

# MARCH: Multi-Agent Reinforced Check for Hallucination

Zhuo Li<sup>1,2\*</sup>, Yupeng Zhang<sup>1\*</sup>, Pengyu Cheng<sup>1†</sup>, Jiajun Song<sup>1</sup>, Mengyu Zhou<sup>1</sup>, Hao Li<sup>1</sup>, Shujie Hu<sup>1</sup>, Yu Qin<sup>1</sup>, Erchao Zhao<sup>1</sup>, Xiaoxi Jiang<sup>1</sup>, Guanjun Jiang<sup>1</sup>

<sup>1</sup>Qwen Large Model Application Team, Alibaba  
<sup>2</sup>The Chinese University of Hong Kong, Shenzhen

## Abstract

Hallucination remains a critical bottleneck for large language models (LLMs), undermining their reliability in real-world applications, especially in Retrieval-Augmented Generation (RAG) systems. While existing hallucination detection methods employ LLM-as-a-judge to verify LLM outputs against retrieved evidence, they suffer from inherent *confirmation bias*, where the verifier inadvertently reproduces the errors of the original generation. To address this, we introduce **Multi-Agent Reinforced self-Check for Hallucination (MARCH)**, a framework that enforces rigorous factual alignment by leveraging deliberate *information asymmetry*. MARCH orchestrates a collaborative pipeline of three specialized agents: a Solver, a Proposer, and a Checker. The Solver generates an initial RAG response, which the Proposer decomposes into claim-level verifiable atomic propositions. Crucially, the Checker validates these propositions against retrieved evidence in isolation, deprived of the Solver’s original output. This well-crafted information asymmetry scheme breaks the cycle of self-confirmation bias. By training this pipeline with multi-agent reinforcement learning (MARL), we enable the agents to co-evolve and optimize factual adherence. Extensive experiments across hallucination benchmarks demonstrate that MARCH substantially reduces hallucination rates. Notably, an 8B-parameter LLM equipped with MARCH achieves performance competitive with powerful closed-source models. MARCH paves a scalable path for factual self-improvement of LLMs through co-evolution. The code is at <https://github.com/Qwen-Applications/MARCH>.

## 1 Introduction

Despite the remarkable success achieved by Large Language Models (LLMs) (OpenAI, 2024; Touvron et al., 2023; Yang et al., 2024; DeepSeek-AI,

2025; Gemini, 2025), the persistence of *hallucinations* remains a fundamental barrier to LLMs’ trustworthy usage (Kalai et al., 2025; Huang et al., 2025; Ji et al., 2023; Maynez et al., 2020). Even enhanced within Retrieval-Augmented Generation (RAG) (Lewis et al., 2020), LLMs frequently generate responses that, while linguistically fluent, contradict the provided evidentiary context (Xu et al., 2025c). Such context-conflicting hallucinations (Liu et al., 2024b) are particularly prevalent in data-intensive tasks, manifesting as inaccurate numbers, distorted temporal sequences, or erroneous arithmetic aggregations, even when the source text explicitly contains the requisite information. In high-stakes domains, such as finance (Kang and Liu, 2023), law (Dahl et al., 2024; Du et al., 2025), and healthcare (Kim et al., 2025; Hu et al., 2022, 2026; Dai et al., 2025), hallucinations are more prohibitive, as users require LLMs to function as meticulous evidence readers rather than creative generators of plausible but unverified narratives.

Among approaches in mitigating hallucinations, Supervised Fine-Tune (SFT) has been primarily considered as effective (Elaraby et al., 2023; Lin et al., 2022; Zhou et al., 2023; Du et al., 2026). However, several practices emphasize that SFT could also amplify hallucinations by fundamentally operating on behavior cloning (Torabi et al., 2018), which induces models to prioritize authoritative stylistic patterns over factual accuracy (Casper et al., 2023; Zhang et al., 2025; Gekhman et al., 2024). Consequently, Reinforcement Learning (RL) has become a necessity for optimizing LLMs towards more truthfulness (Yang et al., 2025; Gemma, 2025). However, the standard outcome-level RL rewards lack the granularity required to supervise fine-grained factual consistency, leading to a fundamental misalignment between the scalar reward and the nuanced evidentiary requirements of RAG.

Specifically, Reinforcement Learning from Hu-

\* These authors contributed equally.

† Corresponding author.

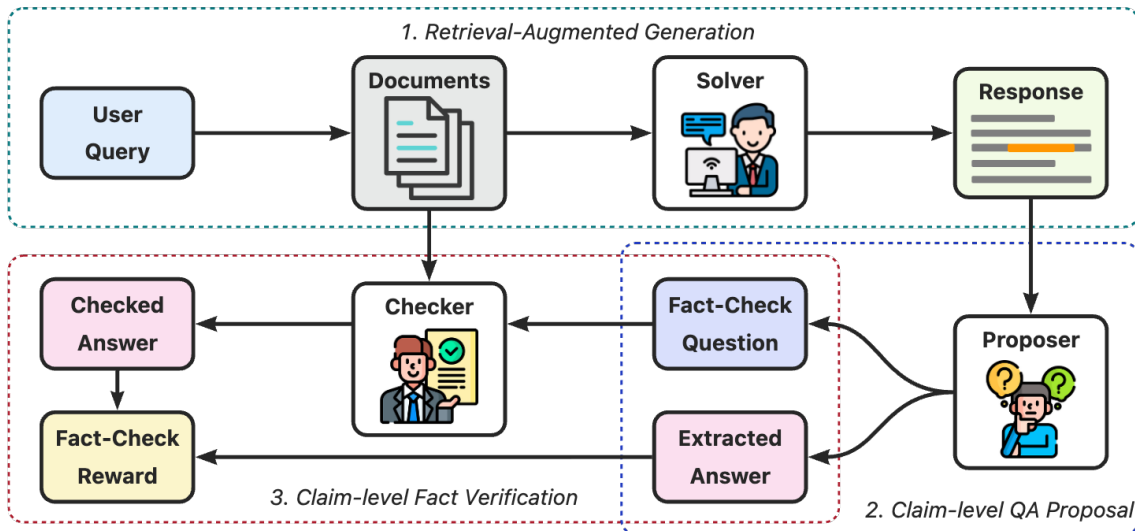


Figure 1: Overview of the MARCH framework. The *Solver* conducts retrieval-augmented generation based on the input query and related documents. The *Proposer* then decomposes this response into atomic propositions (the orange line in the response block) and formulates verifiable question-answer pairs. The *Checker* performs isolated verification by re-answering questions solely based on the retrieved evidence, without access to the *Solver*’s original output. This information-asymmetric pipeline, where the policy model plays all three roles, is optimized via Multi-Agent Reinforcement Learning (MARL) to achieve robust factual alignment.

man Feedback (RLHF) (Ouyang et al., 2022) typically relies on reward models (RMs) (Bradley and Terry, 1952; Li et al., 2025a,d; Liu et al., 2024a; Duan et al., 2026) that compress multifaceted evaluations into coarse-grained scalar scores. This loss of granularity limits the ability to verify individual claims, providing insufficient signals to reinforce the meticulous evidentiary grounding required for data-intensive tasks. Alternatively, while Reinforcement Learning with Verifiable Reward (RLVR) (DeepSeek-AI, 2025) attempts to provide objective feedback via deterministic rules (Shao et al., 2024), rubrics (Gunjal et al., 2025), or external generalist reward model (Liu et al., 2025b), it remains fundamentally constrained by a dual-dependency bottleneck: the scarcity of expert-annotated groundtruth in specialized domains, and the reasoning ceilings of external verifiers. Consequently, unsupported intermediate claims can persist, as models are rewarded for correct final outcomes despite potentially hallucinated reasoning trajectories (Li and Ng, 2025).

While post-hoc metrics and fact-checkers (Tang et al., 2024; Rashkin et al., 2023; Gao et al., 2023; Baker et al., 2025; Guan et al., 2025) provide utility for evaluation, such static tools are seldom integrated into the internal optimization loop to provide actionable, claim-oriented guidance. More critically, the efficacy of such verifiers is fundamentally undermined by the *information leakage*

inherent in current evaluation schemes (Seo et al., 2025). Specifically, concurrent exposure to the query, documents, and generated response induces a confirmation bias (Sharma et al., 2024). In this state, verifiers tend to prioritize validating internal coherence over objective grounding against source evidence, leading to the unwarranted endorsement of erroneous claims. This compromised “blind scrutiny” prevents the model from internalizing the rigorous, document-anchored reasoning essential for high-fidelity generation.

To address the above factuality challenges, we introduce **Multi-Agent Reinforced self-Check for Hallucination (MARCH)**, an RL-based framework that enforces factual alignment through deliberate information asymmetry. MARCH orchestrates a collaborative pipeline of three specialized agents: a *Solver*, a *Proposer*, and a *Checker*, to prevent model-based verifiers from taking cognitive shortcuts through confirmation bias.

As illustrated in Figure 1, the *Solver* first generates a response based on the query and retrieved documents. Subsequently, the *Proposer* acts as a ‘Response Atomizer,’ deconstructing the narrative into discrete, verifiable claims. Finally, the *Checker* validates these claims by re-answering them exclusively based on the source documents, strictly blinded to the *Solver*’s original output to avoid cognitive contamination. Through this asymmetric decoupling, MARCH transforms the verification

process from a circular confirmation of plausible text into a rigorous cross-examination between generated claims and factual evidence.

During training, MARCH promotes agentic self-evolution via a zero-tolerance reward mechanism: any discrepancy between the Proposer’s extracted claims and the Checker’s document-based validation triggers a penalty on the entire response trajectory. This encourages the agents to align their outputs strictly with retrieved evidence, effectively internalizing the capability for rigorous, document-anchored statements. Experimentally, MARCH achieves a substantial reduction in hallucinations compared to its base model, without requiring additional human annotations or external fact-checking tools, thereby demonstrating the efficacy of our asymmetric collaborative reinforcement paradigm. Notably, an 8B-parameter LLM equipped with MARCH achieves performance parity with leading closed-source models on multiple hallucination benchmarks, while simultaneously delivering significant gains in general RAG-QA tasks. These results establish a scalable and verifiable pathway toward trustworthy, agentic self-improvement for large language models.

## 2 Methodology

### 2.1 Problem Formulation and Overview

Given an input query  $x \in \mathcal{X}$  and a set of retrieved documents  $D = [d_1, d_2, \dots, d_l]$ , we aim to optimize a policy model  $\pi_\theta(\mathbf{y}|x, D)$  that generates a response  $\mathbf{y}$  with maximum factual grounding aligned with the provided evidences in  $D$ . Formally, suppose we can extract all the statements of facts  $\{\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_n\}$  from the response  $\mathbf{y}$  using a function  $\mathcal{A}(\cdot)$ , as  $\mathcal{A}(\mathbf{y}) = \{\mathbf{a}_i\}_{i=1}^n$ . Then, from the perspective of reinforcement learning with verifiable reward (RLVR), if an oracle set of ground-truth annotations  $\mathcal{A}^* = \{\mathbf{a}_1^*, \dots, \mathbf{a}_n^*\}$  is ideally available for every atomic statement  $\mathbf{a}_i$  in  $\mathbf{y}$ , the optimization objective is to maximize the expected evidentiary consistency:

$$\max_{\theta} \mathbb{E}_{x \sim \mathcal{X}, \mathbf{y} \sim \pi_\theta(\cdot|x, D)} [R(\{\mathbf{a}_i, \mathbf{a}_i^*\}_{i=1}^n)], \quad (1)$$

where  $R(\cdot)$  is a fact-check function that identifies the exact match between  $\mathcal{A}(\mathbf{y}) = \{\mathbf{a}_i\}_{i=1}^n$  and  $\{\mathbf{a}_i^*\}_{i=1}^n$ . However,  $\mathcal{A}^*$  is practically unavailable in most application domains and requires enormous human annotation.

