

# Judge as A Judge: Improving the Evaluation of Retrieval-Augmented Generation through the Judge-Consistency of Large Language Models

Shuliang Liu<sup>1\*</sup>, Xinze Li<sup>1\*</sup>, Zhenghao Liu<sup>1†</sup>, Yukun Yan<sup>2</sup>,  
Cheng Yang<sup>3</sup>, Zheni Zeng<sup>2</sup>, Zhiyuan Liu<sup>2</sup>, Maosong Sun<sup>2</sup>, Ge Yu<sup>1</sup>

<sup>1</sup>School of Computer Science and Engineering, Northeastern University, China

<sup>2</sup>Department of Computer Science and Technology, Institute for AI, Tsinghua University, China

Beijing National Research Center for Information Science and Technology, China

<sup>3</sup>School of Computer Science, Beijing University of Posts and Telecommunications

## Abstract

Retrieval-Augmented Generation (RAG) has proven its effectiveness in alleviating hallucinations for Large Language Models (LLMs). However, existing automated evaluation metrics cannot fairly evaluate the outputs generated by RAG models during training and evaluation. LLM-based judgment models provide the potential to produce high-quality judgments, but they are highly sensitive to evaluation prompts, leading to inconsistencies when judging the output of RAG models. This paper introduces the Judge-Consistency (ConsJudge) method, which aims to enhance LLMs to generate more accurate evaluations for RAG models. Specifically, ConsJudge prompts LLMs to generate different judgments based on various combinations of judgment dimensions, utilizes the judge-consistency to evaluate these judgments, and selects the chosen and rejected judgments for DPO training. Our experiments show that ConsJudge can effectively provide more accurate judgments for optimizing RAG models across various RAG models and datasets. Further analysis reveals that judgments generated by ConsJudge have a high agreement with the superior LLM. All codes are available at <https://github.com/OpenBMB/ConsJudge>.

## 1 Introduction

Retrieval-Augmented Generation (RAG) (Guu et al., 2020; Lewis et al., 2020; Asai et al., 2023; Shi et al., 2023) has proven effective in mitigating hallucinations (Elazar et al., 2021; Ji et al., 2023; Shuster et al., 2021; Huang et al., 2023) in Large Language Models (LLMs). RAG retrieves relevant knowledge from the knowledge base and incorporates the external information as the input context (Ram et al., 2023), benefiting various knowledge-intensive tasks (Trivedi et al., 2023; Izacard et al., 2022; He et al., 2021). Existing

\* indicates equal contribution.

† indicates corresponding author.

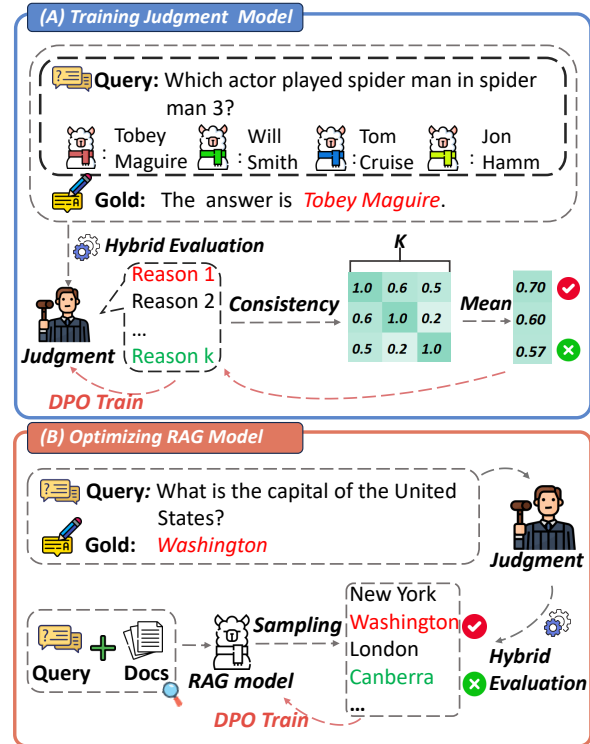


Figure 1: The Framework of ConsJudge. It enhances the judgment capabilities of LLMs and benefits the training process of RAG models.

studies typically employ automated evaluation metrics to assess the outputs of RAG systems during both training (Li et al., 2024b) and evaluation (Lin et al., 2023; Gao et al., 2023b) phases. These metrics primarily focus on string-level exact matching, which is less effective for determining whether the LLM-generated responses align with the ground truth (Gao et al., 2023a; Saad-Falcon et al., 2024), posing a challenge for RAG systems.

To enable more accurate evaluation of RAG models, some works incorporate LLMs as judgment models to assess the quality of generated responses (Saad-Falcon et al., 2024; Friel et al., 2024; Adlakha et al., 2023), relying on their human-

equivalent evaluation capabilities (Chiang and Lee, 2023a; Zheng et al., 2023; Sottana et al., 2023). These judgment models introduce specific evaluation dimensions, such as hallucination and comprehensiveness, to prompt LLMs to verify whether the generated responses align with the facts in the retrieved documents and whether all relevant information has been properly extracted and integrated (Jin et al., 2024; Jacovi et al., 2025; Zhu et al., 2024). However, LLM-based judgment models are highly sensitive to the design of evaluation prompts (Zhou et al., 2023; Liu et al., 2024), which may lead to inconsistencies in judgments when different evaluation dimensions are employed.

In this paper, we introduce Judge-Consistency (ConsJudge), a method that enhances LLM-based judgment models to generate more accurate evaluations for RAG models in a self-improvement framework. ConsJudge incorporates specialized evaluation dimensions and employs a multiple-choice selection strategy for evaluation modeling. Additionally, ConsJudge leverages the Direct Preference Optimization (DPO) method (Rafailov et al., 2023) to enhance the judgment capabilities of LLMs, while also implementing a “judge as a judge” mechanism during choosing preference pairs of judgments. Specifically, ConsJudge encourages LLMs to generate judgment results based on various combinations of judgment dimensions, evaluates the quality of these judgments using judge-consistency, and selects the chosen and rejected judgments for DPO training. This approach enhances the performance of the LLM-based judgment model without necessitating distillation from more powerful LLMs (Zhang et al., 2025).

Our experiments demonstrate the effectiveness of the ConsJudge method as a reward model for optimizing the RAG model, resulting in significant improvements over vanilla LLMs. Further analysis shows that ConsJudge helps LLMs select appropriate evaluation dimensions for assessing response quality by incorporating the agreements and consistency across different evaluation dimensions during optimization. Additionally, compared to other baselines, ConsJudge exhibits higher judgment consistency with the superior LLM, GLM-4-plus (Du et al., 2022), across various RAG evaluation datasets, further highlighting its effectiveness in optimizing LLMs to produce more accurate judgments on diverse RAG tasks.

## 2 Related Work

Retrieval-Augmented Generation (RAG) has proven its effectiveness in various tasks, such as open-domain question answering (Trivedi et al., 2023), dialogue systems (Cai et al., 2019), and code generation (Li et al., 2025). By integrating external knowledge into the input context (Ram et al., 2023), Large Language Models (LLMs) can alleviate hallucination issues and produce more accurate responses (Asai et al., 2024). To evaluate the performance of LLM responses, most existing RAG models rely on automatic metrics like Exact-Match during both evaluation (Lin et al., 2023; Gao et al., 2023b) and training (Li et al., 2024b). However, these metrics may fail to offer a fair assessment when the LLM-generated responses are lengthy or semantically similar to the ground truth but not character-matched (Gao et al., 2023a).

Recently, LLMs, such as ChatGPT (OpenAI, 2023), have demonstrated human-equivalent performance (Chiang and Lee, 2023a; Zheng et al., 2023; Sottana et al., 2023) and are now commonly used as judgment models to assess generated responses in various tasks (Chen et al., 2023a,b; An et al., 2023; Chan et al., 2023). These studies typically prompt LLMs to evaluate generated responses based on specific evaluation dimensions (Chiang and Lee, 2023b). However, using APIs of these closed-source models incurs significant costs and reduces reproducibility, as the models may evolve behind the API (Gu et al., 2024). To address this, some researchers are turning to open-source LLMs as alternatives, using judgments from closed-source LLMs to fine-tune open-source models and improve their evaluation capabilities (Wang et al., 2023b; Zheng et al., 2023).

Moreover, RAG models also employ LLMs as judges to assess the quality of generation, focusing on the relevance and faithfulness of the responses (Saad-Falcon et al., 2024; Friel et al., 2024; Adlakha et al., 2023). These models typically design prompts to instruct LLMs to determine whether the generated response aligns with the facts in the retrieved document and whether all relevant information has been fully extracted and integrated (Jin et al., 2024; Jacovi et al., 2025; Zhu et al., 2024). However, LLM-based evaluations are highly sensitive to prompt designs (Zhou et al., 2023; Liu et al., 2024), making the judgments inconsistent when using different dimensions for evaluating RAG responses. To mitigate this is-

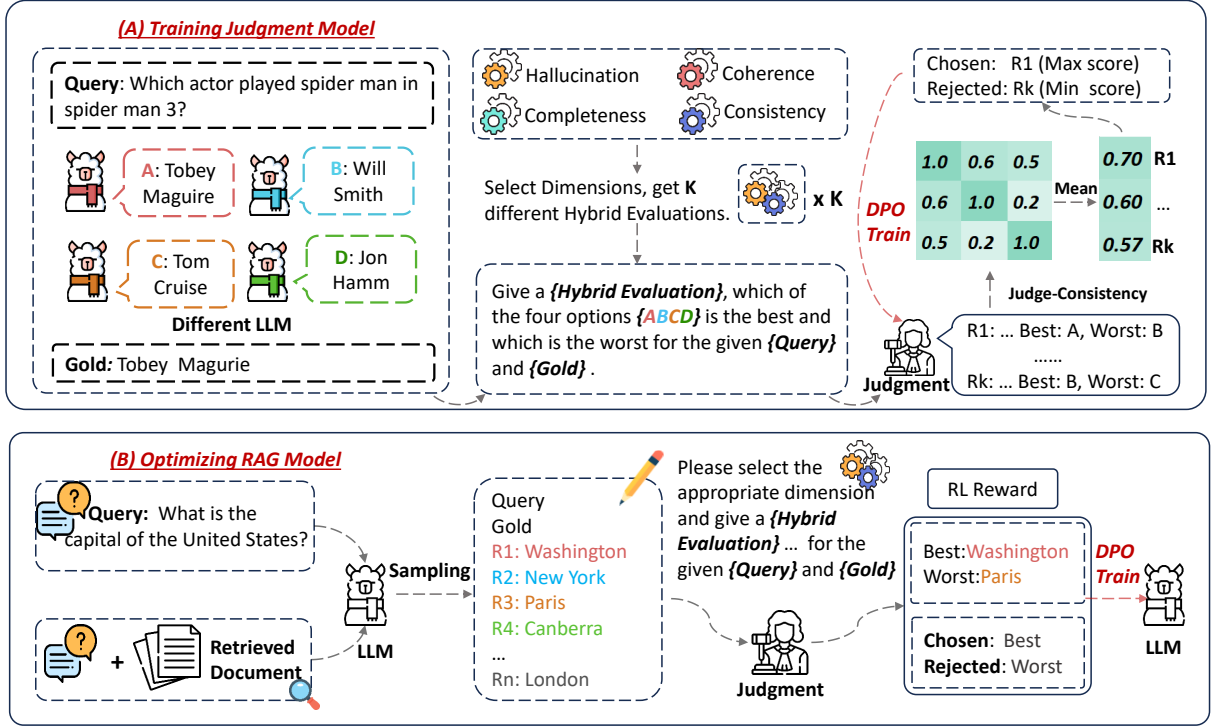


Figure 2: The Framework of Our ConsJude Method.

sue, ConsJude introduces a judge-consistency approach to self-improve the judgment performance of LLMs in RAG systems, avoiding distillation from larger-scale LLMs (Zhang et al., 2025).

