空气柱的振动。也就是我们小学二年级学过的海姆霍兹共振。"}
输出: "

Strong Translation: {"1": "Aha. But you hear that the following part has a falling tone. Listen,", "2": "Ah, let's do it again.", "3": "Of course, the loudest sound when pouring water is this kind of 'ba ba' sound.", "4": "Throughout, there is a feeling of big and small beads falling on a jade plate. It is actually the sound of water entering. It's this,", "5": "Right?", "6": "I should have invited B.Boxer Chief to record this episode.", "

"7": "In short, there are three kinds of sounds when pouring a glass of water, de de ba de ba.", "8": "But scientifically, you can't describe someone's sound as 'de er jin er beng'. It doesn't sound very smart.," "9": "The essence of sound is the vibration of objects. So what are these three sounds vibrating? What about the pitch?," "10": "N is the vibration of the air column above the water surface inside the cup. This is the Helmholtz resonance we learned in second grade."}]"

B GPT-4o Evaluation Prompt

GPT-4o Evaluation Prompt 如下:

请你以公正的评判者的身份，结合参考译文，对翻译文本的效果进行评估。请保持一致的评估标准。你需要按照下面的维度对翻译文本进行评估，根据每个维度的要求描述，给每个维度打出一个1-4的分数。维度1: 准确性: 词句+专有名词翻译是否准确

- 1分: 翻译严重偏离原义，3处以上词语存在事实性错误或语义矛盾。
- 2分: 翻译部分偏离原义，2-3处词语存在事实性错误或语义矛盾。
- 3分: 1处词语存在事实性错误或语义矛盾，或是无词语翻译错误，仅有语义错误。
- 4分: 译文完全准确，无事实性错误，且不影响用户感知。

补充说明: 1. 如果翻译存在漏翻，如果未导致句子主要含义发生变化，则只扣完整性，不扣准确性分数; 2. 原文的错别字、断句不需要结合上下文进行纠正，直接按照原文文本翻译即可，翻译正确即不扣准确性分数。3. 增添部分逻辑合理，不做准确性扣分。维度2: 完整性: 是否漏翻关键词句

- 1分: 句子主要信息和逻辑关系严重缺失，主从句不完整，数词、动词及动词受者大量缺失，达3处以上遗漏问题，严重影响原义理解。
- 2分: 句子主要信息部分缺失主句和从句（非状语、非补语）部分未保留主谓语句，关键成分（数词、动词及其受者）未完整保留，达2-3处遗漏问题，不影响整体理解。
- 3分: 句子主要信息部分缺失，主句和从句（非状语、非补语）部分未保留主谓语句，关键成分（数词、动词及其受者）仅小部分遗漏，仅1处遗漏问题，不影响整体理解。
- 4分: 内容完整性: 确保句子的主要信息和逻辑关系100%完整。主从句完整性: 主句和从句（非状语、非补语）必须保留主谓语句。仅允许进行浓缩或代词替换。其他成分处理: 宾语、补语、插入语和修饰语可以进行压缩性意译或删除。关键成分保留: 数词、动词及其受者不可省略，必须完整保留。

补充说明：1) 请严格对比原文，看主要信息是否均翻译完成，可以接受对语义重复的修饰语做删除，不扣完整性分数；2) 原文的错别字、断句错误导致的问题，不扣完整性分数。维度3：连贯性：上下文衔接是否连贯

- 1分：语法错误频发，上下文转折生硬，逻辑断裂，严重影响用户理解。
- 2分：多处语序混乱或句式杂糅，上下文衔接略生硬，关联词使用不准确，需反复理解
- 3分：整体连贯，仅存在部分的生硬衔接，存在过长修饰语或出现少量语句重复。
- 4分：语句衔接自然，表达顺畅，符合口语表达习惯。

补充说明：1) 连词、独立语、过渡性语句等的省略，如果对句子内容产生影响，则只扣完整性分数，连贯性不扣分；2) 因为原文错别字、断句导致的问题，不对此维度扣分。维度4：风格一致性：上下文叙述风格是否一致

- 1分：翻译内容与视频语境不符，3处以上语法结构、用词与视频风格不一致，使用过于突兀。
- 2分：翻译内容在书面语和口语之间切换，2-3处用词和句法结构不符合语境。
- 3分：上下文风格大体一致，自然流畅，仅1处用词和句法结构不符合语境。
- 4分：上下文内容呈现出统一的口语或书面语的特征，整体内容风格明显。