Therefore, we reformulate equation 1 into a self-contained, verifiable objective by decomposing the

grounding process into three separate agent tasks, as illustrated in Figure 1:

1. **Response Generation:** a Solver agent ordinarily responses  $\mathbf{y}$  based on the input query  $x$  and retrieved context  $D$ .
2. **Claim-level QA Proposal:** a Proposer agent decomposes the response  $\mathbf{y}$  into  $n$  discrete factual claims in a question-answering (QA) format, as  $\mathcal{Q}(\mathbf{y}) = \{(\mathbf{q}_i, \mathbf{a}_i)\}_{i=1}^n$ .
3. **Factual Verification:** a Checker agent answers all the atomic questions  $\{\mathbf{q}_i\}_{i=1}^n$  as  $\{\hat{\mathbf{a}}_i\}_{i=1}^n$  with only the retrieval documents  $D$  provided, without any access to the original response  $\mathbf{y}$  and extracted answers  $\{\mathbf{a}_i\}_{i=1}^n$ .

By substituting the unreachable oracle  $\mathbf{a}_i^*$  with the grounding result  $\hat{\mathbf{a}}_i$  from the multi-agent pipeline, we approximate the factuality optimization as:

$$\max_{\theta} \mathbb{E}_{x \sim \mathcal{X}, \mathbf{y} \sim \pi_\theta(\cdot|x, D)} [R(\{(\mathbf{a}_i, \hat{\mathbf{a}}_i)\}_{i=1}^n)]. \quad (2)$$

By optimizing equation 2, we force  $\pi_\theta$  to internalize a rigorous verification logic, ensuring that the generated content is strictly anchored in the source evidence to enhance overall factual accuracy. We call the proposed learning scheme **Multi-Agent Reinforced self-Check for Hallucination (MARCH)**.

### 2.2 MARCH Modeling

As shown in Figure 1, we instantiate the functional primitives identified above into three specialized agents derived from the same base policy  $\pi_\theta$ : a *Solver*, a *Proposer*, and a *Checker*, which are conditioned to exhibit distinct epistemic behaviors, effectively transforming the RAG task into a structured, asymmetric cooperative game. MARCH begins with the Solver’s policy:

$$\nu_{\text{solve}}(\cdot|x, D) = \pi_\theta(\cdot|x, D, \mathbf{s}_{\text{solve}}), \quad (3)$$

where  $\mathbf{s}_{\text{solve}}$  is the Solver’s system prompt (as shown in Appendix A.1.1) to induce the policy model  $\pi_\theta$  to generate a comprehensive response  $\mathbf{y}$  given the retrieval documents  $D$  and the input  $x$ .

To decouple the continuous response  $\mathbf{y}$  into verifiable components, the Proposer agent functions as a *Response Atomizer* that bridges narrative synthesis and discrete verification:

$$\nu_{\text{propose}}(\cdot|\mathbf{y}) = \pi_\theta(\cdot|\mathbf{y}, \mathbf{s}_{\text{propose}}), \quad (4)$$

where the Proposer’s system prompt  $s_{\text{propose}}$  is in Appendix A.1.2. Specifically,  $\nu_{\text{propose}}$  decomposes the response  $\mathbf{y}$  into a set of  $n$  self-contained factual question-answer (QA) pairs  $\mathcal{Q}(\boldsymbol{\tau}) = \{(\mathbf{q}_i, \mathbf{a}_i)\}_{i=1}^n$ , where each  $\mathbf{q}_i$  represents a targeted inquiry, and each  $\mathbf{a}_i$  denotes the corresponding claim asserted in  $\mathbf{y}$ . The function  $\mathcal{Q}(\cdot)$  extracts the QA pairs from the Proposer’s generation  $\boldsymbol{\tau} \sim \nu_{\text{propose}}(\cdot|\mathbf{y})$ . the Checker operates as the *Blinded Auditor* and the linchpin of the asymmetric cooperation:

$$\nu_{\text{check}}(\cdot|\{\mathbf{q}_i\}_{i=1}^n, \mathbf{D}) = \pi_{\theta}(\cdot|\{\mathbf{q}_i\}, \mathbf{D}, s_{\text{check}}) \quad (5)$$

with the Checker’s prompt  $s_{\text{check}}$  in Appendix A.1.3 To eliminate confirmation bias and break information symmetry, when tasked with answering  $\mathbf{q}_i$ ,  $\nu_{\text{check}}$  is strictly blinded to the Solver’s response  $\mathbf{y}$  and the corresponding claims  $\mathbf{a}_i$ . the *Checker* replies all the queries  $\{\mathbf{q}_i\}_{i=1}^n$  in a response  $\boldsymbol{\lambda} \sim \nu_{\text{check}}(\cdot|\{\mathbf{q}_i\}_{i=1}^n, \mathbf{D})$ , only based on the RAG documents  $\mathbf{D}$ . The corresponding answers  $\mathcal{A}(\boldsymbol{\lambda}) = \{\hat{\mathbf{a}}_i\}_{i=1}^n$  are parsed by the function  $\mathcal{A}(\cdot)$  to serve as the re-checked answers to the QA pairs  $\{(\mathbf{q}_i, \mathbf{a}_i)\}_{i=1}^n$ .

**Overall Objective.** By integrating the conditional dependencies of the *Solver*, the *Proposer*, and the *Checker*, MARCH defines a unified asymmetric cooperative and play objective by focusing on the evidentiary identity between *Solver*’s response  $\mathbf{y}$  and *Checker*’s independent grounding trajectory  $\boldsymbol{\lambda}$ , mediated by the atomic claim decomposition of the *Proposer*. The optimization goal is to maximize the expected reward based on the consistency of claims, marginalized over the discrete atomization performed by  $\nu_{\text{propose}}$ :

$$\begin{aligned} \max_{\pi_{\theta}} \quad & \mathbb{E}_{\mathbf{x} \sim \mathcal{X}, \mathbf{y} \sim \nu_{\text{solve}}(\cdot|\mathbf{x}, \mathbf{D})} [R(\{(\mathbf{a}_i, \hat{\mathbf{a}}_i)\}_{i=1}^n)] \\ \text{s.t.} \quad & \{(\mathbf{q}_i, \mathbf{a}_i)\}_{i=1}^n = \mathcal{Q}(\boldsymbol{\tau}), \boldsymbol{\tau} \sim \nu_{\text{propose}}(\cdot|\mathbf{y}) \\ & \{\hat{\mathbf{a}}_i\}_{i=1}^n = \mathcal{A}(\boldsymbol{\lambda}), \boldsymbol{\lambda} \sim \nu_{\text{check}}(\cdot|\{\mathbf{q}_i\}_{i=1}^n, \mathbf{D}) \end{aligned} \quad (6)$$

where  $R(\cdot)$  measures the evidentiary identity between each asserted claim  $\mathbf{a}_i$  and the independent grounding  $\hat{\mathbf{a}}_i$ . By optimizing equation 6, the policy  $\pi_{\theta}$  is compelled to internalize a rigorous verification logic, ensuring that its narrative synthesis is strictly anchored in verifiable evidence to enhance overall factual accuracy, which also develops a self-consistent grounding capability that eliminates reliance on external supervision.

### 2.3 MARCH Implementation

In our implementation, we specifically prioritize numerical and quantitative verification, driven by the observation that in data-intensive RAG scenarios, LLMs frequently suffer from hallucinations such as misquoting figures or altering numerical values during context integration. By anchoring the verification loop on these high-fidelity tokens, MARCH provides a more rigorous evidentiary baseline for grounding long-form reasoning.

**Zero-Tolerance Reward (ZTR).** To enforce absolute factual grounding, we define a binary Zero-Tolerance Reward function  $R(\{(\mathbf{a}_i, \hat{\mathbf{a}}_i)\}_{i=1}^n)$  in equation 6. In the RAG context, we posit that any single hallucinated claim  $\mathbf{a}_i$  within the response  $\mathbf{y}$  invalidates the entire reasoning chain. Therefore, a trajectory is assigned a successful score only if every claim  $\mathbf{a}_i$  asserted by the *Solver* achieves an identity match with the consensus answer  $\hat{\mathbf{a}}_i$  extracted by the *Checker*:

$$R(\{(\mathbf{a}_i, \hat{\mathbf{a}}_i)\}_{i=1}^n) = \begin{cases} 0, & \text{if every } \mathbf{a}_i = \hat{\mathbf{a}}_i \\ -1, & \text{otherwise} \end{cases} \quad (7)$$

This ‘‘all-or-nothing’’ reward structure prevents the policy  $\pi_{\theta}$  from optimizing for partial correctness or stylistic plausibility, forcing the model to prioritize strict evidentiary identity. Besides, to ensure the *Checker* provides a reliable baseline for the audit and to reduce the variance of the reward signal, we employ a multi-sampling strategy during the cross-validation phase. For each question  $\mathbf{q}_i$  generated by the *Proposer*,  $\nu_{\text{check}}$  generates multiple independent audit trajectories. We apply a majority voting to determine the final consensus answer  $\hat{\mathbf{a}}_i$ , which effectively mitigates stochastic errors in *Checker*’s generation in one independent audit.

**Policy Optimization.** We optimize the shared policy  $\pi_{\theta}$  of the three agents using Proximal Policy Optimization (PPO) (Schulman et al., 2017). Critically, we incorporate both the reasoning path  $\mathbf{y}$  and the audit trajectory  $\boldsymbol{\lambda} \sim \nu_{\text{check}}(\cdot|\{\mathbf{q}_i\}, \mathbf{D})$  into the optimization loop, ensuring that  $\pi_{\theta}$  concurrently evolves as both a reliable generator and a rigorous auditor. Consequently, each original query of a training batch  $\mathbf{x}_i \in \mathcal{B}$  contributes two distinct reward signals to the policy, allowing  $\pi_{\theta}$  to learn the fine-grained correspondence between narrative synthesis and independent document-grounding within a unified training batch. While both trajectories

Table 1: Statistics of the training dataset, where we report the number of training samples, average query length, the number of retrieved documents per query, the average document length, and the proportion of documents that are irrelevant to the query. The large variance in document counts and lengths, together with the high ratio of irrelevant documents, highlights the realism and difficulty of our training setup.

Dataset	Original Dataset	#Samples	Avg. Query Len	#Docs per Query			Total Doc Len (words)			Irrelevant Doc Ratio
				Avg	Max	Min	Avg	Max	Min	
STEM	BioASQ	4721	8.62	9.03	157	1	1314.86	22135	1	30.14%
General	2WikiMultiHopQA	4500	12.87	10	10	10	575.65	2099	95	24.4%
	MuSiQue	4500	15.85	19.99	20	16	11221.72	20741	4112	88.35%

derive their terminal reward from the same evidentiary identity match (equation 7), corresponding advantages  $\hat{A}$  and KL-divergence (Csiszár, 1975) penalties against the reference model  $\pi_{\text{ref}}$  are estimated independently. The joint gradient is computed as the sum of gradients across all individual trajectories  $\nabla_{\theta} \mathcal{L}_{\text{MARCH}}(\theta) \approx$

$$\frac{1}{2|\mathcal{B}|} \sum_{\mathbf{x} \in \mathcal{B}} \left[ \sum_{t=1}^{|\mathbf{y}|} \hat{A}_y^t \nabla_{\theta} \log \pi_{\theta}(y^t | \cdot) + \sum_{t=1}^{|\lambda|} \hat{A}_{\lambda}^t \nabla_{\theta} \log \pi_{\theta}(\lambda^t | *) - \beta \nabla_{\theta} \text{KL}[\pi_{\theta} \| \pi_{\text{ref}}] \right]$$

where  $\hat{A}_{\lambda}^t$  and  $\hat{A}_y^t$  are the  $t$ -th token’s corresponding advantages of the Checker’s output  $\lambda$  and the Solver’s output  $\mathbf{y}$ , estimated via Generalized Advantage Estimation (Schulman et al., 2015). The notation  $(\cdot)$  and  $(*)$  in equation 8 encapsulates the respective contexts for brevity:  $(\mathbf{x}, \mathbf{D}, \mathbf{y}^{<t})$  and  $(\{\mathbf{q}_i\}, \mathbf{D}, \lambda^{<t})$  for the Solver’s and the Checker’s generation prefixes. equation 8 ensures that  $\pi_{\theta}$  concurrently refines its narrative generation and evidentiary precision, internalizing a self-contained verification loop where generative assertions are inherently tethered to the policy’s own auditing capacity. The training process is in Algorithm 1.