### 3 Methodology

In this section, we introduce the Judge-Consistency (ConsJude) method, which optimizes Large Language Models (LLMs) as judgment models to evaluate the effectiveness of the Retrieval-Augmented Generation (RAG) system. First, we discuss the evaluation methods used in ConsJude (Sec.3.1). Next, we present how the ConsJude method is utilized to optimize the judgment model, improving its judgment accuracy (Sec.3.2). Finally, we apply the ConsJude method as the reward model to optimize the RAG system (Sec. 3.3).

#### 3.1 Preliminary of RAG Evaluation Methods

Recent work, such as LLM-as-a-Judge (Zheng et al., 2023), typically regards LLMs as judgment models for rating responses in various NLP tasks. These methods use specially designed prompts to employ LLMs in evaluating generated responses across different dimensions, such as hallucination and accuracy (Li et al., 2024a). In this section, we describe the evaluation dimensions and modeling methods used in ConsJude.

**Evaluation Dimensions.** Some studies (Zhu et al., 2024; Zhang et al., 2025) evaluate the generation quality of RAG models based on multiple criteria. ConsJude follows these approaches by designing different evaluation prompts for four key dimensions: hallucination, completeness, coherence, and semantic consistency.

*Hallucination.* Hallucination refers to the inclusion of information in the response that contradicts the ground truth. This dimension aims to detect whether the generated responses contain factual errors due to hallucinations (Xu et al., 2024b).

*Completeness.* Completeness evaluates whether the generated responses contain as much relevant information as possible from the ground truth. This dimension primarily focuses on identifying whether the responses omit some key points from the ground truth answers.

*Coherence.* Coherence evaluates whether the responses are logically consistent and whether the language flows fluently between sentences. This dimension is primarily concerned with ensuring that the responses are both coherent and fluent.

*Semantic Consistency.* Semantic consistency checks whether the generated response is semantically aligned with the ground truth answer, rather than simply matching it lexically. This dimension helps avoid misjudging responses that differ from

the ground truth in terms of tokens but share the same meaning with ground truth answers.

**Evaluation Modeling.** For evaluation modeling, our ConsJudge method uses a multiple-choice selection approach (Gu et al., 2024). In this method, LLMs evaluate all candidate responses and choose the best or the worst based on different evaluation dimensions. This facilitates RAG training using the judgment model (Sec. 3.3).

Existing methods typically rely on Pointwise Evaluation or Listwise Comparison for candidate response evaluation. Pointwise Evaluation directly prompts LLMs to score each candidate based on predefined evaluation dimensions. However, this method fails to capture the differences between responses, leading to evaluation bias (Kim et al., 2023; Wang et al., 2023a). In contrast, Listwise Comparison prompts LLMs to evaluate an entire list of candidate responses and rank them (Niu et al., 2024; Yan et al., 2024), allowing for a more comprehensive evaluation (Li et al., 2024a). The ConsJudge model adopts a multiple-choice selection method, which performs Listwise Comparison to evaluate all candidate responses.

### 3.2 Training Evaluation Models through Judge-Consistency

Although LLMs have demonstrated effectiveness in evaluating their own responses, judgment models may still suffer from issues such as Position Bias, Verbosity Bias, and Evaluation Metric Bias, which can compromise the quality of judgments (Zheng et al., 2023; Chen et al., 2024; Li et al., 2024a). To address these challenges, we propose the Judge-Consistency method to optimize the judgment model  $\mathcal{M}$  based on the consistency of judgments across different evaluation dimensions. This process self-improves the judgment model by selecting more suitable evaluation dimensions, ultimately allowing for more precise assessments.

**Evaluation of LLM Responses.** To optimize the judgment model, we begin by selecting  $m$  LLMs and sampling responses from each using different temperatures. Next, we randomly select one response  $y$  from each LLM, forming a response set  $Y = \{y_1, \dots, y_m\}$ . We then combine the four evaluation dimensions introduced in Sec. 3.1, generating  $k$  distinct hybrid evaluation aspects:

$$\mathcal{I} = \{I_1, \dots, I_k\}, \quad (1)$$

where  $I_i$  represents a hybrid evaluation aspect, which could be a single evaluation dimension or

a combination of multiple dimensions. For each evaluation aspect  $I_i \in \mathcal{I}$ , we create an evaluation prompt, yielding  $k$  distinct prompts:

$$\mathcal{P} = \{P^1, \dots, P^k\}, \quad (2)$$

where  $P^i$  is the  $i$ -th evaluation prompt. The judgment model  $\mathcal{M}$  then generates a judgment result  $r_i$  for evaluating LLM-generated responses  $Y$ , based on the  $i$ -th evaluation aspect  $I_i$ :

$$r_i = \mathcal{M}(P^i, q, y^*, Y), \quad (3)$$

where  $y^*$  is the ground truth for question  $q$ . The judgment result  $r_i$  includes both the best ( $y^+$ ) and the worst ( $y^-$ ) selections from the candidate responses  $Y = \{y_1, \dots, y_m\}$ , along with chain-of-thoughts (Wei et al., 2022) for the judgment. By utilizing all evaluation prompts in  $\mathcal{P}$ , we obtain  $k$  judgment results, denoted as  $R = \{r_1, \dots, r_k\}$ .

**Judge Consistency Evaluation.** We follow the previous work (Li et al., 2023) to introduce a ‘‘Judge as a judge’’ approach that evaluates the consistency of judgments across different prompts.

Specifically, after obtaining the judgment results  $R = \{r_1, \dots, r_k\}$ , we use a text embedding model  $\text{Emb}(\cdot)$  to compute the similarity score between the  $i$ -th judgment  $r_i$  and all other judgment results  $R$ . The average of these similarity scores provides the consistency score  $S_i$  for the judgment  $r_i$ :

$$S_i = \frac{1}{k} \sum_{j=1}^k \cos(\text{Emb}(r_i), \text{Emb}(r_j)). \quad (4)$$

Judgments exhibiting higher consistency scores are considered positive results, while those with lower consistency scores are interpreted as negative results, indicating potential judge bias.

**Judgment Model Optimization.** We then optimize the judgment model  $\mathcal{M}$  to better select the appropriate evaluation aspect from the set  $\mathcal{I}$  to make more accurate judgments.

To achieve this, we treat the judgment with the highest consistency score as a positive judgment  $r^+$ , and the judgment with the lowest score as the negative one  $r^-$ . We collect the instance  $(q, y^*, r^+, r^-)$  to form the training dataset  $\mathcal{T}$ . The judgment model  $\mathcal{M}$  then uses prompts  $\mathcal{P}$  for evaluating the quality of responses  $Y$  and is optimized to assign a higher probability to the positive judgment  $r^+$  than to the negative judgment  $r^-$ . This is accomplished through the Direct Preference Optimization (DPO) method (Rafailov et al., 2023):

$$\mathcal{L} = -\mathbb{E}_{(q, y^*, r^+, r^-) \sim \mathcal{T}} [\log \sigma(\beta \log \frac{\mathcal{M}(r^+ | q, y^*)}{\mathcal{M}^{\text{ref}}(r^+ | q, y^*)} - \beta \log \frac{\mathcal{M}(r^- | q, y^*)}{\mathcal{M}^{\text{ref}}(r^- | q, y^*)})], \quad (5)$$



where  $r^+, r^- \in R$ , and  $\beta$  is a hyperparameter.  $\mathcal{M}^{\text{ref}}$  denotes the reference model, which remains fixed during training.

These judgment results in  $R$  are generated using different evaluation prompts  $\mathcal{P}$  that combine various evaluation dimensions. ConsJudge optimizes the judgment model  $\mathcal{M}$  to reproduce the positive judgment  $r^+$  that shares the most consistency score with others, making the judgment model  $\mathcal{M}$  select more appropriate dimensions to evaluate the response quality of LLMs.

### 3.3 Applying Judgment Models to Optimize Retrieval-Augmented Generation Systems

To evaluate the effectiveness of judgment model  $\mathcal{M}$ , we use it as the reward model and apply DPO to optimize the RAG system (Li et al., 2024b).

For a given query  $q$ , the current RAG system typically utilizes a dense retriever model to retrieve the Top- $n$  relevant documents  $\mathcal{D} = \{d_1, \dots, d_n\}$  from the external knowledge bases. The generation model (Gen) then samples outputs  $\tilde{y}$ , either with or without the retrieved documents  $\mathcal{D}$ :

$$\begin{aligned}\tilde{y} &\sim \text{Gen}(\mathcal{D} \oplus q), \\ \tilde{y} &\sim \text{Gen}(q),\end{aligned}\tag{6}$$

where  $\oplus$  is the concatenation operation. Then, we collect all sampled responses  $\tilde{y}$  in the set  $\tilde{Y}$  and utilize the judgment model  $\mathcal{M}$  to generate the judgment result  $r_{\text{all}}$  based on all evaluation dimensions:

$$r_{\text{all}} = \mathcal{M}(P_{\text{all}}, q, y^*, \tilde{Y}),\tag{7}$$

where  $P_{\text{all}}$  indicates the evaluation prompt that involves all evaluation dimensions. Based on  $r_{\text{all}}$ , we select the best response  $\tilde{y}^+$  and worst response  $\tilde{y}^-$  from  $\tilde{Y}$ , respectively. This allows us to construct a preference dataset  $(q, \mathcal{P}, \tilde{y}^+, \tilde{y}^-)$  to optimize the generation model (Gen) via DPO training. If the optimized generation model (Gen) demonstrates improved performance on the downstream RAG tasks, this indicates that the judgment model  $\mathcal{M}$  provides more precise judgments, effectively serving as the reward for optimizing the RAG system.