补充说明：首先判断视频文本的场景为书面还是口语场景，后判断语法结构、用词是否符合场景要求。

我们会提供原文、参考译文、待评估的翻译文本，你需要依照前面的每个维度标准，分别给出翻译文本在各个维度的打分并解释理由。请注意，评分应尽可能严格。如果样例在某个维度上同时符合低分和高分的评价标准，应优先选择低分，而非较高的分数。只有当样例完全符合高分标准时，才能给予高分。你的回答以 JSON 格式输出结果，严格按照如下格式返回：

```
{ "维度1": { "解释": "解释", "打分": 1 } }
【以下是原文】 {origin_text}
【以下是参考译文】 {reference_text}
【以下是待评估的翻译文本】 {translate_text}
```

C GPT4o vs Human

To assess the reliability of GPT-4o in evaluating translation quality, we conduct a comparative analysis between human and model-based judgments on the subtitle dataset. Specifically, we invite three human annotators—none of whom are involved in this study—to independently rate the translation outputs. The three volunteers have received a solid higher education, and we provided them with detailed instructions on the annotation task, including scoring criteria and evaluation standards. The average scores assigned by these human evaluators are then compared to those produced by GPT-4o.

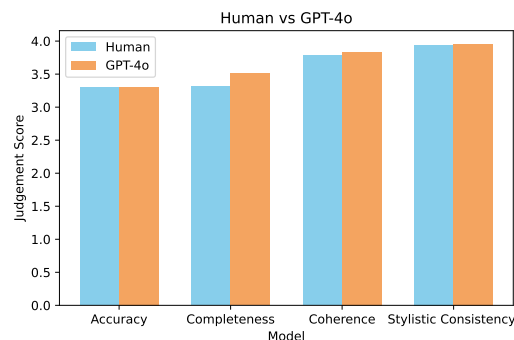


Figure 5: Human Evaluation vs GPT-4o Evaluation

As presented in Figure 5, GPT-4o’s assessments show a high degree of consistency with human judgments, suggesting that it can serve as a reliable proxy for human evaluation in this context. Given the strong alignment and the practicality of automated evaluation at scale, we adopt GPT-4o as the primary evaluator for all translation quality assessments conducted in this work.

D Implementation Details

Below are some specific details of our model training.

Our RM training framework is built on Megatron. We use the Qwen2.5-72B-Chat model as the initialization. The training is conducted with a batch size of 256, using a cosine learning rate scheduler with an initial learning rate of $5e-6$. All models are trained on 64 Huawei’s Ascend 910B NPUs.

Our RL training framework is based on the Verl framework. We use the Qwen2.5-7B-Chat model as the initialization for RL training. During training, we configure a batch size of 16 and perform 16 rollouts per prompt using the GRPO algorithm. The learning rate is initialized at $1e-8$, and a cosine scheduler with warm-up is applied toward the final iteration. Sampling is conducted with a temperature of 1.0, and the maximum generation length is limited to 1,024 tokens. The KL penalty coefficient β is set to 0, effectively removing the KL constraint relative to the reference policy. The PPO clipping range ϵ is fixed at 0.2. All models are trained for one epoch using 8 NVIDIA H800 80G GPUs.

For the subtitle task, the RM training data consists of 60,000 samples, and the RL training data contains 60,000 samples. For the WMT task, the RM training data includes about 50,000 samples, and the RL training data comprises about 30,000 samples.

Subtitle	Accuracy	Completeness	Coherence	Stylistic Consistency	Avg	COMETkiwi	BLEU
RIVAL-Qual	1.71e-10	1.78e-07	0.04	1.76e-06	8.37e-10	9.35e-07	-
WMT	Accuracy	Completeness	Coherence	Stylistic Consistency	Avg	COMETkiwi	BLEU
RIVAL-Qual	6.32e-9	2.58e-06	1.78e-07	3.63e-04	7.93e-10	9.48e-07	8.14e-07
RIVAL-Qual+Quant	2.57e-8	1.35e-08	9.73e-03	2.95e-05	4.32e-9	4.85e-07	1.79e-07

Table 3: P value for RIVAL.

E Statistical Significance Testing

To more comprehensively evaluate the performance of our RIVAL method, we incorporated **p-values** alongside conventional metrics and mean values. The consistent statistical significance observed across all experimental results—as clearly demonstrated in the accompanying table—strongly attests to the robustness of our approach. These findings not only provide compelling evidence that our method substantially outperforms the baseline but also highlight its reliability and generalizability under varied conditions.

F MAE vs MSE case

Source: 请将以下内容翻译为英文: 他强调,遵照秘鲁宪法和国际人权文件,秘鲁政府不会强迫个人接受生殖健康和计划生育方案的服务。

Strong Translation: He emphasized that in keeping with the country’s Constitution and with the international documents on human rights, the Government of Peru did not coerce individuals who availed themselves of the services of the reproductive health and family planning programme.