### 3 Experiment

We first evaluate the effectiveness of MARCH in mitigating hallucinations. Then, we show that the improved factual consistency ability helps downstream reasoning tasks like Question-Answering.

#### 3.1 Training Setup

**Training Datasets and Base Model.** We train our framework covering: a STEM (Science-Technology-Engineering-Math) setting based on BioASQ (Krithara et al., 2023), and a General setting based on 2WikiMultiHopQA (Ho et al., 2020) and MuSiQue (Trivedi et al., 2022). Crucially, training is performed using only queries and retrieved documents, without access to ground-truth

answers or annotations. Table 1 shows the number of retrieved documents per query and the document lengths vary widely across datasets, and a substantial fraction of the documents are irrelevant to the query (e.g., 30.1% for STEM and up to 88.4% for MuSiQue). Such a high level of noise and heterogeneity in the retrieval docs makes the training setting highly realistic but also extremely challenging, as the agents must learn to perform document-grounded fact checking in the presence of many noisy documents. We consider Meta-Llama3.1-8B-Instruct (Touvron et al., 2023) as the training start.

**Training Details.** The training framework is built upon the VerL (Sheng et al., 2025), and we employ Fully Sharded Data Parallel (FSDP) (Zhao et al., 2023) for efficient multi-node multi-GPU training. During the training rollout, we utilize the vLLM (Kwon et al., 2023) for efficient sample generation with the temperature of 0.6. We give the detailed PPO training recipes in Appendix B.1.

#### 3.2 Experiments on Hallucination

**Evaluation Benchmark.** We evaluate MARCH along the aspect of *hallucination and factual grounding* on the RAGTruth (Niu et al., 2024), FaithBench (Tamber et al., 2025), Contextual-JudgeBench (Xu et al., 2025a) and Facts Grounding benchmark (Jacovi et al., 2025), which cover QA, data-to-text writing, summarization, and diverse domains such as medical, financial, and legal. More details are provided in Appendix B.2.1. The best performance is marked with **Red** and the second-best with **Blue**.

**Evaluation Protocol and Metrics.** Following Niu et al. (2024); Tamber et al. (2025); Jacovi et al. (2025); Xu et al. (2025a), we focus on *response-level* factuality eval. Given generation, retrieved documents, and ground-truth reference, we deploy a Qwen3-235B-A22B (Yang et al., 2025) as a judge model to decide whether the response is factually consistent/accurate with

Table 2: Consistency Rate (%) on RAGTruth and FaithBench benchmark. Baseline results are cited from [https://github.com/vectara/FaithJudge/tree/main/eval\\_results](https://github.com/vectara/FaithJudge/tree/main/eval_results).

Model	RAGTruth			FaithBench	Average
	Summary	Data2Txt	QA		
o4-mini-high-2025-04-16	76.67	64.00	82.01	45.83	67.13
o4-mini-medium-2025-04-16	76.67	67.33	88.49	55.56	71.01
o4-mini-low-2025-04-16	70.00	71.33	83.45	55.56	70.09
GLM-4-9B-Chat	87.77	52.67	90.67	61.11	73.06
Phi-4-mini-Instruct	78.67	24.67	94.96	38.89	59.30
Qwen2.5-14B-Instruct	82.00	51.33	93.53	45.83	68.17
Ministral-8B-Instruct	83.33	49.33	88.49	43.06	66.05
Meta-Llama3.1-8B-Instruct	71.33	48.67	63.31	37.50	55.20
<b>MARCH-STEM</b>	86.67	72.67	83.45	56.94	74.93
Δ	↑15.34	↑24.00	↑20.14	↑19.44	↑19.73
<b>MARCH-General</b>	92.67	72.00	83.45	52.78	75.23
Δ	↑21.34	↑23.33	↑20.14	↑15.28	↑20.03

the evidence/ground-truth factual labels. We mitigate the randomness during evaluation by drawing eight independent generations for each test query; the judge model then evaluates each sampling individually, and the final verdict is determined by a majority vote across these eight trajectories.

**Results on RAGTruth and FaithBench.** Table 2 shows that our method brings substantial and consistent improvements over the base Llama3.1-8B-Instruct across all settings. With STEM data (*MARCH-STEM*), the learned policy improves the average accuracy from 55.20 to 74.93 (+19.73) and with General QA data (*MARCH-General*), the average accuracy further rises to 75.23 (+20.03), which suggest that MARCH reliably strengthens the base model’s ability to judge and mitigate hallucinations, rather than overfitting to a particular domain. When compared to other larger or proprietary LLMs, our policies also achieve an advanced average performance. Together, these results demonstrate that our framework can turn a small-sized base LLM into a highly competitive factuality judge simply by training on noisy, document-grounded RAG data.

**Results on Facts Grounding.** FACTS Grounding evaluates whether model responses are fully supported by a single source document. Figure 2 shows that our policies achieve large gains over the base model, reaching a factuality score of 85.23% (*MARCH-STEM*) and 80.12% (*MARCH-General*), compared with 57.09% for the base Llama3.1-8B-Instruct. Our method successfully lifts an 8B backbone into the comparable performance band as leading proprietary models, placing *MARCH-STEM* just behind the top Gemini 2.5/1.5 Flash models on the public leaderboard while surpassing popular models such as GPT-4o and several larger open models. The strong results indicate that MARCH

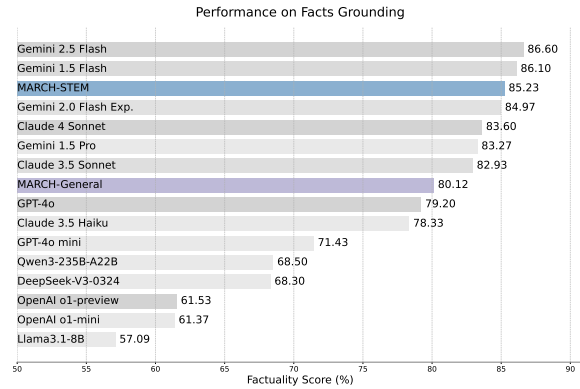


Figure 2: Comparison of performance on the Facts Grounding. MARCH (highlighted in blue and purple) demonstrates competitive performance against a range of leading open and proprietary models.

substantially improves long-document grounding robustness across heterogeneous domains.

**Results on ContextualJudgeBench.** Table 3 reports results on ContextualJudgeBench, which evaluates multiple dimensions of judge behavior, including refusal, faithfulness, completeness, and conciseness. Starting from Llama3.1-8B, our framework again yields large gains across all eight dimensions under both training regimes. Specifically, the policy trained on General QA data (*MARCH-General*) improves the average accuracy from 29.7 to 51.6 (+21.9), while the policy trained on STEM data (*MARCH-STEM*) reaches 52.3 (+22.6). The improvements are particularly pronounced on faithfulness and completeness (both QA and summarization), where accurate document-grounded judgment is most critical. Based on an 8B LLM, our policies surprisingly reach or surpass the performance of many much larger “Instruct + Reasoning” judge models, and approach the best specialized judge models on several dimensions. Table 3 indicates that MARCH not only improves binary hallucination detection but also equips LLM with a more nuanced ability to assess answer quality along multiple axes, such as whether key information is covered and whether responses remain concise and well-calibrated in context.

### 3.3 Experiments on Multi-Hop QA

Intuitively, the capacity for precise factual extraction and verification reasonably serves as the fundamental bedrock for complex reasoning, as Question Answering tasks usually require a model to sequentially bridge information across disparate documents. In this section, we empirically find

Table 3: Accuracy (%) on ContextualJudgeBench. Baseline results are cited from Xu et al. (2025a). More comprehensive results are in Appendix B.2.2.

Model	Refusal (Ans.)	Refusal (Unans.)	Faithfulness (QA)	Faithfulness (Summ.)	Completeness (QA)	Completeness (Summ.)	Conciseness (QA)	Conciseness (Summ.)	Average
<i>Large Judge Model</i>									
Prometheus-2-8x7b (Kim et al., 2024)	22.0	29.6	22.4	29.6	20.4	39.8	10.2	18.4	24.1
Skywork-70B (Liu et al., 2024a)	82.4	11.2	48.0	47.6	36.8	41.4	21.6	27.9	39.6
ST-Eval-70B (Wang et al., 2024a)	50.0	42.0	51.2	45.6	40.8	39.4	36.1	29.9	41.9
SFRJudge-70B (Wang et al., 2025)	87.6	32.4	60.8	<b>54.8</b>	40.8	53.4	<b>44.7</b>	36.1	51.4
<i>Instruct + Reasoning Model</i>									
Meta-Llama-3.1-70B (Touvron et al., 2023)	59.6	48.0	58.0	48.4	38.0	51.8	15.7	27.5	43.4
Meta-Llama-3.3-70B (Touvron et al., 2023)	71.6	42.4	68.0	48.4	42.0	51.8	20.8	30.7	47.0
R1-Distill-Llama-3.3-70B (DeepSeek-AI, 2025)	<b>89.6</b>	50.4	<b>74.0</b>	48.4	42.4	57.4	19.2	29.5	51.4
GPT-4o-mini	71.2	22.8	45.6	42.4	33.2	54.2	11.8	29.5	38.8
GPT-4o	64.0	52.0	68.0	50.8	39.6	56.2	12.9	22.5	45.8
DeepSeek-R1 (DeepSeek-AI, 2025)	<b>92.0</b>	52.0	<b>72.0</b>	50.4	41.2	60.6	20.4	26.2	<b>51.9</b>
<i>OURS</i>									
Meta-Llama-3.1-8B (Touvron et al., 2023)	28.0	43.2	34.8	34.8	23.2	41.0	11.4	21.3	29.7
<b>MARCH-General</b>	37.2	57.8	52.6	<b>54.8</b>	<b>48.0</b>	<b>63.9</b>	<b>48.1</b>	<b>50.1</b>	51.6
Δ	↑ 9.2	↑ 14.6	↑ 17.8	↑ 20.0	↑ 24.8	↑ 22.9	↑ 36.7	↑ 28.8	↑ 21.9
<b>MARCH-STEM</b>	38.1	<b>64.0</b>	53.6	<b>58.0</b>	44.9	<b>65.5</b>	<b>44.7</b>	<b>49.5</b>	<b>52.3</b>
Δ	↑ 10.1	↑ 20.8	↑ 18.8	↑ 23.2	↑ 21.7	↑ 24.5	↑ 33.3	↑ 28.2	↑ 22.6

that the improved factual grounding brought by MARCH can directly translate into better performance on reasoning-intensive downstream tasks.

**Evaluation Benchmarks.** Following Xu et al. (2025b); Trivedi et al. (2023); Shao et al. (2023), we evaluate the performance of our MARCH on three popular multi-hop QA datasets, including HotpotQA (Yang et al., 2018), 2WikiMulti-HopQA (Ho et al., 2020) and MuSiQue (Trivedi et al., 2022). More details about datasets are provided in Appendix B.2.3.

**Evaluation Protocol and Metrics.** Following Xu et al. (2025b); Shao et al. (2023); Trivedi et al. (2023), we sample the first 500 examples from the development set of each dataset as our evaluation subset to reduce the evaluation overhead while maintaining statistical reliability. For each question, the model receives the query and its associated context passages and produces a corresponding answer. We report accuracy (%) with respect to the GT reference answers judged by a Qwen3-235B-A22B, using the official evaluation scripts when available. We also draw eight independent generations for each test query and then adopt a majority voting to decide the final verdict.

**Comparison with Existing Baselines.** Table 4 compares MARCH-Best, which incorporates joint optimization of both the Solver and the Checker alongside CoT prompting, with various specialized RAG methods and proprietary models. Within the 8B-parameter scale, MARCH achieves higher accuracy on the HotpotQA and MuSiQue benchmarks compared to methods such as RAFT (Dong et al., 2023) and GenGround (Shi et al., 2024). Notably,

Table 4: Comparison of MARCH-Best against RAG baselines. MARCH-Best denotes our best-performing model. Underline indicates the best performance among 8B-class models, and **Bold** indicates the best overall performance. Baseline results are cited from Xu et al. (2025b).