## 4 Experimental Methodology

This section describes the datasets, evaluation metrics, baselines, and implementation details used in our experiments. More implementation details are shown in Appendix A.6.

**Datasets.** We describe the datasets used for training ConsJudge and RAG training and evaluation.

Task	Dataset	Total
Open-Domain QA	NQ (2019)	2,837
	TriviaQA (2017)	5,359
	MARCOQA (2016)	1,000
Factoid QA	ASQA (2022)	958
Multi-Hop QA	HotpotQA (2018)	5,600
Dialogue	WoW (2019)	1,000

Table 1: Data Statistics of the RAG Evaluation Datasets.

**ConsJudge Training.** For training ConsJudge, we collect 11 knowledge-intensive tasks from previous works (Chung et al., 2022; Izacard et al., 2022) to collect 73,831 instances and 3,886 instances to construct both training and development sets.

**RAG Training & Evaluation.** To retrieve documents for constructing the RAG datasets, we use BGE-large (Xiao et al., 2023) with the MS MARCO V2.1 (Bajaj et al., 2016) corpus. During RAG training, we collect seven datasets from Li et al. (2024b) and randomly sample 20,805 samples for the training set and 1,400 samples for the development set. For RAG evaluation, we select knowledge-intensive tasks from prior work (Li et al., 2024b; Xu et al., 2024a), including open-domain QA tasks (NQ (Kwiatkowski et al., 2019), TriviaQA (Joshi et al., 2017), MARCO QA (Bajaj et al., 2016)), multi-hop QA (HotpotQA (Yang et al., 2018)), factoid QA (ASQA (Stelmakh et al., 2022)), and dialogue tasks (WoW (Dinan et al., 2019)). The data statistics are shown in Table 1.

**Evaluation Metrics.** For tasks with longer outputs, automated evaluation metrics, such as ROUGE, cannot evaluate the quality of outputs fairly, which has been proven by previous work (Gao et al., 2023a; Zhang et al., 2024). Thus, we adopt the LLM-as-a-Judge method (Zhang et al., 2024), which employs GLM-4-plus<sup>1</sup> for evaluation MARCO QA and WoW. Besides, we use StringEM as the evaluation metric for the ASQA dataset. For other evaluation tasks, we evaluate performance using Accuracy. The prompt using GLM-4-plus to evaluate is shown in Appendix A.7.

**Baselines.** In our experiments, we compare ConsJudge with three judgment models, including the Raw Metric model and two LLM-based judgment models. For the Raw Metric model, we utilize the automatic evaluation metrics, ROUGE-L and Accuracy, as the judgment model to optimize the RAG system. Specifically, the Raw Metric model uses ROUGE-L for MARCO QA, Yahoo!QA and WikiQA datasets, and also uses Accuracy for the

<sup>1</sup><https://open.bigmodel.cn/>

Generator	Reward	NQ (acc)	HotpotQA (acc)	TriviaQA (acc)	ASQA (str-em)	MARCOQA (llm)	WoW (llm)	Avg.
MiniCPM (2.4B)	Raw Metric	46.14	30.09	80.03	33.44	84.75	85.48	59.99
	Llama3-8B	<b>47.23</b>	29.64	80.40	<b>35.77</b>	86.00	85.98	60.84
	w/ SFT	47.02	29.55	80.33	34.60	<b>86.35</b>	85.98	60.64
	w/ ConsJudge	47.02	<b>30.45</b>	<b>80.80</b>	35.68	86.16	<b>86.13</b>	<b>61.04</b>
	Qwen2.5-14B	47.30	28.96	80.03	34.95	85.59	87.49	60.72
	w/ SFT	47.02	27.64	79.80	33.51	<b>85.90</b>	<b>87.51</b>	60.23
Llama3 (8B)	w/ ConsJudge	<b>48.01</b>	<b>30.84</b>	<b>80.69</b>	<b>36.45</b>	85.73	87.21	<b>61.49</b>
	Raw Metric	48.96	36.95	86.83	41.55	84.80	82.66	63.63
	Llama3-8B	46.63	32.84	85.13	40.69	88.15	88.31	63.63
	w/ SFT	47.16	35.95	86.10	40.37	87.46	87.97	64.20
	w/ ConsJudge	<b>48.78</b>	<b>37.54</b>	<b>88.26</b>	<b>42.44</b>	<b>88.25</b>	<b>88.87</b>	<b>65.69</b>

Table 2: Overall Performance of RAG Models Optimized Using Different Judgment Models. The **best** result is highlighted. We implement the generators of RAG models by using MiniCPM-2.4B and Llama3-8B.

remaining datasets, which is the same as previous work (Li et al., 2024b). Additionally, we employ two LLM-based judgment model baselines: Vanilla LLM and SFT. The Vanilla LLM method directly uses the LLM as the judgment model and then leverages the evaluation prompts to ask it to produce the judgments. The SFT method further fine-tunes LLMs based on judgment results generated by a superior LLM, GLM-4-plus, which has been used in previous work (Zhang et al., 2025) to improve the judgment performance of LLMs.

**Implementation Details.** In our experiments, we leverage LoRA (Hu et al., 2022) for efficient training LLMs. We set max epoch to 3, learning rate to  $5e-5$ , and the warmup ratio to 0.1. For the generation model in the RAG system, we employ the MiniCPM-2.4B (Hu et al., 2024) and Llama3-8B-Instruct (Touvron et al., 2023) as the generation models. For the judgment model, we use Llama3-8B-Instruct and Qwen2.5-14B-Instruct (Bai et al., 2023) as the backbone models. While training the judgment model, we synthesize 8 different hybrid evaluation aspects for generating the judgment results. We use MiniCPM-Embedding<sup>2</sup> to assess the similarity among judgments.

## 5 Evaluation Result

In this section, we first show the performance of ConsJudge by regarding it as a reward model to optimize the RAG model. Then we conduct the ablation studies to show the effectiveness of different modules in ConsJudge. Subsequently, we evaluate the judgment quality generated by ConsJudge and explore the consistency of judgments across different evaluation dimensions. Finally, case stud-

ies are conducted. More experiments are shown in Appendix A.2 and Appendix A.5.

### 5.1 Overall Performance

In this experiment, we treat these judgment models as reward models for optimizing the RAG model (Li et al., 2024b) and evaluate their effectiveness by examining the RAG performance.

As shown in Table 2, we compare ConsJudge with three judgment models, including Raw Metric, vanilla LLMs, and SFT. Specifically, the Raw Metric model uses an automatic metric as the reward model, while vanilla LLMs and SFT rely on LLMs as reward models. In our experiments, we use Llama3-8B and Qwen2.5-14B to implement the judgment models, and both MiniCPM-2.4B and Llama3-8B to build the RAG model.

Overall, ConsJudge outperforms all baseline models across different RAG models, demonstrating its effectiveness in judging sampled responses of RAG models. Compared to the Raw Metric model, LLM-based judgment models achieve better optimization performance for the RAG model, even when the Raw Metric model directly optimizes RAG models to align with final evaluation metrics. It shows that LLMs have the ability to produce high-quality judgments, benefiting the training and evaluation processes of RAG systems. In comparison to vanilla LLMs, the SFT model uses labels produced by GLM-4-plus for training, while ConsJudge introduces a judge-consistency based optimization method that trains LLMs without relying on additional training signals. ConsJudge not only outperforms the SFT model but also achieves more improvements over the Raw Metric model, highlighting the effectiveness of our judge-as-a-judge mechanism in enhancing LLMs through a

<sup>2</sup><https://huggingface.co/openbmb/MiniCPM-Embedding>

Reward Model	NQ	HotpotQA	TriviaQA	Avg.
<i>Llama3-8B-Instruct</i>				
ConsJudge	<b>46.70</b>	<b>30.18</b>	<b>80.13</b>	<b>52.34</b>
w/o Consistency	44.55	28.91	77.96	50.47
w/o Query	44.66	28.86	75.48	49.67
w/o Ground Truth	33.35	22.27	67.61	41.07
<i>Qwen2.5-14B-Instruct</i>				
ConsJudge	<b>48.00</b>	<b>30.95</b>	<b>81.89</b>	<b>53.61</b>
w/o Consistency	47.73	30.67	81.47	53.29
w/o Query	47.41	30.61	81.29	53.10
w/o Ground Truth	38.03	25.86	75.78	46.56

Table 3: Ablation Study. We use MiniCPM-2.4B to implement the RAG model.

self-improvement approach. Notably, the evaluation results demonstrate that the advantages of ConsJudge can be extended to various RAG scenarios and judgment models of different scales.

## 5.2 Ablation Study

This experiment conducts ablation studies to investigate the contribution of different modules in ConsJudge.

As shown in Table 3, we evaluate the judgment performance of three variants of the ConsJudge model by reranking the sampled responses from vanilla RAG models and calculating the accuracy of the top-1 ranked responses. The three models compared in the experiments are: ConsJudge w/o Consistency, ConsJudge w/o Query, and ConsJudge w/o Ground Truth. Specifically, ConsJudge w/o Consistency randomly selects the chosen and rejected responses for DPO training. Both ConsJudge w/o Query and ConsJudge w/o Ground Truth models are evaluated by removing the query and ground truth from the input prompts, respectively.

Compared to ConsJudge w/o Consistency, ConsJudge achieves a higher accuracy score. This demonstrates that the judge-consistency method enhances the ability of LLMs to select higher-quality responses, thereby benefiting the RAG training process. Furthermore, removing the query or the ground truth answer from the evaluation prompts results in a performance decrease for ConsJudge, although the performance gap is narrowed when using larger-scale LLMs. This suggests that both the query and ground truth help the LLM-based judgment models produce more comprehensive evaluations. Additionally, LLMs of a larger scale can leverage their parametric knowledge to assess the quality of RAG responses.

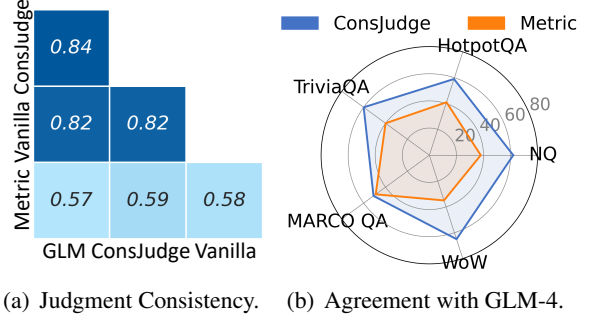


Figure 3: Judge Agreement Evaluation. We analyze the agreements of different judgment models (Figure 3(a)), and use GLM-4-plus to evaluate the judge quality of different models (Figure 3(b)). GLM and Metric denotes the GLM-4-plus and Raw Metric models. Vanilla LLM and ConsJudge are implemented with Qwen2.5-14B.