Weak Translation: He emphasized that, in accordance with Peruvian constitution and international human rights documents, the Peruvian government would not force individuals to accept services related to reproductive health and family planning.
strong bleu score = 1

weak bleu score = 0.6404765601431773

mae strong bleu score = 0.9680025577545166

mae weak bleu score = 0.6667302250862122

mse strong bleu score = 1.941943645477295

mse weak bleu score = 1.9110959768295288

Source: 请将以下内容翻译为中文: 235. At the first regular session of 1999 of the Administrative Committee on Coordination, in April, the members concluded that, to meet the challenges of globalization, the United Nations system needed to cooperate more effectively with the private sector and civil society, as well as with Governments.

Strong Translation: 235. 行政协调会成员

在1999年4月第一届常会上总结指出,为应付全球化的挑战,联合国系统必须更有效地与私营部门、民间社会和各国政府合作。

Weak Translation: 在1999年4月召开的协调委员会第一次定期会议上,成员们得出结论认为,为了应对全球化带来的挑战,联合国系统需要更加有效地与私营部门、民间社会以及各国政府进行合作。

strong bleu score = 1

weak bleu score = 0.10885796200376416

mae strong bleu score = 1.09672212600708

mae weak bleu score = 0.2919744849205017

mse strong bleu score = 1.3286265134811401

mse weak bleu score = 1.2083359956741333

	Quantitative Error	Qualitative Acc
MAE	0.19	99.5
MSE	0.93	97.0

Table 4: Comparison Between MAE and MSE for Training the RM.

G Related works

Reinforcement Learning for Machine Translation Early works employed RL to bridge the gap between training on token-level log-likelihoods and evaluating on corpus-level, non-differentiable metrics like BLEU (Ranzato et al., 2016; Edunov et al., 2018), while also mitigating exposure bias in autoregressive decoders (Wang and Sennrich, 2020). Subsequent research repurposed RL for domain or user adaptation by replacing engineered objectives with human feedback in human-in-the-loop workflows (Sokolov et al., 2016; Nguyen et al., 2017). Recent methods add reward shaping (Wu et al., 2018; Goyal et al., 2019; Devidze et al., 2022), injecting intermediate signals alongside final-step BLEU, yet partial or token-wise BLEU still fails to capture subtle semantic and contextual differences, making it a poor choice for fine-grained reward design.

Nowadays, ground-breaking RL-based LLMs typified by OpenAI’s O1 (OpenAI et al., 2024) and DeepSeek-R1 (DeepSeek-AI et al., 2025)

deliver state-of-the-art results on diverse benchmarks, drawing considerable attention from the research community. Extending this progress, R1-Translator (He et al., 2025) is the first to incorporate human-aligned chain-of-thought reasoning into general machine translation through RL. MT-R1-Zero (Feng et al., 2025) pushes the paradigm further by introducing zero-shot RL with a hybrid rule-and-metric reward for translation tasks. Meanwhile, ReMedy (Tan and Monz, 2025) reconceptualizes machine-translation evaluation as a reward-modeling problem, training on pairwise preferences to avoid the noise of absolute quality ratings. Complementary work (Ramos et al., 2025) replaces sparse sentence-level feedback with severity-weighted, token-level rewards computed by XCOMET, yielding a more informative learning signal.

Reward Hacking in Reinforcement Learning Reward hacking (Weng, 2024; Everitt et al., 2017; Everitt and Hutter, 2019; di Langosco et al., 2021; Pan et al., 2022) occurs when a reinforcement-learning (RL) agent exploits flaws in its reward signal and optimizes the proxy rather than the true task goal. This vulnerability stems from the practical difficulty of defining a complete, noise-free reward in complex or partially observed environments. Documented cases span robotics (Lehman et al., 2020; Christiano et al., 2017a), Atari-style games (Bansal et al., 2017; Gleave et al., 2019), and RLHF pipelines for language models (Paulus et al., 2017; Denison et al., 2024), with agents looping around waypoints, rewriting unit tests, or gaming automatic metrics to inflate scores without real progress. Recent research therefore frames reward hacking as a safety and alignment challenge and proposes layered defences: (i) reward-side fixes—adversarial or ensemble reward models, potential-based shaping, reward capping, and “trip-wire” signals that trigger intervention (Amodei et al., 2016); (ii) policy-side safeguards—model look-ahead, adversarial blinding, or indifference techniques to block exploits (Uesato et al., 2020); and (iii) post-hoc monitoring—decoupled human approval, anomaly detection on trajectories, and systematic data audits to surface misalignment early (Pan et al., 2022; Revel et al., 2025). Together, these strategies emphasise adaptive, multi-stage protection rather than any single remedy.