Method	Backbone	HotpotQA	MuSiQue	2WikiMQA
<i>Advanced RAG Frameworks with Proprietary Models</i>				
CoT (Wei et al., 2022)	GPT-4o	56.2	26.8	58.4
RAG-Standard (Lewis et al., 2020)	GPT-4o	64.0	29.8	57.8
IRCoT (Trivedi et al., 2022)	GPT-4o	66.4	<b>44.2</b>	<b>78.0</b>
Iter-RetGen (Shao et al., 2023)	GPT-4o	62.6	42.0	71.4
RAG-Star (Jiang et al., 2024)	GPT-4o	57.0	40.0	63.0
RQ-RAG (Chan et al., 2024)	GPT-4o	62.0	31.6	60.2
<i>Specialized RAG Micro-tuning (White-box 7B/8B)</i>				
Baseline	Llama3.1-8B	35.0	5.6	17.4
GenGround (Shi et al., 2024)	7B-class	47.2	20.2	45.6
RAFT (Dong et al., 2023)	Llama3.1-8B	51.2	22.0	44.6
IRCoT (Trivedi et al., 2022)	Llama3.1-8B	64.0	29.2	56.8
MARCH (10-Shots)	Llama3.1-8B	<b>73.6</b>	<b>40.8</b>	69.4
MARCH (CoT)	Llama3.1-8B	70.6	36.2	70.6

on the HotpotQA benchmark, which contains a high density of noisy documents, MARCH reaches an accuracy of 71.2%, exceeding the reported results of standard GPT-4o RAG (64.0%) and the IRCoT (Trivedi et al., 2023) framework using GPT-4o (66.4%). Table 4 suggests that MARCH can effectively enhance the factual reliability of smaller models in multi-hop reasoning scenarios, leading to performance that is comparable to and even higher than much larger proprietary models.

### 3.4 Ablation Study and Further Investigation

#### Performance Analysis of Joint Optimization.

Figure 3 illustrates that the joint optimization of the Solver and the Checker consistently achieves the highest accuracy across all training strategies in both STEM and General benchmarks. A key observation is that while the Solver-only update provides a base improvement, the inclusion of the Checker significantly amplifies the model’s reason-

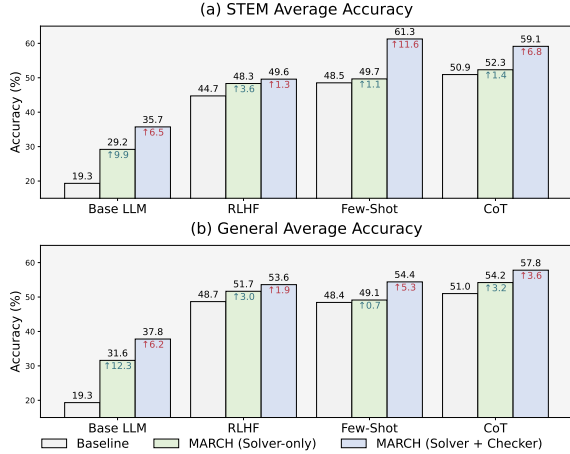


Figure 3: Averaged results on MuSuQueQA, 2WikiMultiHopQA and HotpotQA benchmarks across STEM (a) and General (b) training datasets based on Llama3.1-8B-Instruct. Comprehensive results are in Table 8.

ing capabilities, particularly in complex scenarios. Specifically, in the *Few-Shot* setting for STEM, the joint model achieves a remarkable absolute gain of 11.6% over the Solver-only version. These results suggest that the Checker’s auditing signal is crucial for leveraging the full potential of advanced prompting techniques like *few-shot* and *CoT*, as it effectively mitigates factual drift and ensures logical consistency throughout multi-hop reasoning trajectories. More detailed results and analysis are provided in Appendix B.4 and Table 8.

**Ablation Study on Reward Design.** We evaluate two reward formulations: (i) Error Rate Reward (ERR), a proportional penalty based on the ratio of incorrect claims, and (ii) Zero-Tolerance Reward (ZTR), the binary constraint that enforces absolute consistency. As shown in Table 5, while all RL-tuned configurations significantly outperform the vanilla baselines, ZTR consistently surpasses ERR (e.g., 61.25% vs 55.46% on Llama3.1 joint-training), suggesting that strict logical consistency is more effective for complex RAG tasks. Within the ZTR, the penalty-based scalar (-1/0) yields superior results compared to the incentive-based (0/1) setting, as the former provides a stronger corrective gradient to suppress hallucinations. Detailed analyses and task-specific breakdowns are provided in Appendix B.5.

**Performance on Different Model Family.** As shown in the last group in Table 5, the MARCH framework exhibits robust generalizability across different model families, with Qwen3-8B reflecting the same performance trajectory observed in

Table 5: Ablation study on reward mechanisms and configurations. **Hallu.** denotes the average of RAGTruth and Facts Grounding; **QA** denotes the average of MuSiQue, 2Wiki, and HotpotQA.

Configuration	Reward	Hallu.	QA	Avg.
Llama3.1-8B Baseline	-	55.10	19.33	39.77
<i>Ablation: Reward Function (General)</i>				
Solver-only	ERR	69.03	29.13	51.93
Solver-only	ZTR (-1/0)	79.55	31.60	59.00
MARCH (Full)	ERR	75.81	28.33	55.46
<b>MARCH (Full)</b>	<b>ZTR (-1/0)</b>	<b>79.59</b>	<b>37.80</b>	<b>61.25</b>
<i>Ablation: Reward Scalar (STEM)</i>				
Solver-only	ZTR (0/1)	68.12	20.63	49.38
Solver-only	ZTR (-1/0)	73.19	22.53	53.87
MARCH (Full)	ZTR (0/1)	69.53	24.93	50.42
<b>MARCH (Full)</b>	<b>ZTR (-1/0)</b>	<b>76.56</b>	<b>35.73</b>	<b>59.06</b>
<i>Ablation: Base LLM</i>				
Qwen3-8B Baseline	-	54.02	67.30	56.84
+MARCH (Full, General)	ZTR (-1/0)	65.05	71.77	67.90
+MARCH Full, STEM)	ZTR (-1/0)	65.24	71.93	68.11

Llama-based experiments. Specifically, MARCH elevates Qwen3-8B to average accuracies of 67.90% and 68.11% under General and STEM settings respectively, marking a significant absolute gain of over 11% compared to its vanilla baseline (56.84%) and demonstrating that the dual-trajectory optimization mechanism is model-agnostic and consistently internalizes superior evidentiary grounding across diverse architectures.

**How Would Proposed Question Numbers Influence Performance?** A potential reward hacking of the MARCH framework could be that the system may attempt to inflate its reward by reducing the number of proposed questions, adhering to a “less said, less leaked” strategy where fewer claims minimize the probability of being audited as a hallucination. As illustrated in Figure 4, we observe a gradual decline in the average number of proposed questions as training progresses across all configurations. Notably, the configuration involving joint updates of the Solver + Checker exhibits a more pronounced downward trend compared to the Solver Only baseline, suggesting that the ZTR on the Checker’s trajectory imposes a stricter consistency constraint, further incentivizing the *Proposer* to prune potentially risky claims.

However, we find that this quantitative shrinkage can be effectively mitigated through simple instructional constraints. By explicitly requiring the *Proposer* to generate no fewer than  $k$  questions within the prompt (e.g.,  $k = 3$ ), the decline in question density is significantly arrested. As shown by the solid lines in Figure 4, models with

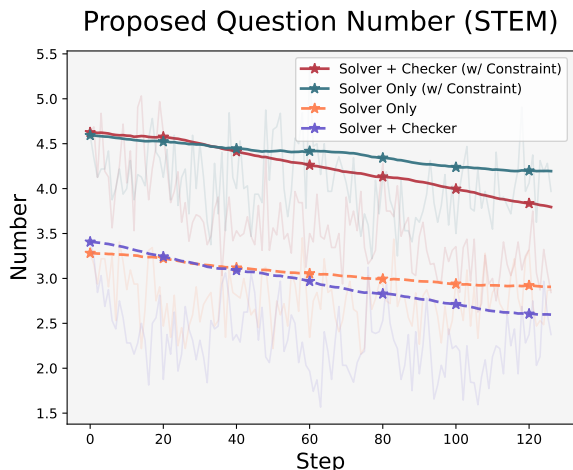


Figure 4: Average number of proposed questions per step on the STEM training dataset. The dashed lines represent the MARCH framework without maintaining the question number, while the solid lines (w/ Constraint) indicate the results after incorporating instructional constraints to maintain informational density.

such constraints in the legend (solid lines marked with “w/ constraint”) maintain a much more stable and higher volume of factual inquiries throughout the self-evolution process, which indicates that the MARCH framework remains robust when balanced with basic task-specific instructions, ensuring that the drive for factual accuracy does not come at the cost of informational density.

### 3.5 Additional Analysis

We show MARCH can be further boosted with advanced methods to achieve higher performance (Appendix B.6). We visualize the PPO training dynamics and analyze the time consumption (Appendix B.7), where we find the consistently increasing rewards and acceptable trade-off between training consuming and improved performance. We present cases to show outputs of each agent during one training step (Appendix C).

## 4 Related Work

**Hallucinations and Fact-Check in RAG.** While RAG reduces fabrications by anchoring responses in external context (Lewis et al., 2020), LLMs remain susceptible to context-conflicting hallucinations (Shuster et al., 2021). Current mitigation efforts largely revolve around outcome-level supervision (Shi et al., 2023; Gekhman et al., 2024; Yu et al., 2024a; Asai et al., 2023), post-hoc reasoning refinements and utilizing external verifiers to filter outputs during inference (Dhuliawala et al., 2023;

Madaan et al., 2023; Min et al., 2023; Rashkin et al., 2023; Gao et al., 2023; Tang et al., 2024). However, these methods often treat LLM as a black box and fail to internalize factual constraints within the training objective itself. Moving beyond, we propose a framework where the model learns to prioritize factual adherence by minimizing logical discrepancies identified through an internal audit.

**Multi-Agent Systems for LLMs.** Multi-agent systems have emerged as a robust paradigm for enhancing LLM performance through collaborative or adversarial interactions (Du et al., 2023; Li et al., 2023; Hu et al., 2025; Li et al., 2025c,b), which typically assign distinct personas to one / multiple model instances to facilitate code generation (Huang et al., 2023; Islam et al., 2024), complex reasoning (Li et al., 2023; He et al., 2023; Talebirad and Nadiri, 2023), decision-making (Shinn et al., 2023; Yu et al., 2024b), recommendation (Wang et al., 2024b), and debate-driven consensus (Du et al., 2023; Liang et al., 2024). While effective at inference time, most existing architectures treat agents as separate entities, often overlooking the potential for joint policy optimization across different functional roles. In addition, previous work shows that multi-agent system sometimes amplifies the hallucination of LLMs (Hong et al., 2023). Our framework differs by instantiating functionally distinct roles within a single shared policy. By structurally enforcing an information barrier between agents, we facilitate the co-evolution of generative and evaluative capabilities within a unified training loop, ensuring that the model’s internal verification logic directly informs its generative precision.

## 5 Conclusion

This paper presented MARCH to mitigate hallucinations in data-intensive RAG tasks. By transforming a shared policy into a coordinated hierarchy of the Solver, the Proposer, and the Checker, MARCH enables LLMs to concurrently evolve reasoning synthesis and independent grounding capabilities. Our results across various benchmarks demonstrate that the Zero-Tolerance Reward and dual-trajectory PPO effectively internalize a self-contained verification loop, yielding significant performance gains over strong baselines. MARCH’s generalization across different model families underscores the potential as a scalable solution for high-fidelity information synthesis in complex knowledge domains.