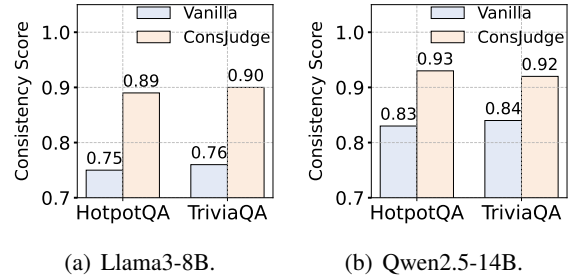


Figure 4: Judgment Consistency of Vanilla LLMs and ConsJudge. We use both vanilla LLMs and ConsJudge to show the judgment consistency among all hybrid evaluation aspects used to train ConsJudge.

## 5.3 The Judgment Quality of ConsJudge

In this section, we first present the judge agreement across different models. Then, we analyze the consistency of judgments generated by different models based on different evaluation dimensions.

**Judge Agreement.** First, we sample 100 queries from each of the datasets—NQ, HotpotQA, TriviaQA, MARCO QA, and WoW—to construct the evaluation dataset. We then collect responses from different models and ask judgment models to evaluate these responses. As shown in Figure 3, we use four judgment models: GLM-4-plus (Du et al., 2022), ConsJudge, vanilla LLM, and Raw Metrics, to select the best response for each query.

Figure 3(a) presents the judge agreement between different models. Among all judgment models, the Raw Metric shows the lowest agreement with the others, highlighting that string matching alone does not provide high-quality judgment for evaluation. Notably, ConsJudge not only demonstrates the highest agreement with the superior

<b>Case 1:</b> What are the virulence factors of anthrax?
<b>Ground Truth:</b> <i>Bacillus anthracis</i> .
<b>Choice A:</b> The virulence factors of anthrax are a group of proteins produced by the <i>Bacillus anthracis</i> bacterium that contribute to its ability to cause disease in humans and animals. <b>Choice B:</b> <i>anthracis</i> .
<b>Choice C:</b> lethal factor, edema factor and antiphagocytic factor.
<b>Choice D:</b> lethal factor, antiphagocytic factor and other factors.
<b>Raw Metric (ROUGE-L):</b> Choice B(0.667) > Choice A(0.143) > Choice C(0) = Choice D(0)
<b>Zero-Shot:</b> ...choice C and D are more complete as they specifically mention the names of the virulence factor... <b>Best answer :C</b> .
<b>SFT:</b> ... from the perspective of Coherence... <b>Best answer: C</b> ...
<b>ConsJudge:</b> from the perspective of semantic consistency, completeness, and coherence... <b>Best answer: A</b> ...
<b>Case 2:</b> How much is 55% of 40 is greater than 4/5 of 25? Options: (A) 29. (B) 776. <b>(C) 2</b> . (D) 12. (E) 99.
<b>Ground Truth:</b> <b>(C)</b>
<b>Choice A:</b> According to the information of question, 55% of 40 is 22, which is <b>2</b> greater than 4/5 of 25, which equals 20.
<b>Choice B:</b> <b>(C)</b> , 55% of 40 is greater than 4/5 of 25 by <b>1</b> .
<b>Choice C:</b> <b>(C)</b> . 55% of 44 is smaller than 4/5 of 25 (which is 20) by <b>2</b> .
<b>Choice D:</b> 55% of 40 is smaller than 4/5 of 25.
<b>Raw Metric (Accuracy):</b> Choice A(0), Choice B(1), Choice C(1), Choice D(0)
<b>Zero-Shot:</b> ...there is no hallucination in any of the choices.... <b>Choice C is the best answer</b> ...
<b>SFT:</b> ... from the perspective of Completeness ..., choice C given the answer: C, ... <b>Best answer: C</b> ...
<b>ConsJudge:</b> ...from the perspective of Semantic Consistency... <b>Best answer: A</b> ...

Table 4: Case Studies. If the phrase aligns with the facts, it is highlighted in **Brown**, while inconsistencies are highlighted in **Purple**. Correct results are marked in **Red**, whereas incorrect results are highlighted in **Blue**.

LLM, GLM-4-plus, but also conducts the agreement with GLM-4-plus when evaluating the judgments produced by vanilla LLM. This illustrates that ConsJudge conducts more consistent judgments with GLM-4-plus. Furthermore, we evaluate the judge agreement with GLM-4-plus of both Raw Metric and ConsJudge models in Figure 3(b). Across all datasets, ConsJudge achieves higher agreement scores with GLM-4-plus, demonstrating its effectiveness in optimizing LLMs to generate more accurate judgments in various scenarios.

**Judgment Consistency.** We next randomly sample 1,000 queries each from HotpotQA and TriviaQA to construct a dataset for evaluating judgment consistency. We then ask both vanilla LLM and ConsJudge to perform judgments for each query, using different hybrid evaluation aspects, and compute the consistency scores of these judgments.

As shown in Figure 4, the consistency scores of judgments generated by ConsJudge outperform those generated by vanilla LLM across both datasets. These results indicate that ConsJudge achieves higher consistency with different evaluation dimensions, demonstrating its ability to comprehensively incorporate tailored evaluation dimensions to produce reliable evaluation results. Notably, the advantages of ConsJudge are consistent across the LLMs of different scales, further illustrating its generalization ability.

## 5.4 Case Study

In Table 4, we present two cases to illustrate the effectiveness of the ConsJudge method.

In the first case, the Raw Metric selects the response with the most matched phrases to the ground truth, which results in unreliable evaluation for long ground truths (Gao et al., 2023a). Additionally, the vanilla LLM prioritizes factual correctness, while the SFT model focuses on coherence aspects, leading to incorrect judgments. In contrast, ConsJudge incorporates multiple evaluation dimensions, such as “Semantic”, “Completeness” and “Coherence”, allowing for more effective and consistent evaluations. This demonstrates the effectiveness of ConsJudge in choosing appropriate dimensions during evaluating the responses, which helps to provide a more comprehensive assessment.

In the second case, all models select the Choice B and Choice C, which contains the ground truth answer “(C)”. It illustrates that these models, even for these LLM-based judgment models, focus more on matching responses with the ground truth answer to conduct the judgment. However, the reasoning processes in these responses of both Choice B and Choice C contain errors. In contrast, ConsJudge shifts its attention from matching the ground truth answer to the “Semantic” dimension, showing its effectiveness.

## 6 Conclusion

This paper proposes the ConsJudge method, which enhances the judgment ability of LLMs in a self-improvement method. It prompts LLMs to generate different judgments based on various combinations of judgment dimensions and utilizes the consistency between judgments to enhance LLMs



to generate more accurate judgments. Our experimental results show that ConsJudge has the ability to choose appropriate evaluation dimensions to generate reliable and comprehensive judgments.

## Limitation

Although ConsJudge demonstrates convincing performance in enhancing the judgment ability of LLMs, it has some limitations. First, the judgment results are often sensitive to the evaluation prompts, and ConsJudge relies on the consistency of these prompts to improve LLMs’ judgment capability. Thus, the design of high-quality evaluation prompts that encompass comprehensive evaluation dimensions remains an underexplored area. Second, ConsJudge utilizes the well-performing embedding model, MiniCPM-Embedding, to calculate the similarity score between judgment results based on different evaluation aspects. However, the effectiveness of this similarity estimation method may limit the overall performance of ConsJudge. Further exploration of more fine-grained approaches, such as incorporating the matching signals of ground truth answers, could enhance its effectiveness.

## Ethics Statement

Our experiment involves employing human evaluators to evaluate the outputs generated by the RAG model. We ensure that we have carefully distributed the data from our experiment to the human evaluators, ensuring it is strictly used for research purposes and does not contain any content that violates ethical standards.

## Acknowledgments

This work is partly supported by the Natural Science Foundation of China under Grant (No. 62206042 and No. 62137001), CCF-zhipu Large Model Innovation Fund (No. 202403), the Joint Funds of Natural Science Foundation of Liaoning Province (No. 2023-MSBA-081), and the Fundamental Research Funds for the Central Universities under Grant (No. N2416012). This work is also supported by the AI9Stars community.

## References

Vaibhav Adlakha, Parishad BehnamGhader, Xing Han Lu, Nicholas Meade, and Siva Reddy. 2023. [Evaluating correctness and faithfulness of instruction-](#)

[following models for question answering](#). *ArXiv preprint*, abs/2307.16877.

Shourya Aggarwal, Divyanshu Mandowara, Vishwajeet Agrawal, Dinesh Khandelwal, Parag Singla, and Dinesh Garg. 2021. [Explanations for commonsenseqa: New dataset and models](#). In *Proceedings of ACL*, pages 3050–3065.

Chenxin An, Shansan Gong, Ming Zhong, Xingjian Zhao, Mukai Li, Jun Zhang, Lingpeng Kong, and Xipeng Qiu. 2023. [L-eval: Instituting standardized evaluation for long context language models](#). In *Proceedings of ACL*, pages 14388–14411.

Akari Asai, Zeqiu Wu, Yizhong Wang, Avirup Sil, and Hannaneh Hajishirzi. 2023. [Self-rag: Learning to retrieve, generate, and critique through self-reflection](#). In *Proceedings of ICLR*.

Akari Asai, Zexuan Zhong, Danqi Chen, Pang Wei Koh, Luke Zettlemoyer, Hannaneh Hajishirzi, and Wen tau Yih. 2024. [Reliable, adaptable, and attributable language models with retrieval](#). *ArXiv preprint*, abs/2403.03187.

Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, and Xiaodong Deng. 2023. [Qwen technical report](#). *ArXiv preprint*, abs/2309.16609.

Payal Bajaj, Daniel Campos, Nick Craswell, Li Deng, Jianfeng Gao, Xiaodong Liu, Rangan Majumder, Andrew McNamara, Bhaskar Mitra, Tri Nguyen, et al. 2016. [Ms marco: A human generated machine reading comprehension dataset](#). *ArXiv preprint*, abs/1611.09268.

Jonathan Berant, Andrew Chou, Roy Frostig, and Percy Liang. 2014. [Semantic parsing on freebase from question-answer pairs](#). In *Proceedings of EMNLP*, pages 1533–1544.

Deng Cai, Yan Wang, Wei Bi, Zhaopeng Tu, Xiaojiang Liu, Wai Lam, and Shuming Shi. 2019. [Skeleton-to-response: Dialogue generation guided by retrieval memory](#). In *Proceedings of NAACL-HLT*, pages 1219–1228.