## 6 Limitations

Despite the performance gains, our study has several limitations that warrant further investigation: (1) Focus on Factuality over Creativity: Our optimization objective, particularly the Zero-Tolerance Reward (ZTR), prioritizes rigorous factual consistency and numerical precision. (2) Evaluation Scope: Our evaluation focused primarily on English-language benchmarks and numerical-heavy domains. The framework’s cross-lingual transferability and performance on subjective or low-resource language tasks remain to be explored in future work. (3) Reliance on LLM-as-a-Judge: Our evaluation employs LLM-as-a-judge (e.g., Qwen3-235B-A22B), which may carry inherent biases, such as a preference for certain linguistic styles or the “self-preference” bias.

## References

- Andrei Alexandru, Antonia Calvi, Henry Broomfield, Jackson Golden, Kyle Dai, Mathias Leys, Maurice Burger, Max Bartolo, Roman Engeler, Sashank Pisu-pati, Toby Drane, and Young Sun Park. 2025. [Atla selene mini: A general purpose evaluation model](#). *Preprint*, arXiv:2501.17195.
- Akari Asai, Zeqiu Wu, Yizhong Wang, Avirup Sil, and Hannaneh Hajishirzi. 2023. [Self-rag: Learning to retrieve, generate, and critique through self-reflection](#). *Preprint*, arXiv:2310.11511.
- Bowen Baker, Joost Huizinga, Leo Gao, Zehao Dou, Melody Y. Guan, Aleksander Madry, Wojciech Zaremba, Jakub Pachocki, and David Farhi. 2025. [Monitoring reasoning models for misbehavior and the risks of promoting obfuscation](#). *Preprint*, arXiv:2503.11926.
- 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.
- Stephen Casper, Xander Davies, Claudia Shi, Thomas Krendl Gilbert, Jérémy Scheurer, Javier Rando, Rachel Freedman, Tomasz Korbak, David Lindner, Pedro Freire, Tony Wang, Samuel Marks, Charbel-Raphaël Segerie, Micah Carroll, Andi Peng, Phillip Christoffersen, Mehul Damani, Stewart Slocum, Usman Anwar, and 13 others. 2023. [Open problems and fundamental limitations of reinforcement learning from human feedback](#). *Preprint*, arXiv:2307.15217.
- Chi-Min Chan, Chunpu Xu, Ruibin Yuan, Hongyin Luo, Wei Xue, Yike Guo, and Jie Fu. 2024. [Rq-rag: Learning to refine queries for retrieval augmented generation](#). *Preprint*, arXiv:2404.00610.
- Imre Csiszár. 1975. I-divergence geometry of probability distributions and minimization problems. *The annals of probability*, pages 146–158.
- Matthew Dahl, Varun Magesh, Mirac Suzgun, and Daniel E Ho. 2024. Large legal fictions: Profiling legal hallucinations in large language models. *Journal of Legal Analysis*, 16(1):64–93.
- Chongyuan Dai, Jinpeng Hu, Hongchang Shi, Zhuo Li, Xun Yang, and Meng Wang. 2025. [Psyche-r1: Towards reliable psychological llms through unified empathy, expertise, and reasoning](#). *Preprint*, arXiv:2508.10848.
- DeepSeek-AI. 2025. [Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning](#). *Preprint*, arXiv:2501.12948.
- Darshan Deshpande, Selvan Sunitha Ravi, Sky CH-Wang, Bartosz Mielczarek, Anand Kannappan, and Rebecca Qian. 2024. [Glider: Grading llm interactions and decisions using explainable ranking](#). *Preprint*, arXiv:2412.14140.
- Shehzaad Dhuliawala, Mojtaba Komeili, Jing Xu, Roberta Raileanu, Xian Li, Asli Celikyilmaz, and Jason Weston. 2023. [Chain-of-verification reduces hallucination in large language models](#). *Preprint*, arXiv:2309.11495.
- Hanze Dong, Wei Xiong, Deepanshu Goyal, Yihan Zhang, Winnie Chow, Rui Pan, Shizhe Diao, Jipeng Zhang, Kashun Shum, and Tong Zhang. 2023. [Raft: Reward ranked finetuning for generative foundation model alignment](#). *Preprint*, arXiv:2304.06767.
- Yilun Du, Shuang Li, Antonio Torralba, Joshua B. Tenenbaum, and Igor Mordatch. 2023. [Improving factuality and reasoning in language models through multiagent debate](#). *Preprint*, arXiv:2305.14325.
- Yuhao Du, Zhuo Li, Pengyu Cheng, Zhihong Chen, Yuejiao Xie, Xiang Wan, and Anningzhe Gao. 2026. [Rlhf in an sft way: From optimal solution to reward-weighted alignment](#). *Preprint*, arXiv:2502.11026.
- Yuhao Du, Zhuo Li, Pengyu Cheng, Xiang Wan, and Anningzhe Gao. 2025. [Atoxia: Red-teaming large language models with target toxic answers](#). In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 3251–3266, Albuquerque, New Mexico. Association for Computational Linguistics.
- Zhibin Duan, Guowei Rong, Zhuo Li, Bo Chen, Mingyuan Zhou, and Dandan Guo. 2026. [Mitigating reward hacking in rlhf via bayesian non-negative reward modeling](#). *Preprint*, arXiv:2602.10623.
- Mohamed Elaraby, Mengyin Lu, Jacob Dunn, Xueying Zhang, Yu Wang, Shizhu Liu, Pingchuan Tian, Yuping Wang, and Yuxuan Wang. 2023. [Halo: Estimation and reduction of hallucinations in open-source weak large language models](#). *Preprint*, arXiv:2308.11764.

- Luyu Gao, Zhuyun Dai, Panupong Pasupat, Anthony Chen, Arun Tejasvi Chaganty, Yicheng Fan, Vincent Zhao, Ni Lao, Hongrae Lee, Da-Cheng Juan, and Kelvin Guu. 2023. [RARR: Researching and revising what language models say, using language models](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 16477–16508, Toronto, Canada. Association for Computational Linguistics.
- Zorik Gekhman, Gal Yona, Roei Aharoni, Matan Eyal, Amir Feder, Roi Reichart, and Jonathan Herzig. 2024. [Does fine-tuning llms on new knowledge encourage hallucinations?](#) *Preprint*, arXiv:2405.05904.
- Gemini. 2025. [Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities](#). *Preprint*, arXiv:2507.06261.
- Gemma. 2025. [Gemma 3 technical report](#). *Preprint*, arXiv:2503.19786.
- Melody Y. Guan, Miles Wang, Micah Carroll, Zehao Dou, Annie Y. Wei, Marcus Williams, Benjamin Arnav, Joost Huizinga, Ian Kivlichan, Mia Glaese, Jakub Pachocki, and Bowen Baker. 2025. [Monitoring monitorability](#). *Preprint*, arXiv:2512.18311.
- Anisha Gunjal, Anthony Wang, Elaine Lau, Vaskar Nath, Yunzhong He, Bing Liu, and Sean Hendryx. 2025. [Rubrics as rewards: Reinforcement learning beyond verifiable domains](#). *Preprint*, arXiv:2507.17746.
- Zhitao He, Pengfei Cao, Yubo Chen, Kang Liu, Ruopeng Li, Mengshu Sun, and Jun Zhao. 2023. [Lego: A multi-agent collaborative framework with role-playing and iterative feedback for causality explanation generation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 9142–9163.
- Xanh Ho, Anh-Khoa Duong Nguyen, Saku Sugawara, and Akiko Aizawa. 2020. [Constructing a multi-hop qa dataset for comprehensive evaluation of reasoning steps](#). *arXiv preprint arXiv:2011.01060*.
- Sirui Hong, Mingchen Zhuge, Jonathan Chen, Xiawu Zheng, Yuheng Cheng, Jinlin Wang, Ceyao Zhang, Zili Wang, Steven Ka Shing Yau, Zijuan Lin, and 1 others. 2023. [Metagpt: Meta programming for a multi-agent collaborative framework](#). In *The Twelfth International Conference on Learning Representations*.
- Jinpeng Hu, Zhuo Li, Zhihong Chen, Zhen Li, Xiang Wan, and Tsung-Hui Chang. 2022. [Graph enhanced contrastive learning for radiology findings summarization](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4677–4688, Dublin, Ireland. Association for Computational Linguistics.
- Jinpeng Hu, Hongchang Shi, Chongyuan Dai, Zhuo Li, Peipei Song, and Meng Wang. 2025. [Beyond emotion recognition: A multi-turn multimodal emotion understanding and reasoning benchmark](#). In *Proceedings of the 33rd ACM International Conference on Multimedia*, pages 5814–5823.
- Jinpeng Hu, Ao Wang, Qianqian Xie, Zhuo Li, Hui Ma, and Dan Guo. 2026. [Agentmental: An interactive multi-agent framework for explainable and adaptive mental health assessment](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pages 31050–31058.
- Dong Huang, Jie M Zhang, Michael Luck, Qingwen Bu, Yuhao Qing, and Heming Cui. 2023. [Agent-coder: Multi-agent-based code generation with iterative testing and optimisation](#). *arXiv preprint arXiv:2312.13010*.
- Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, and Ting Liu. 2025. [A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions](#). *ACM Transactions on Information Systems*, 43(2):1–55.
- Md Ashraf Islam, Mohammed Eunus Ali, and Md Rizwan Parvez. 2024. [Mapcoder: Multi-agent code generation for competitive problem solving](#). *arXiv preprint arXiv:2405.11403*.
- Alon Jacovi, Andrew Wang, Chris Alberti, Connie Tao, Jon Lipovetz, Kate Olszewska, Lukas Haas, Michelle Liu, Nate Keating, Adam Bloniarz, and 1 others. 2025. [The facts grounding leaderboard: Benchmarking llms’ ability to ground responses to long-form input](#). *arXiv preprint arXiv: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*, 55(12):1–38.
- Jinhao Jiang, Jiayi Chen, Junyi Li, Ruiyang Ren, Shijie Wang, Wayne Xin Zhao, Yang Song, and Tao Zhang. 2024. [Rag-star: Enhancing deliberative reasoning with retrieval augmented verification and refinement](#). *Preprint*, arXiv:2412.12881.
- Adam Tauman Kalai, Ofir Nachum, Santosh S Vempala, and Edwin Zhang. 2025. [Why language models hallucinate](#). *arXiv preprint arXiv:2509.04664*.
- Haoqiang Kang and Xiao-Yang Liu. 2023. [Deficiency of large language models in finance: An empirical examination of hallucination](#). *arXiv preprint arXiv:2311.15548*.
- 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 the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 4334–4353, Miami, Florida, USA. Association for Computational Linguistics.
- Yubin Kim, Hyewon Jeong, Shan Chen, Shuyue Stella Li, Chanwoo Park, Mingyu Lu, Kumail Alhamoud, Jimin Mun, Cristina Grau, Minseok Jung, and 1 others. 2025. Medical hallucinations in foundation models and their impact on healthcare. *arXiv preprint arXiv:2503.05777*.
- Anastasia Krithara, Anastasios Nentidis, Konstantinos Bougiatiotis, and Georgios Paliouras. 2023. Bioasqqa: A manually curated corpus for biomedical question answering. *Scientific Data*, 10(1):170.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient memory management for large language model serving with pagedattention. *Preprint*, arXiv:2309.06180.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, and 1 others. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in neural information processing systems*, 33:9459–9474.
- Guohao Li, Hasan Hammoud, Hani Itani, Dmitrii Khizbullin, and Bernard Ghanem. 2023. Camel: Communicative agents for "mind" exploration of large language model society. *Advances in Neural Information Processing Systems*, 36:51991–52008.
- Junyi Li and Hwee Tou Ng. 2025. Reasoning models hallucinate more: Factuality-aware reinforcement learning for large reasoning models. *Preprint*, arXiv:2505.24630.
- Zhuo Li, Pengyu Cheng, Zhechao Yu, Feifei Tong, Anningzhe Gao, Tsung-Hui Chang, Xiang Wan, Erchao Zhao, Xiaoxi Jiang, and Guanjun Jiang. 2025a. Eliminating inductive bias in reward models with information-theoretic guidance. *Preprint*, arXiv:2512.23461.
- Zhuo Li, Yuhao Du, Jinpeng Hu, Xiang Wan, and Anningzhe Gao. 2025b. Self-instructed derived prompt generation meets in-context learning: Unlocking new potential of black-box LLMs. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1840–1857, Vienna, Austria. Association for Computational Linguistics.
- Zhuo Li, Yuhao Du, Xiaoqi Jiao, Steven Y. Guo, Yuege Feng, Xiang Wan, Anningzhe Gao, and Jinpeng Hu. 2025c. Add-one-in: Incremental sample selection for large language models via a choice-based greedy paradigm. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 5321–5340, Suzhou, China. Association for Computational Linguistics.
- Zhuo Li, Yuege Feng, Dandan Guo, Jinpeng Hu, Anningzhe Gao, and Xiang Wan. 2025d. APLOT: Robust reward modeling via adaptive preference learning with optimal transport. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 5524–5538, Suzhou, China. Association for Computational Linguistics.
- Tian Liang, Zhiwei He, Wenxiang Jiao, Xing Wang, Yan Wang, Rui Wang, Yujiu Yang, Shuming Shi, and Zhaopeng Tu. 2024. Encouraging divergent thinking in large language models through multi-agent debate. In *Proceedings of the 2024 conference on empirical methods in natural language processing*, pages 17889–17904.
- Stephanie Lin, Jacob Hilton, and Owain Evans. 2022. Truthfulqa: Measuring how models mimic human falsehoods. *Preprint*, arXiv:2109.07958.
- Chris Yuhao Liu, Liang Zeng, Jiakai Liu, Rui Yan, Jujie He, Chaojie Wang, Shuicheng Yan, Yang Liu, and Yahui Zhou. 2024a. Skywork-reward: Bag of tricks for reward modeling in llms. *Preprint*, arXiv:2410.18451.
- Chris Yuhao Liu, Liang Zeng, Yuzhen Xiao, Jujie He, Jiakai Liu, Chaojie Wang, Rui Yan, Wei Shen, Fuxiang Zhang, Jiacheng Xu, Yang Liu, and Yahui Zhou. 2025a. Skywork-reward-v2: Scaling preference data curation via human-ai synergy. *Preprint*, arXiv:2507.01352.
- Nelson F Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2024b. Lost in the middle: How language models use long contexts. *Transactions of the Association for Computational Linguistics*, 12:157–173.
- Zijun Liu, Peiyi Wang, Runxin Xu, Shirong Ma, Chong Ruan, Peng Li, Yang Liu, and Yu Wu. 2025b. Inference-time scaling for generalist reward modeling. *Preprint*, arXiv:2504.02495.
- Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegrefe, Uri Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, Shashank Gupta, Bodhisattwa Prasad Majumder, Katherine Hermann, Sean Welleck, Amir Yazdanbakhsh, and Peter Clark. 2023. Self-refine: Iterative refinement with self-feedback. *Preprint*, arXiv:2303.17651.
- Joshua Maynez, Shashi Narayan, Bernd Bohnet, and Ryan McDonald. 2020. On faithfulness and factuality in abstractive summarization. *Preprint*, arXiv:2005.00661.
- Sewon Min, Kalpesh Krishna, Xinxin Lyu, Mike Lewis, Wen-tau Yih, Pang Koh, Mohit Iyyer, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2023. FActScore: Fine-grained atomic evaluation of factual precision in long form text generation. In *Proceedings of the*

- 2023 Conference on Empirical Methods in Natural Language Processing, pages 12076–12100, Singapore. Association for Computational Linguistics.
- Cheng Niu, Yuanhao Wu, Juno Zhu, Siliang Xu, Kashun Shum, Randy Zhong, Juntong Song, and Tong Zhang. 2024. Ragtruth: A hallucination corpus for developing trustworthy retrieval-augmented language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 10862–10878.
- OpenAI. 2024. [Gpt-4 technical report](#). *Preprint*, arXiv:2303.08774.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, and 1 others. 2022. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744.
- Junsoo Park, Seungyeon Jwa, Ren Meiying, Daeyoung Kim, and Sanghyuk Choi. 2024. [OffsetBias: Leveraging debiased data for tuning evaluators](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 1043–1067, Miami, Florida, USA. Association for Computational Linguistics.
- Hannah Rashkin, Vitaly Nikolaev, Matthew Lamm, Lora Aroyo, Michael Collins, Dipanjan Das, Slav Petrov, Gaurav Singh Tomar, Iulia Turc, and David Reitter. 2023. Measuring attribution in natural language generation models. *Computational Linguistics*, 49(4):777–840.
- John Schulman, Philipp Moritz, Sergey Levine, Michael Jordan, and Pieter Abbeel. 2015. High-dimensional continuous control using generalized advantage estimation. *arXiv preprint arXiv:1506.02438*.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. [Proximal policy optimization algorithms](#). *Preprint*, arXiv:1707.06347.
- Wooseok Seo, Seungju Han, Jaehun Jung, Benjamin Newman, Seungwon Lim, Seungbeen Lee, Ximing Lu, Yejin Choi, and Youngjae Yu. 2025. [Verifying the verifiers: Unveiling pitfalls and potentials in fact verifiers](#). *Preprint*, arXiv:2506.13342.
- Zhihong Shao, Yeyun Gong, Yelong Shen, Minlie Huang, Nan Duan, and Weizhu Chen. 2023. Enhancing retrieval-augmented large language models with iterative retrieval-generation synergy. *arXiv preprint arXiv:2305.15294*.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. 2024. [Deepseekmath: Pushing the limits of mathematical reasoning in open language models](#). *Preprint*, arXiv:2402.03300.
- Mrinank Sharma, Meg Tong, Tomasz Korbak, David Duvenaud, Amanda Askell, Samuel R. Bowman, Newton Cheng, Esin Durmus, Zac Hatfield-Dodds, Scott R. Johnston, Shauna Kravec, Timothy Maxwell, Sam McCandlish, Kamal Ndousse, Oliver Rausch, Nicholas Schiefer, Da Yan, Miranda Zhang, and Ethan Perez. 2024. Towards understanding sycophancy in language models. In *The Twelfth International Conference on Learning Representations*.
- Guangming Sheng, Chi Zhang, Zilingfeng Ye, Xibin Wu, Wang Zhang, Ru Zhang, Yanghua Peng, Haibin Lin, and Chuan Wu. 2025. [Hybridflow: A flexible and efficient rlhf framework](#). In *Proceedings of the Twentieth European Conference on Computer Systems*, EuroSys '25, page 1279–1297. ACM.
- 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](#). *Preprint*, arXiv:2301.12652.
- Zhengliang Shi, Shuo Zhang, Weiwei Sun, Shen Gao, Pengjie Ren, Zhumin Chen, and Zhaochun Ren. 2024. Generate-then-ground in retrieval-augmented generation for multi-hop question answering. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7339–7353.
- Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. 2023. Reflexion: Language agents with verbal reinforcement learning. *Advances in Neural Information Processing Systems*, 36:8634–8652.
- Kurt Shuster, Spencer Poff, Moya Chen, Douwe Kiela, and Jason Weston. 2021. [Retrieval augmentation reduces hallucination in conversation](#). *Preprint*, arXiv:2104.07567.
- Yashar Talebirad and Amirhossein Nadiri. 2023. Multi-agent collaboration: Harnessing the power of intelligent llm agents. *arXiv preprint arXiv:2306.03314*.
- Manveer Singh Tamber, Forrest Bao, Chenyu Xu, Ge Luo, Suleman Kazi, Minseok Bae, Miaoran Li, Ofer Mendelevitch, Renyi Qu, and Jimmy Lin. 2025. Benchmarking llm faithfulness in rag with evolving leaderboards. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing: Industry Track*, pages 799–811.
- Liyan Tang, Philippe Laban, and Greg Durrett. 2024. Minicheck: Efficient fact-checking of llms on grounding documents. *arXiv preprint arXiv:2404.10774*.
- Faraz Torabi, Garrett Warnell, and Peter Stone. 2018. Behavioral cloning from observation. *arXiv preprint arXiv:1805.01954*.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard

- Grave, and Guillaume Lample. 2023. [Llama: Open and efficient foundation language models](#). *Preprint*, arXiv:2302.13971.
- 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*, 10:539–554.
- 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 the 61st annual meeting of the association for computational linguistics (volume 1: long papers)*, pages 10014–10037.
- Peifeng Wang, Austin Xu, Yilun Zhou, Caiming Xiong, and Shafiq Joty. 2025. [Direct judgement preference optimization](#). *Preprint*, arXiv:2409.14664.
- Tianlu Wang, Ilia Kulikov, Olga Golovneva, Ping Yu, Weizhe Yuan, Jane Dwivedi-Yu, Richard Yuanzhe Pang, Maryam Fazel-Zarandi, Jason Weston, and Xian Li. 2024a. [Self-taught evaluators](#). *Preprint*, arXiv:2408.02666.
- Zhefan Wang, Yuanqing Yu, Wendi Zheng, Weizhi Ma, and Min Zhang. 2024b. Macrec: A multi-agent collaboration framework for recommendation. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2760–2764.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, and 1 others. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.
- Austin Xu, Srijan Bansal, Yifei Ming, Semih Yavuz, and Shafiq Joty. 2025a. Does context matter? contextual-judgebench for evaluating llm-based judges in contextual settings. *arXiv preprint arXiv:2503.15620*.
- Ran Xu, Wenqi Shi, Yuchen Zhuang, Yue Yu, Joyce C Ho, Haoyu Wang, and Carl Yang. 2025b. Collab-rag: Boosting retrieval-augmented generation for complex question answering via white-box and black-box llm collaboration. *arXiv preprint arXiv:2504.04915*.
- Ziwei Xu, Sanjay Jain, and Mohan Kankanhalli. 2025c. [Hallucination is inevitable: An innate limitation of large language models](#). *Preprint*, arXiv:2401.11817.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, and 41 others. 2025. [Qwen3 technical report](#). *Preprint*, arXiv:2505.09388.
- An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, and 1 others. 2024. Qwen2.5 technical report. *arXiv preprint arXiv:2412.15115*.
- 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 *Proceedings of the 2018 conference on empirical methods in natural language processing*, pages 2369–2380.
- Wenhao Yu, Hongming Zhang, Xiaoman Pan, Kaixin Ma, Hongwei Wang, and Dong Yu. 2024a. [Chain-of-note: Enhancing robustness in retrieval-augmented language models](#). *Preprint*, arXiv:2311.09210.
- Yangyang Yu, Zhiyuan Yao, Haohang Li, Zhiyang Deng, Yuechen Jiang, Yupeng Cao, Zhi Chen, Jordan Suchow, Zhenyu Cui, Rong Liu, and 1 others. 2024b. Fincon: A synthesized llm multi-agent system with conceptual verbal reinforcement for enhanced financial decision making. *Advances in Neural Information Processing Systems*, 37:137010–137045.
- Yue Zhang, Yafu Li, Leyang Cui, Deng Cai, Lemao Liu, Tingchen Fu, Xinting Huang, Enbo Zhao, Yu Zhang, Chen Xu, Yulong Chen, Longyue Wang, Anh Tuan Luu, Wei Bi, Freda Shi, and Shuming Shi. 2025. [Siren’s song in the ai ocean: A survey on hallucination in large language models](#). *Preprint*, arXiv:2309.01219.
- Yanli Zhao, Andrew Gu, Rohan Varma, Liang Luo, Chien-Chin Huang, Min Xu, Less Wright, Hamid Shojanazeri, Myle Ott, Sam Shleifer, Alban Desmaison, Can Balioglu, Pritam Damania, Bernard Nguyen, Geeta Chauhan, Yuchen Hao, Ajit Mathews, and Shen Li. 2023. [Pytorch fsdp: Experiences on scaling fully sharded data parallel](#). *Preprint*, arXiv:2304.11277.
- Chunting Zhou, Pengfei Liu, Puxin Xu, Srinu Iyer, Jiao Sun, Yuning Mao, Xuezhe Ma, Avia Efrat, Ping Yu, Lili Yu, Susan Zhang, Gargi Ghosh, Mike Lewis, Luke Zettlemoyer, and Omer Levy. 2023. [Lima: Less is more for alignment](#). *Preprint*, arXiv:2305.11206.