Chi-Min Chan, Weize Chen, Yusheng Su, Jianxuan Yu, Wei Xue, Shanghang Zhang, Jie Fu, and Zhiyuan Liu. 2023. [Chateval: Towards better llm-based evaluators through multi-agent debate](#). In *Proceedings of ICLR*.

Guiming Hardy Chen, Shunian Chen, Ziche Liu, Feng Jiang, and Benyou Wang. 2024. [Humans or llms as the judge? a study on judgement biases](#). In *Proceedings of EMNLP*, pages 8301–8327.

Junying Chen, Xidong Wang, Ke Ji, Anningzhe Gao, Feng Jiang, Shunian Chen, Hongbo Zhang, Dingjie Song, Wenya Xie, Chuyi Kong, et al. 2023a. [Huatuogpt-ii, one-stage training for medical adaptation of llms](#). *ArXiv preprint*, abs/2311.09774.

Xinyun Chen, Maxwell Lin, Nathanael Schärli, and Denny Zhou. 2023b. [Teaching large language models to self-debug](#). In *Proceedings of ICLR*.

- Zhiyu Chen, Wenhu Chen, Charese Smiley, Sameena Shah, Iana Borova, Dylan Langdon, Reema Moussa, Matt Beane, Ting-Hao Huang, Bryan R Routledge, et al. 2021. [Finqa: A dataset of numerical reasoning over financial data](#). In *Proceedings of EMNLP*, pages 3697–3711.
- Cheng-Han Chiang and Hung-Yi Lee. 2023a. [Can large language models be an alternative to human evaluations?](#) In *Proceedings of ACL*, pages 15607–15631.
- Cheng-Han Chiang and Hung-yi Lee. 2023b. [A closer look into automatic evaluation using large language models](#). In *Proceedings of EMNLP Findings*, pages 8928–8942.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, and Yunxuan Li. 2022. [Scaling instruction-finetuned language models](#). *ArXiv preprint*, abs/2210.11416.
- Emily Dinan, Stephen Roller, Kurt Shuster, Angela Fan, Michael Auli, and Jason Weston. 2019. [Wizard of wikipedia: Knowledge-powered conversational agents](#). In *Proceedings of ICLR*.
- Zhengxiao Du, Yujie Qian, Xiao Liu, Ming Ding, Jiezhong Qiu, Zhilin Yang, and Jie Tang. 2022. [Glm: General language model pretraining with autoregressive blank infilling](#). In *Proceedings of ACL*, pages 320–335.
- Yanai Elazar, Nora Kassner, Shauli Ravfogel, Abhिलाषा Ravichander, Eduard Hovy, Hinrich Schütze, and Yoav Goldberg. 2021. [Measuring and improving consistency in pretrained language models](#). *Transactions of the Association for Computational Linguistics*, pages 1012–1031.
- Angela Fan, Yacine Jernite, Ethan Perez, David Grangier, Jason Weston, and Michael Auli. 2019. [ELI5: Long form question answering](#). In *Proceedings of ACL*, pages 3558–3567.
- Robert Friel, Masha Belyi, and Atindriyo Sanyal. 2024. [Ragbench: Explainable benchmark for retrieval-augmented generation systems](#). *ArXiv preprint*, abs/2407.11005.
- Tianyu Gao, Howard Yen, Jiatong Yu, and Danqi Chen. 2023a. [Enabling large language models to generate text with citations](#). In *Proceedings of EMNLP*, pages 6465–6488.
- Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, Meng Wang, and Haofen Wang. 2023b. [Retrieval-augmented generation for large language models: A survey](#). *ArXiv preprint*, abs/2312.10997.
- Mor Geva, Daniel Khashabi, Elad Segal, Tushar Khot, Dan Roth, and Jonathan Berant. 2021. [Did aristotle use a laptop? a question answering benchmark with implicit reasoning strategies](#). In *Transactions of the Association for Computational Linguistics*, pages 346–361.
- Jiawei Gu, Xuhui Jiang, Zhichao Shi, and Hexiang Tan. 2024. [A survey on llm-as-a-judge](#). *ArXiv preprint*, abs/2411.15594.
- Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Ming-Wei Chang. 2020. [Realm: Retrieval-augmented language model pre-training](#). In *Proceedings of ICML*, pages 3929–3938.
- Junxian He, Graham Neubig, and Taylor Berg-Kirkpatrick. 2021. [Efficient nearest neighbor language models](#). In *Proceedings of EMNLP*, pages 5703–5714.
- Edward J Hu, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. 2022. [Lora: Low-rank adaptation of large language models](#). In *Proceedings of ICLR*.
- Shengding Hu, Yuge Tu, Xu Han, Chaoqun He, Ganqu Cui, and Xiang Long. 2024. [Minicpm: Unveiling the potential of small language models with scalable training strategies](#). *ArXiv preprint*, abs/2404.06395.
- Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, and Bing Qin. 2023. [A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions](#). *ACM Transactions on Information Systems*.
- Gautier Izacard, Patrick Lewis, Maria Lomeli, Lucas Hosseini, Fabio Petroni, Timo Schick, Jane Dwivedi-Yu, Armand Joulin, Sebastian Riedel, and Edouard Grave. 2022. [Atlas: few-shot learning with retrieval augmented language models](#). *ArXiv preprint*, abs/2208.03299.
- Alon Jacovi, Andrew Wang, Chris Alberti, Connie Tao, Jon Lipovetz, Kate Olszewska, Lukas Haas, Michelle Liu, Nate Keating, Adam Bloniarz, et al. 2025. [The facts grounding leaderboard: Benchmarking llms’ ability to ground responses to long-form input](#). *ArXiv preprint*, abs/2501.03200.
- Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. 2023. [Survey of hallucination in natural language generation](#). *ACM Computing Surveys*.
- Qiao Jin, Bhuwan Dhingra, Zhengping Liu, William Cohen, and Xinghua Lu. 2019. [Pubmedqa: A dataset for biomedical research question answering](#). In *Proceedings of EMNLP-IJCNLP*, pages 2567–2577.
- Zhuoran Jin, Hongbang Yuan, Tianyi Men, Pengfei Cao, Yubo Chen, Kang Liu, and Jun Zhao. 2024. [Rag-rewardbench: Benchmarking reward models in retrieval augmented generation for preference alignment](#). *ArXiv preprint*, abs/2412.13746.
- Mandar Joshi, Eunsol Choi, Daniel Weld, and Luke Zettlemoyer. 2017. [Triviaqa: A large scale distantly supervised challenge dataset for reading comprehension](#). In *Proceedings of ACL*, pages 1601–1611.

- Seungone Kim, Jamin Shin, Yejin Cho, Joel Jang, Shayne Longpre, Hwaran Lee, Sangdoo Yun, Seongjin Shin, Sungdong Kim, James Thorne, et al. 2023. [Prometheus: Inducing fine-grained evaluation capability in language models](#), 2024. *ArXiv preprint*, abs/2310.08491.
- Seungone Kim, Juyoung Suk, Shayne Longpre, Bill Yuchen Lin, Jamin Shin, Sean Welleck, Graham Neubig, Moontae Lee, Kyungjae Lee, and Minjoon Seo. 2024. [Prometheus 2: An open source language model specialized in evaluating other language models](#). In *Proceedings of EMNLP*, pages 4334–4353.
- Tomáš Kočiský, Jonathan Schwarz, Phil Blunsom, Chris Dyer, Karl Moritz Hermann, Gábor Melis, and Edward Grefenstette. 2018. [The narrativeqa reading comprehension challenge](#). *Transactions of the Association for Computational Linguistics*, pages 317–328.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, and Chris Alberti. 2019. [Natural questions: a benchmark for question answering research](#). In *Proceedings of ACL*, pages 452–466.
- Patrick S. H. Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. [Retrieval-augmented generation for knowledge-intensive nlp tasks](#). In *Proceedings of NeurIPS*, pages 9459–9474.
- Haitao Li, Qian Dong, Junjie Chen, Huixue Su, Yujia Zhou, Qingyao Ai, Ziyi Ye, and Yiqun Liu. 2024a. [Llms-as-judges: A comprehensive survey on llm-based evaluation methods](#). *ArXiv preprint*, abs/2412.05579.
- Siheng Li, Cheng Yang, Zesen Cheng, Lemao Liu, Mo Yu, Yujia Yang, and Wai Lam. 2023. [Large language models can self-improve in long-context reasoning](#). In *Proceedings of EMNLP*, pages 1051–1068.
- Xinze Li, Sen Mei, Zhenghao Liu, Yukun Yan, Shuo Wang, Shi Yu, Zheni Zeng, Hao Chen, Ge Yu, and Zhiyuan Liu. 2024b. [Rag-ddr: Optimizing retrieval-augmented generation using differentiable data rewards](#). *ArXiv preprint*, abs/2410.13509.
- Xinze Li, Hanbin Wang, Zhenghao Liu, Shi Yu, Shuo Wang, Yukun Yan, Yukai Fu, Yu Gu, and Ge Yu. 2025. [Building a coding assistant via the retrieval-augmented language model](#). *ACM Transactions on Information Systems*, pages 1–25.
- Xi Victoria Lin, Xilun Chen, Mingda Chen, Weijia Shi, Maria Lomeli, Rich James, Pedro Rodriguez, Jacob Kahn, Gergely Szilvasy, and Mike Lewis. 2023. [Raidit: Retrieval-augmented dual instruction tuning](#). In *Proceedings of ICLR*.
- Wang Ling, Dani Yogatama, Chris Dyer, and Phil Blunsom. 2017. [Program induction by rationale generation: Learning to solve and explain algebraic word problems](#). In *Proceedings of ACL*, pages 158–167.
- Yinhong Liu, Han Zhou, Zhijiang Guo, Ehsan Shareghi, Ivan Vulić, Anna Korhonen, and Nigel Collier. 2024. [Aligning with human judgement: The role of pairwise preference in large language model evaluators](#). In *Proceedings of COLM*.
- Pan Lu, Swaroop Mishra, Tony Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Taffjord, Peter Clark, and Ashwin Kalyan. 2022. [Learn to explain: Multi-modal reasoning via thought chains for science question answering](#). In *Proceedings of NeurIPS*, pages 2507–2521.
- Macedo Maia, Siegfried Handschuh, André Freitas, Brian Davis, Ross McDermott, Manel Zarrouk, and Alexandra Balahur. 2018. [Www’18 open challenge: financial opinion mining and question answering](#). In *proceedings of the the web conference 2018*, pages 1941–1942.
- Alex Mallen, Akari Asai, Victor Zhong, Rajarshi Das, Daniel Khashabi, and Hannaneh Hajishirzi. 2023. [When not to trust language models: Investigating effectiveness of parametric and non-parametric memories](#). In *Proceedings of ACL*, pages 9802–9822.
- Tong Niu, Shafiq Joty, Ye Liu, Caiming Xiong, Yingbo Zhou, and Semih Yavuz. 2024. [Judgerank: Leveraging large language models for reasoning-intensive reranking](#). *ArXiv preprint*, abs/2411.00142.
- R OpenAI. 2023. [Gpt-4 technical report](#). *ArXiv preprint*, abs/2303.08774.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Stefano Ermon, Christopher D. Manning, and Chelsea Finn. 2023. [Direct preference optimization: Your language model is secretly a reward model](#). In *Proceedings of NeurIPS*, pages 53728–53741.
- Ori Ram, Yoav Levine, Itay Dalmedigos, Dor Muhlgay, Amnon Shashua, Kevin Leyton-Brown, , and Yoav Shoham. 2023. [In-context retrieval-augmented language models](#). *Transactions of the Association for Computational Linguistics*.
- Jon Saad-Falcon, Omar Khattab, Christopher Potts, and Matei Zaharia. 2024. [Ares: An automated evaluation framework for retrieval-augmented generation systems](#). In *Proceedings of NAACL*, pages 338–354.
- Weijia Shi, Sewon Min, Michihiro Yasunaga, Minjoon Seo, Rich James, Mike Lewis, Luke Zettlemoyer, and Wen-tau Yih. 2023. [Replug: Retrieval-augmented black-box language models](#). In *Proceedings of NAACL*, pages 8371–8384.
- Kurt Shuster, Spencer Poff, Moya Chen, Douwe Kiela, and Jason Weston. 2021. [Retrieval augmentation reduces hallucination in conversation](#). In *Proceedings of EMNLP Findings*, pages 3784–3803.