## A Methodology

### A.1 Prompts

#### A.1.1 System Prompt of the *Solver* Agent

**Instruction:**

You must respond based strictly on the information in provided passages. Do not incorporate any external knowledge or infer any details beyond what is given in the passages.

**Passages  $D$ :**

{Document-1}

...

{Document-K}

**Query  $x$ :**

{User Query  $x$ }

**Output:**

#### A.1.2 System Prompt of the *Proposer* Agent

**You are a data annotation expert.** Your task is to analyze a model's response to a user query. For every critical number found in the model's response, you must create a question where the answer is precisely and exactly that number.

**Rule 1. Clarity and Specificity of Questions:** Each question must be clear, complete, and unambiguous. It must contain enough context from the response to ensure the number is the only possible correct answer. For example, if the model's response contains "In 2025, 50 people will take the bar exam in Beijing," an effective question would be "How many people will take the bar exam in the Beijing area in 2025?", and the answer is 50. A bad example would be "How many people will take the bar exam in 2025?". This is ambiguous because it lacks the specific region ('Beijing area').

**Rule 2. Format of the Answer:** The answer must be a pure number, either an integer or a decimal. It must NOT include any other characters or formats, such as percentage signs (%), 'k' to denote thousands (e.g., 10k), ranges (e.g., 10-20), or words (e.g., fifty).

Please provide your output as an unordered list using the following format for each question-answer pair:

- Question: xxx [Answer: n]

- Question: xxx [Answer: n]

Now, generate questions for the user query and model response below. Do not output any other content or explanations.

**Model Response:**

{response  $y$ }

### A.1.3 System Prompt of the *Checker Agent*

**You are a data validation expert.** You will be given a set of reference materials and a list of user questions. Your task is to retrieve information from the reference materials to answer all questions. Before providing each answer, you must state the evidence for it.

You must adhere strictly to the content of the reference materials and must not fabricate or infer any information that is not explicitly mentioned. If the relevant information cannot be found in the materials, you must state “Cannot answer”. It is important to note that all answers are pure numbers. Therefore, your answer must also be a pure number and must not contain any other content, such as percentage signs (%), 'k' to denote thousands, ranges, or numbers written as words.

Here is an example:

**Reference Materials:**

Document 1: ...in 2024, 50 people will take the bar exam in Beijing.

Document 2: ...(omitted)

- Question 1: How many people will take the bar exam in Beijing in 2024?

- Question 2: How many people will take the bar exam in Beijing in 2025?

**Example Response:**

1. Evidence: Document 1 states that 50 people will take the bar exam in Beijing in 2024.

[Answer: 50]

2. Evidence: The materials do not contain information about the bar exam in Beijing for 2025, therefore I cannot answer.

[Answer: Cannot answer]

Now, based on the reference materials below, answer all the following questions. Strictly adhere to the example format provided.

**Reference Materials:**

{docs  $D$ }

**Questions:**

{questions  $q^{\text{aux}}$ }

### A.2 Training Algorithm

---

**Algorithm 1: MARCH Training Process.**

---

1 **Input:** An LLM Policy  $\pi_\theta$ ; A Training Batch  $\mathcal{B} = \{(\mathbf{x}_i, \mathbf{D}_i)\}_{i=1}^{|\mathcal{B}|}$ ; Solver, Proposer and Checker  
System Prompts  $\{\mathbf{s}_{\text{solve}}, \mathbf{s}_{\text{propose}}, \mathbf{s}_{\text{check}}\}$ ;  
// Phase 1: Multi-Agent Collaboration and Rollout  
2 **for** each training sample  $(\mathbf{x}, \mathbf{D}) \in \mathcal{B}$  **do**  
    // Step 1: Generate reasoning path  $\mathbf{y}$  from Solver  
3     Sample  $\mathbf{y} \sim \nu_{\text{solve}}(\cdot | \mathbf{x}, \mathbf{D}) = \pi_\theta(\cdot | \mathbf{x}, \mathbf{D}, \mathbf{s}_{\text{solve}})$ ;  
    // Step 2: Atomize response  $\mathbf{y}$  by Proposer and Parse QA pairs by  $\mathcal{Q}(\cdot)$   
4     Sample  $\boldsymbol{\tau} \sim \nu_{\text{propose}}(\cdot | \mathbf{y}) = \pi_\theta(\cdot | \mathbf{y}, \mathbf{s}_{\text{propose}})$ ;  
5     Extract QA pairs  $\{(\mathbf{q}_i, \mathbf{a}_i)\}_{i=1}^n = \mathcal{Q}(\boldsymbol{\tau})$ ;  
    // Step 3: Generate audit trajectory  $\boldsymbol{\lambda}$  from Checker and Parse re-checked  
    answers by  $\mathcal{A}(\cdot)$   
6     Sample  $\boldsymbol{\lambda} \sim \nu_{\text{check}}(\cdot | \{\mathbf{q}_i\}_{i=1}^n, \mathbf{D}) = \pi_\theta(\cdot | \{\mathbf{q}_i\}_{i=1}^n, \mathbf{D}, \mathbf{s}_{\text{check}})$ ;  
7     Parse re-checked answers  $\{\hat{\mathbf{a}}_i\}_{i=1}^n = \mathcal{A}(\boldsymbol{\lambda})$ ;  
    // Step 4: Unified reward labeling  
8     Calculate Zero-Tolerance Reward  $R_i$  by equation 7 (based on the match between  $\mathbf{a}$  and  $\hat{\mathbf{a}}$ );  
    // Step 5: Store both reasoning and auditing trajectories  
9     Add  $\mathbf{y}_i$  and  $\boldsymbol{\lambda}_i$  to the unified batch trajectory set  $\mathcal{T}_{\mathcal{B}}$ ;  
10    Associate both  $\mathbf{y}_i$  and  $\boldsymbol{\lambda}_i$  with terminal reward  $R_i$ ;  
11 **end**  
    // Phase 2: Joint Policy Optimization  
12 **for** each trajectory  $\tau' \in \mathcal{T}_{\mathcal{B}} = \{\mathbf{y}_i, \boldsymbol{\lambda}_i\}$  **do**  
13    Estimate Advantage  $\hat{A}_{\tau'}$  using GAE based on  $R_i$ ;  
14    Compute  $\text{KL}(\pi_\theta || \pi_{\text{ref}})$  for each token in  $\tau'$ ;  
15 **end**  
16 Update  $\pi_\theta$  by computing the joint gradient  $\nabla_\theta \mathcal{L}_{\text{MARCH}}(\theta)$  according to equation 8  
$$\frac{1}{2|\mathcal{B}|} \sum_{\mathbf{x} \in \mathcal{B}} \left[ \sum_{t=1}^{|\mathbf{y}|} \hat{A}_y^t \nabla_\theta \log \pi_\theta(y^t | \cdot) + \sum_{t=1}^{|\boldsymbol{\lambda}|} \hat{A}_\lambda^t \nabla_\theta \log \pi_\theta(\lambda^t | *) - \beta \nabla_\theta \text{KL}[\pi_\theta || \pi_{\text{ref}}] \right];$$
  
17 where the notation  $(\cdot)$  and  $(*)$  in equation 8 encapsulates the respective contexts for brevity:  
 $(\mathbf{x}, \mathbf{D}, \mathbf{y}^{<t})$  and  $(\{\mathbf{q}_i\}, \mathbf{D}, \boldsymbol{\lambda}^{<t})$  for the Solver’s and the Checker’s generation prefixes;

---

## B Experiment

### B.1 Training Details

Training framework is built upon the VerL (Sheng et al., 2025), and we employ Fully Sharded Data Parallel (FSDP) (Zhao et al., 2023) for efficient multi-node multi-GPU training. During the training rollout, we utilize the vLLM (Kwon et al., 2023) for efficient sample generation with the temperature of 0.6. The PPO (Schulman et al., 2017) training is conducted for a single epoch with a global batch size of 32, with a max prompt length of 24,567, max response length of 8,192. We set learning rates for the actor and the critic as  $1 \times 10^{-6}$  and  $1 \times 10^{-5}$ , respectively. We save and evaluate a checkpoint every 20 steps. We report results from the single best-performing checkpoint selected based on the highest average score across all evaluation benchmarks. We give more detailed training hyper-parameters in the Table 6.

Table 6: Key hyperparameters for our RL training, derived from the execution script.

Data & Model Configuration	Max Prompt Length	24,567
	Max Response Length	8,192
	Filter Overlong Prompts	True
Optimizer & Learning Rates	Actor Learning Rate	$1 \times 10^{-6}$
	Critic Learning Rate	$1 \times 10^{-5}$
	Actor LR Warmup Steps	5
	Critic LR Warmup Steps	10
PPO Algorithm & Rewards	$\gamma$	0.998
	KL Divergence Coefficient	$1 \times 10^{-3}$
Training Schedule	Total Epochs	1
	Global Batch Size / Mini-batch Size	32
Rollout (Training)	Inference Engine	vLLM
	Temperature	0.6
Rollout (Evaluation)	Inference Engine	vLLM
	Temperature	0.6
	Batch Size	32/32
	Top-P	0.95
	Top-K	20
	Rollout Number	8

### B.2 Evaluation Details

#### B.2.1 Hallucination and Factual Grounding Benchmarks

**RAGTruth and FaithBench** are public benchmarks for evaluating hallucinations in retrieval-augmented generation, which contain three task types: (i) QA, (ii) Data2Txt Writing, and (iii) Summarization, each with human-annotated hallucinations at the response level. For each example, the model is given a query and a set of retrieved passages, and asked to produce an answer, a data-to-text description, or a summary.

**Contextual JudgeBench** is a judge benchmark with 2,000 challenging response pairs across eight splits inspired by real-world contextual evaluation scenarios, which is built with a multi-pronged data construction pipeline that leverages both existing human annotations and model-based perturbations. The evaluation task requires a model to act as a judge, choosing the better of two provided responses (A and B) according to specific, context-dependent criteria.

**Facts Grounding** is a large-scale factuality benchmark spanning five domains (e.g., medical, financial, legal, and others) and nine task types, including summarization, fact finding, comparison, and definition. Each instance consists of a query, a collection of documents, and a reference output with fine-grained

Table 7: Accuracy (%) on ContextualJudgeBench. Baseline results are cited from Xu et al. (2025a).