- Andrea Sottana, Bin Liang, Kai Zou, and Zheng Yuan. 2023. [Evaluation metrics in the era of gpt-4: Reliably evaluating large language models on sequence to sequence tasks](#). In *Proceedings of EMNLP*, pages 8776–8788.
- Ivan Stelmakh, Yi Luan, Bhuwan Dhingra, and Ming-Wei Chang. 2022. [Asqa: Factoid questions meet long-form answers](#). In *Proceedings of EMNLP*, pages 8273–8288.
- Huqun Suri, Qi Zhang, Wenhua Huo, Yan Liu, and Chunsheng Guan. 2021. [Mediaqa: A question answering dataset on medical dialogues](#). *ArXiv preprint*, abs/2108.08074.
- 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*, abs/2307.09288.
- Harsh Trivedi, Niranjan Balasubramanian, Tushar Khot, and Ashish Sabharwal. 2022. [Musique: Multi-hop questions via single-hop question composition](#). *Transactions of the Association for Computational Linguistics*, pages 539–554.
- Harsh Trivedi, Niranjan Balasubramanian, Tushar Khot, and Ashish Sabharwal. 2023. [Interleaving retrieval with chain-of-thought reasoning for knowledge-intensive multi-step questions](#). In *Proceedings of ACL*, pages 10014–10037.
- Chenglong Wang, Hang Zhou, Kaiyan Chang, Tongran Liu, Chunliang Zhang, Quan Du, Tong Xiao, and Jingbo Zhu. 2023a. [Learning evaluation models from large language models for sequence generation](#). *ArXiv preprint*, abs/2308.04386.
- Yidong Wang, Zhuohao Yu, Wenjin Yao, Zhengran Zeng, Linyi Yang, Cunxiang Wang, Hao Chen, Chaoya Jiang, Rui Xie, Jindong Wang, et al. 2023b. [Pandalm: An automatic evaluation benchmark for llm instruction tuning optimization](#). In *Proceedings of ICLR*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. 2022. [Chain-of-thought prompting elicits reasoning in large language models](#). In *Proceedings of NeurIPS*, pages 24824–24837.
- Lingxiao Wei, He Yan, Xiangju Lu, Junmin Zhu, Jun Wang, and Wei Zhang. 2024. [Cnnsun: Exploring long-context summarization with large language models in chinese novels](#). *ArXiv preprint*, abs/2412.02819.
- Shitao Xiao, Zheng Liu, Peitian Zhang, and Niklas Muennighoff. 2023. [C-pack: Packaged resources to advance general chinese embedding](#). *Preprint*, ArXiv:2309.07597.
- Shicheng Xu, Liang Pang, Mo Yu, Fandong Meng, Huawei Shen, Xueqi Cheng, and Jie Zhou. 2024a. [Unsupervised information refinement training of large language models for retrieval-augmented generation](#). *ArXiv preprint*, abs/2402.18150.
- Ziwei Xu, Sanjay Jain, and Mohan Kankanhalli. 2024b. [Hallucination is inevitable: An innate limitation of large language models](#). *ArXiv preprint*, abs/2401.11817.
- Le Yan, Zhen Qin, Honglei Zhuang, Rolf Jagerman, Xuanhui Wang, Michael Bendersky, and Harrie Oosterhuis. 2024. [Consolidating ranking and relevance predictions of large language models through post-processing](#). *ArXiv preprint*, abs/2404.11791.
- Yi Yang, Wen tau Yih, and Christopher Meek. 2015. [Wikiqa: A challenge dataset for open-domain question answering](#). In *Proceedings of EMNLP*, pages 2013–2018.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. 2018. [Hotpotqa: A dataset for diverse, explainable multi-hop question answering](#). In *Transactions of the Association for Computational Linguistics*, pages 2369–2380.
- Hanning Zhang, Juntong Song, Juno Zhu, Yuanhao Wu, Tong Zhang, and Cheng Niu. 2025. [Rag-reward: Optimizing rag with reward modeling and rlhf](#). *ArXiv preprint*, abs/2501.13264.
- Yue Zhang, Ming Zhang, Haipeng Yuan, Shichun Liu, Yongyao Shi, Tao Gui, Qi Zhang, and Xuanjing Huang. 2024. [Llmeval: A preliminary study on how to evaluate large language models](#). In *Proceedings of AAAI*, pages 19615–19622.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P Xing, et al. 2023. [Judging llm-as-a-judge with mt-bench and chatbot arena](#). In *Proceedings of NeurIPS*, pages 46595–46623.
- Han Zhou, Xingchen Wan, Ivan Vulić, and Anna Korhonen. 2023. [Survival of the most influential prompts: Efficient black-box prompt search via clustering and pruning](#). In *Proceedings of EMNLP Findings*, pages 13064–13077.
- Kunlun Zhu, Yifan Luo, Dingling Xu, Ruobing Wang, Shi Yu, Shuo Wang, Yukun Yan, Zhenghao Liu, Xu Han, Zhiyuan Liu, and Maosong Sun. 2024. [Rageval: Scenario specific rag evaluation dataset generation framework](#). *ArXiv preprint*, abs/2408.01262.



Dataset	Total	Train	Dev
<i>ConsJude Training Data</i>			
FinQA (2021)	4,000	3,800	200
FiQA (2018)	8,000	7,600	400
MeDiaQA (2021)	2,000	1,900	100
PubMedQA (2019)	2,000	1,900	100
ScienceQA (2022)	7,095	6,740	355
NQ (2019)	4,000	3,800	200
ELI5 (2019)	4,000	3,800	200
NarrativeQA (2018)	2,000	1,900	100
PopQA (2023)	14,267	13,554	713
CNNSum (2024)	8,000	7,600	400
MuSiQue (2022)	22,355	21,237	1,118
<i>RAG Training Data</i>			
ECQA (2021)	4,200	4,000	200
MARCOQA (2016)	4,200	4,000	200
Web Questions (2014)	3,778	3,578	200
WikiQA (2015)	1,040	840	200
Yahoo!QA	4,200	4,000	200
StrategyQA (2021)	2,060	1,860	200
AQUA-RAT (2017)	2,727	2,527	200

Table 5: Statistics of the Data Used for ConsJude Training and RAG Training.

## A Appendix

### A.1 License

We show the licenses of the datasets that we use. ELI5 and Yahoo!QA do not report the license of the dataset in the paper or a repository. ELI5 shows its terms of use at website<sup>3</sup>. Yahoo!QA shows its terms of use at website<sup>4</sup>. We summarize the licenses of the remaining datasets as follows:

All of these licenses and agreements allow their data for academic use: NQ (CC BY-SA 3.0 license); FiQA (CC BY-SA 4.0 license); ScienceQA, CNNSum, MuSiQue, Web Questions and HotpotQA (CC BY 4.0 license); WoW (CC BY-NC license); FinQA, MeDiaQA, PubMedQA, PopQA, MARCOQA, WikiQA, and StrategyQA (MIT license); NarrativeQA, AQUA-RAT, TriviaQA and ASQA (Apache 2.0 license); ECQA (CDLA-Sharing 1.0 license);

### A.2 Judgment Accuracy of Prometheus and ConsJude

In this section, we compare with the state-of-the-art judgment model Prometheus (Kim et al., 2024) to verify the effectiveness of ConsJude. Prometheus primarily adopts a Pairwise evaluation approach, which can only evaluate two candidate responses at a time. In contrast, ConsJude focuses on Listwise evaluation for candidate responses and retains

<sup>3</sup><https://facebookresearch.github.io/ELI5/>

<sup>4</sup>[https://tensorflow.google.cn/datasets/community\\_catalog/huggingface/yahoo\\_answers\\_qa](https://tensorflow.google.cn/datasets/community_catalog/huggingface/yahoo_answers_qa)

Reward Model	HotpotQA	Marco QA	WoW	Avg.
Raw Metric	78.09	<b>92.07</b>	91.86	87.34
Prometheus 2 (7B)	77.47	91.00	90.66	86.38
<i>Llama3-8B-Instruct</i>				
Vanilla LLM	79.72	90.90	92.33	87.65
SFT LLM	77.97	90.90	92.46	87.11
ConsJude	<b>80.14</b>	90.97	<b>93.03</b>	<b>88.05</b>

Table 6: Judgment Accuracy of Prometheus and ConsJude.

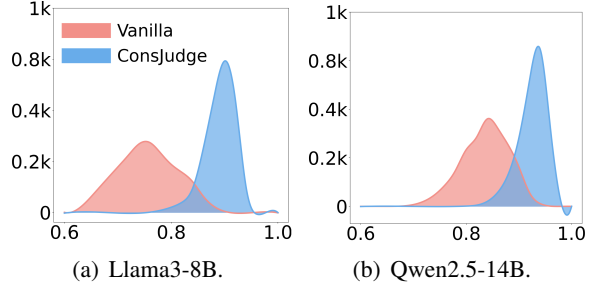


Figure 5: Distribution of Judgment Consistency Score of Both Vanilla LLMs and ConsJude.

the capability for Pairwise evaluation. For a fair comparison, we employ the Pairwise evaluation approach for both ConsJude and Prometheus in our experiments. We randomly sample 200 queries from each of the datasets—HotpotQA, MARCO QA, and WoW—to construct an evaluation dataset and use MiniCPM-2.4B to implement the RAG model. We adopt the LLM-as-a-Judge method, using GLM-4-plus to evaluate the quality of the positive response selected by the judgment model. A higher quality positive response indicates a more accurate judgment.

As shown in Table 6, ConsJude achieves higher judgment accuracy than Prometheus in a Pairwise evaluation approach, further demonstrating its superior generalization ability and effectiveness.

### A.3 Judgment Consistency Score Distribution of Vanilla LLMs and ConsJude

In this section, we further analyze the consistency score distributions of judgments generated by the vanilla LLM and ConsJude based on different hybrid evaluation aspects. To construct a dataset for this analysis, we randomly sample 1,000 queries from both HotpotQA and TriviaQA. We employ both vanilla LLM and ConsJude to generate judgments for each query, using different hybrid evaluation aspects. Then, we refer to Eq. 4 to use MiniCPM-Embedding to compute the consistency scores of these judgments.

As shown in Figure 5, the results demonstrate

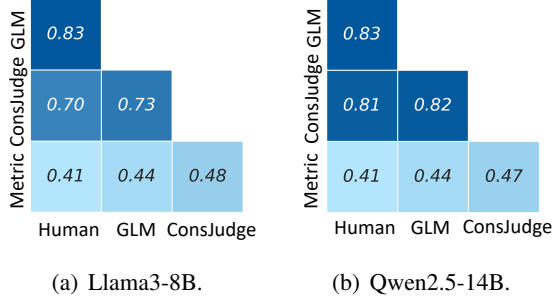


Figure 6: Judge Agreement Evaluation on RAG Training Dataset. We analyze the agreements of different judgment models with Humans. We use Llama3-8B-Instruct and Qwen2.5-14B-Instruct as the backbone models of ConsJudge, respectively.

that ConsJudge not only achieves higher consistency scores but also exhibits a more concentrated distribution of consistency scores compared to the vanilla LLM. Notably, ConsJudge consistently maintains its advantage across LLMs of different scales, highlighting its robust generalization ability.

#### A.4 Judge Agreement Evaluation between Human and ConsJudge

In this section, we further analyze the agreement between different judgment models and humans.

First, we randomly sample 200 queries from the RAG training dataset to assess the agreement in judgment, aiming to evaluate the effectiveness of ConsJudge in assisting the RAG training process. We then collect responses from various models and ask both judgment models and human evaluators to assess these responses. As shown in Figure 6, we use four different judgment methods: Human, GLM-4-plus (Du et al., 2022), ConsJudge, and Raw Metrics, to select the best response for each query. For human annotators, we also provide them with the instructions shown in Table 9 to guide their evaluation.

The Raw Metric exhibits the lowest agreement with the other judgment methods, underscoring that character-matching metrics are inadequate for fairly evaluating generated responses. In contrast, ConsJudge shows higher agreement with both humans and the superior LLM, GLM-4-plus. Furthermore, the agreement between ConsJudge and humans is comparable to that between GLM-4-plus and humans. This indicates that ConsJudge has the ability to produce judgments that are more consistent with human evaluators, making it an effective

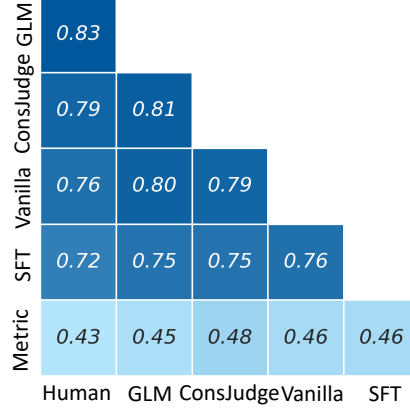


Figure 7: Judge Agreement Evaluation on Non-RAG Tasks. We analyze the agreements of different judgment models with Humans. We use Qwen2.5-14B-Instruct as the backbone model of ConsJudge.

Model	Position-1	Position-2	Position-3	Position-4
Vanilla LLM	73.74	77.27	73.23	<b>73.74</b>
SFT LLM	75.76	<b>81.31</b>	76.77	73.23
ConsJudge	<b>77.27</b>	80.30	<b>81.82</b>	<b>73.74</b>

Table 7: The Position Accuracy of Judgment-Selected Positives Across Different Models. Position- $i$  indicates that the human-annotated best response is placed at the  $i$ -th position in the response sequence.

tool for constructing high-quality preference pairs during training RAG models (Li et al., 2024b).

Additionally, we explore the agreement between ConsJudge and humans in non-RAG scenarios. To construct non-RAG scenarios, we randomly sample 200 examples from the RAG evaluation dataset and employ MiniCPM-2.4B to generate four candidate responses for each example, relying solely on the internal knowledge of the model. As shown in Figure 7, even in non-RAG scenarios, ConsJudge maintains a high level of agreement with human judgments, consistent with its performance in RAG scenarios. The experimental results further demonstrate the generalization ability of ConsJudge.

#### A.5 Mitigating Position Biases in Evaluation

This section shows the robustness of ConsJudge to position bias. We randomly sample 200 queries from the RAG evaluation dataset, with ConsJudge implemented by Qwen2.5-14B.

For each query, we use Qwen2.5-14B to sample four candidate responses, among which we manually annotate the best response. Subsequently, we place the best response in the 1st, 2nd, 3rd, or 4th position among the four choices in the prompt, with the remaining three positions filled by other can-

Quantity	Hybrid Evaluation Aspects
Single	Hallucination, Completeness, Coherence, Semantic Consistency
Two	Hallucination + Completeness, Coherence + Semantic Consistency
Three	Hallucination + Completeness + Semantic Consistency
Four	Hallucination + Completeness + Coherence + Semantic Consistency

Table 8: Statistics of the Hybrid Evaluation Aspects.

didate responses. In this way, we construct four test scenarios in which the best response occupies a different position in each. By comparing the judgment accuracy of Vanilla LLM, SFT LLM, and ConsJudge across these positional scenarios, we evaluate their capability to mitigate position bias. Consistently higher judgment accuracy across all positional conditions indicates reduced vulnerability to position bias during evaluation.

As shown in Table 7, ConsJudge outperforms Vanilla LLM and SFT LLM in position accuracy, demonstrating its effectiveness in reducing position bias impacts.

#### A.6 More Experimental Details

In this section, we introduce more details of our experiments. We first show the details of the training ConsJudge. Then, we describe the details of applying ConsJudge to optimize the RAG model.

**ConsJudge Training.** To construct the ConsJudge Training dataset, as shown in Table 5, we collect multiple queries from these datasets and use four different LLMs, MiniCPM-2.4B (Hu et al., 2024), MiniCPM3-4B (Hu et al., 2024), Llama3-8B-Instruct (Touvron et al., 2023) and Qwen1.5-14B-Chat (Bai et al., 2023) to generate responses for each query. Specifically, each LLM generates three responses using three different temperatures, 0.5, 0.6, and 0.7, and we randomly sample one response from them, resulting in a total of four responses for each query. Furthermore, we combine the four different evaluation dimensions, resulting in eight hybrid evaluation aspects. As shown in Table 8, these include four individual evaluation dimensions, two combinations of two dimensions, one combination of three dimensions, and one that integrates all evaluation dimensions. For the hybrid evaluation aspects combining two evaluation dimensions, one integrates Coherence and Semantic Consistency, focusing on evaluating the logical coherence and fluency of the response, while another combines Hallucination and Completeness,

emphasizing whether the response is factually accurate and complete. For the hybrid aspects involving three dimensions, we exclude the Coherence, as it is less relevant in the RAG scenario than the other evaluation dimensions.

**RAG Training.** To construct the DPO training data to optimize the RAG models, we employ ConsJudge to select the best and worst responses from the sampling responses generated by RAG models. As shown in Table 5, we collect multiple queries from these datasets and use bge-large (Xiao et al., 2023) to retrieve top-5 relevant documents for each query. To enhance sampling diversity, RAG models generate responses under two different input conditions: the query alone (without RAG) and the query with the top-5 retrieved documents. RAG model samples two responses for each different input, yielding a total of four sampled responses. After that, we use the judgment model to select the best and worst responses from them to construct the DPO training dataset.

#### A.7 Prompt Templates Used in Experiment

In this section, we present the prompt templates used in our experiment.

First, we present the prompt designed for ConsJudge to evaluate the responses generated by RAG models, as shown in Table 9. Next, as illustrated in Figures 8 and 9, we introduce the prompts used for training and evaluating the RAG models. These prompt templates are based on RA-DIT (Lin et al., 2023) and RAG-DDR (Li et al., 2024b), specifically tailored to different LLMs and tasks to facilitate the generation of more effective responses. Additionally, the prompt designed for evaluating the performance of RAG models across the MARCO QA and WoW datasets using the GLM-4-plus model is displayed in Figure 10. Finally, Figure 11 presents the prompt used to instruct the GLM-4-plus to compare the judge quality between different judgment models.

#### A.8 More Case Studies

In this section, we present two more detailed cases for analysis.

First, as illustrated in Figure 12, in this case we show how ConsJudge selects more appropriate evaluation dimensions than baseline models, leading to more accurate judgments. The Vanilla model primarily relies on Completeness for evaluation, leading to an incorrect judgment. Meanwhile, the SFT model considers all dimensions to evaluate

the responses. In contrast, ConsJudge focuses on Hallucination, Coherence, and Semantic Consistency, and evaluates different responses using the most suitable dimension and provides more accurate judgments.

Then, to better understand the four defined evaluation dimensions, we present a specific case that uses four separate dimensions to evaluate. As shown in Figure 13, we can obtain diverse judgments by evaluating from different dimensions to enhance the judgment performance of the model.



<i>Dimensions Descriptions</i>
<p><b>Hallucination:</b> Hallucination refers to the presence of information in the option that contradicts ground truth, it is an incorrect answer to the question.</p> <p><b>Completeness:</b> Completeness refers to whether the choice contains as complete information as possible from the ground truth. it did not fully answer the question correctly.</p> <p><b>Coherence:</b> coherence refers to whether the choice is logically coherent and whether the language between each sentence is fluent.</p> <p><b>Semantic Consistency:</b> Semantic Consistency refers to whether the choice is semantically consistent with the ground truth, rather than just having lexical repetition.</p>
<i>Prompt</i>
<p>You are an excellent evaluation expert. Please select the best answer and the worst answer from four choices based on the ground truth and the query from the <i>{Name of the hybrid evaluation aspects.}</i> aspect.</p> <p><i>{Here is the descriptions of the hybrid evaluation aspects.}</i></p> <p>Note: your result format must strictly be "COT:{.there is your analysis}.</p> <p>Answer : Best answer:{ a choice must be one of[A,B,C,D]}.</p> <p>Worst answer :{ a choice must be one of[A, B, C, D]}".</p> <p>Output the content of COT first and then output the Answer.</p> <p>Here is the query:{query}, Here is the ground truth:{Ground Truth}</p> <p>Here is the A choice:{choiceA}, Here is the B choice:{choiceB},</p> <p>Here is the C choice:{choiceC}, Here is the D choice:{choiceD}.</p> <p>Result:</p>

Table 9: The Prompt Templates Used in ConsJudge.

<p><b>[TASK]</b> : Web Questions , Wiki QA , Yahoo! QA , MS MARCO QA ,</p> <p><b>[Instruction]</b> : Q : {Question}</p> <p>A:</p> <p><b>[TASK]</b> : Algebra QA with Rationales , Explanations for CommonsenseQA</p> <p><b>[Instruction]</b> : Please answer multiple choice question and choose the best answer choice first. Then give your explanation between [&lt;COT] and [COT&gt;].</p> <p>question : {Question}</p> <p>Answer :</p> <p><b>[TASK]</b> : Strategy QA</p> <p><b>[Instruction]</b> : Please answer the question. Then give your explanation between [&lt;COT] and [COT&gt;].</p> <p>question : {Question}</p> <p>Answer :</p>	<p><b>[TASK]</b> : Web Questions , Wiki QA , Yahoo! QA , MS MARCO QA ,</p> <p><b>[Instruction]</b> : Q : {Question}</p> <p>A:</p> <p><b>[TASK]</b> : Algebra QA with Rationales , Explanations for CommonsenseQA</p> <p><b>[Instruction]</b> : Please answer multiple choice questions below and output only the choice.</p> <p>{Question}</p> <p>Answer:</p> <p><b>[TASK]</b> : Strategy QA</p> <p><b>[Instruction]</b> : Please answer the question and only output the answer. Then give your explanation between [&lt;COT] and [COT&gt;].</p> <p>question : {Question}</p> <p>Answer :</p>
(a) MiniCPM-2.4B.	(b) Llama3-8B.

Figure 8: The Prompt Templates Used in Training Processes of RAG Models.

[TASK] : MARCO QA , NQ , TriviaQA , HotpotQA  
[Instruction] : Q : {Question}    A:

[TASK] : WoW  
[Instruction] : Q : {Question\_1}  
A : {Answer\_1}  
...  
Q : {Question\_n}  
A:

[TASK] : ASQA  
[Instruction] : Answer the following question. The question may be ambiguous and have multiple correct answers, and in that case, you have to provide a long-form answer including all correct answers.  
Question : {Question}  
Answer :

Figure 9: The Prompt Templates Used in Evaluation Processes of RAG Models.

You are a professional evaluator. Your task is to evaluate the quality of the given text based on query and the gold answer from **Accuracy, Informativeness, Fluency, Logicality and Harmlessness** dimensions.

**Accuracy:**1 star means Completely wrong,2 stars means Partially correct,3 stars means Completely correct;

**Informativeness:**1 star means Lack of necessary information or off-topic response,2 stars means Insufficient information to provide a complete answer,3 stars means Information is correct and sufficient;

**Fluency:**1 star means Many grammar errors,2 stars means Generally fluent but with some grammar errors,3 stars means Language is fluent and in accordance with human conventions;

**Logicality:** 1 star means Confused and full of logical flaws,2 stars means Some logical issues present,3 stars means Logically sound;

**Harmlessness:** 1 star means Violates ethics or is objectionable,2 stars means Mostly compliant but with a few flaws,3 stars means Fully compliant with public morality;

Please give the star for each dimension and the final average star, such as 'Accuracy: a number, Informativeness: a number, Fluency: a number, Logicality: a number, Harmlessness: a number. Average: a number'.

Here is the given {query}, Here is the given {gold}, Here is the given {text}

Figure 10: The Prompt Templates Used for GLM-4-plus to Evaluate the Performance of RAG Models on the MARCO QA and WoW Datasets.

You are a professional evaluator. Please evaluate the quality of the following two texts, based on the information in the question and ground truth, the stronger the logic of the text, the better the coherence, and the closer it is to the facts, the higher the quality of the text. Please only give your choice (A or B).

Question: {question}  
Ground Truth: {ground truth}  
Text A: {text 1}  
Text B: {text 2}

Figure 11: The Prompt Templates Used for GLM-4-plus to Evaluate the Judge Quality of the Different Judgment Models.

<p><b>Query:</b> " What is pseudo mean? "</p> <p><b>Ground truth:</b> " A combining form meaning false, pretended, unreal."</p> <p><b>Choice A:</b> " Pseudo means ‘fake’ or ‘imitation.’ In the context of statistics, a pseudo-mean is a value that is calculated using data from a sample rather than the actual population. It is used for estimation and approximation purposes, often in situations where it is not feasible to obtain data from the entire population or when an exact estimate is not necessary."</p> <p><b>Choice B:</b> " Pseudo means ‘fake’ or ‘imitation’ in Latin. It is not a commonly used term in the English language for any specific meaning. If you are referring to a specific context, please provide more information so I can give a more accurate answer."</p> <p><b>Choice C:</b> " A person who makes deceitful pretenses"</p> <p><b>Choice D:</b> " a person who makes deceitful pretenses"</p>	
<p><b>Raw Metric:</b> Choice C(0.154) &gt; Choice B(0.044) &gt; Choice A(0) = Choice D(0)</p> <p><b>Vanilla LLM:</b> "From the perspective of Hallucination and Completeness <b>Hallucination:</b> Choice A introduces the concept of a "pseudo-mean" in statistics, which is not mentioned in the ground truth and thus constitutes hallucination. <b>Completeness:</b> Choice B provides a definition but does not specify that "pseudo" means "false, pretended, unreal," making it incomplete. Choices C and D introduce unrelated concepts about people making deceitful pretenses, which do not align with the ground truth.[Answer]: Best answer: B. Worst answer: D."</p>	
<p><b>SFT LLM:</b> "From the perspective of <b>Hallucination Completeness, Coherence and similarity consistency</b>, Choice A introduces the concept of a statistical term unrelated to the ground truth, indicating significant hallucination. Choice B accurately describes a person who deceives, aligning with the ground truth. Choice C misrepresents pseudo as a term for "fake" or "imitation," which is not the primary meaning. Choice D avoiding hallucination. Choice D is the best; Choice A is the worst. Choice B closely matches the semantic content of the ground truth. Choice C partially aligns but introduces semantic discrepancies. Choice B is the best in semantic consistency; Choice A is the worst.[Answer]Best answer: D. Worst answer: B."</p>	
<p><b>ConsJudge:</b> "From the perspective of <b>Hallucination</b>, the option A has a slight hallucination as it defines pseudo as ‘fake’ or ‘imitation’, which is partially correct but also includes additional information about its use in statistics, whereas the ground truth only mentions the combining form meaning ‘false, pretended, unreal’. Option B lacks <b>coherence</b> as it does not provide a clear definition of pseudo. Options C and D are coherent but lack <b>semantic consistency</b> with the ground truth as they define a pseudo as a person, whereas the ground truth defines it as a combining form.[Answer]: Best answer: A. Worst answer: B."</p>	

Figure 12: An Example of Evaluation Results of ConsJudge Compared to Other Baseline Models.

<p><b>Query:</b> " how big is auburndale florida?"</p> <p><b>Ground truth:</b> "According to the U.S Census estimates of 2005, the city had a population of 12,381.."</p> <p><b>Choice A:</b> " The U.S. Census estimates for Auburndale, Florida, show that the city had a population of 13,381 in 2005, indicating the size of the city."</p> <p><b>Choice B:</b> " According to the 2005 U.S. Census data, the population of Auburndale, Florida, was recorded as 12,381."</p> <p><b>Choice C:</b> "Auburndale’s population is estimated to be 1,238,100 according to the U.S. Census of 2005."</p> <p><b>Choice D:</b> "U.S. has a land area of 12,381 square miles according to the U.S. Census estimates"</p>
<p><b>Hallucination:</b> From the perspective of the <b>Hallucination</b> dimension, I analyze the options provided. Option A has a hallucination because it states the population of Auburndale, Florida is 13,381, which contradicts the ground truth of 12,381. Option C has a hallucination because it gives an extremely large population of 1,238,100, which is not supported by the ground truth. Option D has a hallucination because it mentions the land area of the U.S., not Auburndale, Florida..[A]:<b>Best answer:B</b>, which accurately reports the population of Auburndale, Florida, as per the 2005 U.S. Census data.[<b>Worst answer:A,C,D</b>].</p>
<p><b>Completeness:</b> "From the perspective of <b>completeness</b> dimension, I analyzed the options provided and compared them to the ground truth. Option B provides the most accurate information about the population of Auburndale, Florida in 2005, which matches the ground truth exactly. Options A and B provide similar information, but option A includes an error in the population figure. Option C is significantly off from the ground truth, providing an incorrect population figure that is orders of magnitude larger. Option D is also incorrect, as it provides information about the land area of the United States, not Auburndale, Florida.[Answer]:<b>Best answer:B.Worst answer:C</b>."</p>
<p><b>Coherence:</b> "From the perspective of <b>coherence</b>, I analyze the options based on logical consistency and fluency of language. The options provide information about the population of Auburndale, Florida, with some accuracy issues. * Option A has an error in the population figure, which is not consistent with the ground truth. * Option B accurately reports the population figure of 12,381, which matches the ground truth. * Option C provides an incorrect population figure, which is significantly higher than the actual value. * Option D is unrelated to the population of Auburndale, Florida, and instead provides information about the land area of the United States.[Answer]:<b>Best answer:B.Worst answer:A</b>."</p>
<p><b>Semantic Consistency:</b> "From the perspective of <b>Semantic Consistency</b> dimension, I analyzed the options based on the provided ground truth that the population of Auburndale, Florida in 2005 was 12,381. * Option A has a slight deviation from the ground truth, as it states the population as 13,381 instead of 12,381. * Option B accurately reflects the ground truth by stating the exact population figure. * Option C is semantically inconsistent as it provides an incorrect population figure (1,238,100) that is not even close to the actual value. * Option D is semantically inconsistent as it mentions the land area of the U.S., which is not relevant to the question about the size of Auburndale, Florida.[Answer]:<b>Best answer:B.Worst answer:C</b>."</p>

Figure 13: An Example of Evaluation Results of ConsJudge Evaluates Separately Using Four Dimensions.