Model	Refusal (Ans.)	Refusal (Unans.)	Faithfulness (QA)	Faithfulness (Summ.)	Completeness (QA)	Completeness (Summ.)	Conciseness (QA)	Conciseness (Summ.)	Average
<i>Small Judge Model</i>									
Glider-3.8B (Deshpande et al., 2024)	12.0	8.8	45.6	9.2	20.8	28.7	5.1	4.1	16.8
Prometheus-2-7b (Kim et al., 2024)	12.4	44.0	27.2	32.0	24.0	42.6	6.7	29.5	27.3
Llama-3-OffsetBias-8B (Park et al., 2024)	64.8	11.2	34.0	26.4	33.2	21.1	46.3	23.0	32.6
Skywork-8B (Liu et al., 2024a)	60.8	12.0	38.8	31.6	38.4	26.7	29.4	21.3	32.4
Alta-Selene-8B (Alexandru et al., 2025)	74.4	26.4	40.8	32.8	32.4	34.7	23.1	32.0	37.1
SFRJudge-8B (Wang et al., 2025)	70.8	22.0	40.4	38.8	40.4	43.4	27.5	31.1	39.3
SFRJudge-12B (Wang et al., 2025)	68.4	28.4	45.2	43.6	28.0	51.0	16.1	29.5	38.8
<i>Large Judge Model</i>									
Prometheus-2-8x7b (Kim et al., 2024)	22.0	29.6	22.4	29.6	20.4	39.8	10.2	18.4	24.1
Skywork-70B (Liu et al., 2024a)	82.4	11.2	48.0	47.6	36.8	41.4	21.6	27.9	39.6
ST-Eval-70B (Wang et al., 2024a)	50.0	42.0	51.2	45.6	40.8	39.4	36.1	29.9	41.9
SFRJudge-70B (Wang et al., 2025)	87.6	32.4	60.8	54.8	40.8	53.4	44.7	36.1	51.4
<i>Instruct + Reasoning Model</i>									
Llama-3.1-70B (Touvron et al., 2023)	59.6	48.0	58.0	48.4	38.0	51.8	15.7	27.5	43.4
Llama-3.3-70B (Touvron et al., 2023)	71.6	42.4	68.0	48.4	42.0	51.8	20.8	30.7	47.0
R1-Distill-Llama-3.3-70B (DeepSeek-AI, 2025)	89.6	50.4	74.0	48.4	42.4	57.4	19.2	29.5	51.4
GPT-4o-mini	71.2	22.8	45.6	42.4	33.2	54.2	11.8	29.5	38.8
GPT-4o	64.0	52.0	68.0	50.8	39.6	56.2	12.9	22.5	45.8
DeepSeek-R1 (DeepSeek-AI, 2025)	92.0	52.0	72.0	50.4	41.2	60.6	20.4	26.2	51.9
<i>OURS</i>									
Llama-3.1-8B	28.0	43.2	34.8	34.8	23.2	41.0	11.4	21.3	29.7
MARCH-General	37.2	57.8	52.6	54.8	48.0	63.9	48.1	50.1	51.6
Δ	↑ 9.2	↑ 14.6	↑ 17.8	↑ 20.0	↑ 24.8	↑ 22.9	↑ 36.7	↑ 28.8	↑ 21.9
MARCH-STEM	38.1	64.0	53.6	58.0	44.9	65.5	44.7	49.5	52.3
Δ	↑ 10.1	↑ 20.8	↑ 18.8	↑ 23.2	↑ 21.7	↑ 24.5	↑ 33.3	↑ 28.2	↑ 22.6

factual annotations, providing a broad and challenging testbed for document-grounded evaluation. During test, we adopt the official JSON judgment prompt provided by Jacovi et al. (2025); Li et al. (2025b).

## B.2.2 Comprehensive Results on ContextualJudgeBench

Table 7 gives the comprehensive results on ContextualJudgeBench.

### B.2.3 QA Benchmarks

**HotpotQA.** HotpotQA (Yang et al., 2018) is a large-scale multi-hop QA dataset that requires reasoning over multiple Wikipedia passages to answer fact-based questions. Each question is paired with multiple supporting paragraphs and a short, factoid-style answer.

**2WikiMultiHopQA.** 2WikiMultiHopQA (Ho et al., 2020) extends HotpotQA with questions that require reasoning across two different Wikipedia articles, emphasizing cross-document retrieval and multi-hop reasoning over disjoint evidence.

**MuSiQue.** MuSiQue (Trivedi et al., 2022) is a challenging multi-hop QA dataset with questions that explicitly require reasoning across multiple sentences scattered across different documents, often involving compositional reasoning over several intermediate facts.

## B.3 Baseline Methods

**Baselines.** The raw backbone model used in all settings, evaluated in a zero-shot manner with task-specific instructions but without additional reasoning or RL fine-tuning.

**+CoT.** A prompting baseline where we encourage step-by-step reasoning by prepending chain-of-thought instructions (e.g., “let’s think step by step”) to the prompt. The model is still the same base LLM; only the decoding prompt is changed.

**+Few-Shot.** An in-context learning baseline where we augment the prompt with a small number of annotated QA examples (few-shot demonstrations) drawn from the training distribution, following standard practice in multi-hop QA and factuality evaluation.

**+RLHF.** A version of the base LLM further optimized with reinforcement learning from human feedback, using general-purpose reward models that favor helpful and harmless responses but are not

Table 8: Main results on Multi-Hop Question Answering benchmarks across two training datasets based on Llama3.1-8B-Instruct. We report the accuracy rate (%) by comparing the model output with the corresponding ground-truth label. We produce baseline results on our own environment.

Training Dataset			STEM				General			
Method	Solver	Checker	MuSiQue	2Wiki	Hotpot	Average	MuSiQue	2Wiki	Hotpot	Average
<i>Compare with Instruct Baseline</i>										
Baseline			5.60	17.40	35.00	19.33	5.60	17.40	35.00	19.33
+MARCH	✓	-	12.80 <sup>↑ 7.20</sup>	25.40 <sup>↑ 8.00</sup>	49.40 <sup>↑ 14.40</sup>	29.20 <sup>↑ 9.87</sup>	12.20 <sup>↑ 6.60</sup>	27.80 <sup>↑ 10.40</sup>	54.80 <sup>↑ 19.80</sup>	31.60 <sup>↑ 12.27</sup>
	✓	✓	16.60 <sup>↑ 3.80</sup>	35.40 <sup>↑ 10.00</sup>	55.20 <sup>↑ 5.80</sup>	35.73 <sup>↑ 6.53</sup>	18.40 <sup>↑ 6.20</sup>	38.60 <sup>↑ 10.80</sup>	56.40 <sup>↑ 1.60</sup>	37.80 <sup>↑ 6.20</sup>
<i>Compare with RLHF Baseline</i>										
RLHF			24.20	49.60	60.40	44.73	24.20	56.40	65.40	48.67
+MARCH	✓	-	29.00 <sup>↑ 4.80</sup>	49.60 <sup>↑ 0.00</sup>	66.40 <sup>↑ 6.00</sup>	48.33 <sup>↑ 3.60</sup>	30.80 <sup>↑ 6.60</sup>	56.40 <sup>↑ 0.00</sup>	67.80 <sup>↑ 2.40</sup>	51.67 <sup>↑ 3.00</sup>
	✓	✓	29.40 <sup>↑ 0.40</sup>	53.20 <sup>↑ 3.60</sup>	66.20 <sup>↓ 0.20</sup>	49.60 <sup>↑ 1.27</sup>	30.80 <sup>↑ 0.00</sup>	61.00 <sup>↑ 4.60</sup>	69.00 <sup>↑ 1.20</sup>	53.60 <sup>↑ 1.93</sup>
<i>Boosting with Few Shots (10-Shots)</i>										
RLHF + Few Shots			29.20	55.20	61.20	48.53	29.40	55.20	60.60	48.43
+MARCH	✓	-	28.60 <sup>↓ 0.60</sup>	58.60 <sup>↑ 3.40</sup>	61.80 <sup>↑ 0.60</sup>	49.67 <sup>↑ 1.14</sup>	28.00 <sup>↓ 1.40</sup>	57.20 <sup>↑ 2.20</sup>	62.20 <sup>↑ 1.60</sup>	49.13 <sup>↑ 0.70</sup>
	✓	✓	40.80 <sup>↑ 12.80</sup>	69.40 <sup>↑ 10.80</sup>	73.60 <sup>↑ 11.80</sup>	61.27 <sup>↑ 10.60</sup>	37.20 <sup>↑ 9.20</sup>	62.00 <sup>↑ 4.80</sup>	64.00 <sup>↑ 1.80</sup>	54.40 <sup>↑ 5.27</sup>
<i>Boosting with Chain-of-Thought (CoT)</i>										
RLHF + CoT			29.40	59.00	64.40	50.93	31.60	58.60	62.80	51.00
+MARCH	✓	-	32.40 <sup>↑ 3.00</sup>	60.60 <sup>↑ 1.60</sup>	64.00 <sup>↓ 0.40</sup>	52.33 <sup>↑ 1.40</sup>	32.40 <sup>↑ 0.80</sup>	62.00 <sup>↑ 3.40</sup>	68.20 <sup>↑ 5.40</sup>	54.20 <sup>↑ 3.20</sup>
	✓	✓	36.20 <sup>↑ 3.80</sup>	70.60 <sup>↑ 10.00</sup>	70.60 <sup>↑ 6.60</sup>	59.13 <sup>↑ 6.80</sup>	38.40 <sup>↑ 6.00</sup>	63.80 <sup>↑ 1.80</sup>	71.20 <sup>↑ 3.00</sup>	57.80 <sup>↑ 3.60</sup>

specifically tuned for document-grounded factuality. We adopt Skywork-Reward-V2-Llama-3.1-8B (Liu et al., 2025a) as the reward model.

#### B.4 Analysis of Internal Performance Gains

Table 8 demonstrates that joint optimization of the *Solver* and the *Checker* outperforms both baseline configurations and the *Solver*-only updates across most benchmarks. A key observation is the stability provided by the *Checker* role in complex reasoning scenarios. Specifically, while a *Solver*-only update can lead to performance regressions under 10-shot prompting (e.g., dropping from 29.20 to 28.60 on MuSiQue STEM), the joint update effectively reverses this trend, reaching 40.80, highlighting that the *Checker* acts as a critical anchor by preventing factual drift when the model is tasked with extended prompts. Notably, the benefits of MARCH are most pronounced when integrated with advanced prompting strategies like Few-shot and CoT. The average STEM score improvement (Joint vs. the *Solver*-only) increases from 1.27 points in the RLHF setting to 10.60 points in the Few-shot setting and 6.80 points in the CoT setting. In summary, Table 8 suggests that by internalizing a claim-level auditing logic, MARCH significantly amplifies the efficacy of in-context learning and multi-step reasoning, ensuring that increased model capacity translates into verifiable factual accuracy rather than hallucinated complexity.

#### B.5 Extended Ablation Analysis

In this section, we provide a comprehensive analysis of the reward mechanisms and model-specific performance gains reported in Table 5 and Table 9.

**Influence of Reward Formulations.** We investigate two primary reward designs:

- **Error Rate (ER):** Defined as  $R_{ER} = -\frac{N_{err}}{N_{total}}$ , where  $N_{total}$  is the total number of claims generated and  $N_{err}$  is the number of claims identified as incorrect by the Checker. This provides a granular, proportional feedback signal.
- **Zero-Tolerance Reward (ZTR):** Defined as  $R_{ZTR} = -1$  if  $N_{err} > 0$ , and 0 otherwise. This enforces a rigorous binary constraint, treating any single factual error as a failure of the entire response.

Our results indicate that ZTR consistently outperforms ER across both *Solver*-only and Joint training configurations. For instance, in the joint training setting, ZTR improves the average accuracy from 55.46% (ER) to 61.25% (ZTR). This confirms that for RAG tasks requiring high factual precision, a

Table 9: Ablation study on reward functions, scalars and base LLM across multiple RAG and reasoning benchmarks. We compare Error Rate Reward (ERR) and Zero-Tolerance Reward (ZTR) under different training configurations.

Reward Function	Reward Scalar	Solver	Checker	RAGTruth			Facts Grounding	MuSiQue	2Wiki	Hotpot	Average
				Summary	Data2Txt	QA					
Llama3.1-8B Baseline				71.33	28.67	63.31	57.09	5.60	17.40	35.00	39.77
<i>Ablation Study on Reward Function, General Training Dataset</i>											
ERR	-	✓	-	80.00	49.33	74.82	71.98	10.60	24.80	52.00	51.93
ZTR	-1/0	✓	-	88.67	66.67	82.73	80.12	12.20	27.80	54.80	59.00
ERR	-	✓	✓	87.33	62.67	76.26	76.98	9.80	25.60	49.60	55.46
ZTR	-1/0	✓	✓	89.33	66.67	82.73	76.63	18.40	38.60	56.40	61.25
<i>Ablation Study on Reward Scalar, STEM Training Dataset</i>											
ZTR	0/1	✓	-	78.00	42.00	75.54	68.84	7.4	25.00	48.85	49.38
	-1/0	✓	-	80.58	51.33	80.50	77.09	12.80	25.40	49.40	53.87
	0/1	✓	✓	74.67	54.67	79.14	69.65	6.8	25.20	42.80	50.42
	-1/0	✓	✓	83.33	58.67	82.73	81.51	16.60	35.40	55.20	59.06
<i>Ablation Study on Different Base LLM</i>											
Qwen3-8B Baseline				54.00	28.50	61.15	52.41	48.00	72.40	81.50	56.84
General	ZTR	✓	✓	72.00	47.67	75.90	64.45	53.20	76.90	85.20	67.90
STEM	ZTR	✓	✓	74.67	48.00	74.10	64.19	53.60	76.20	86.00	68.11

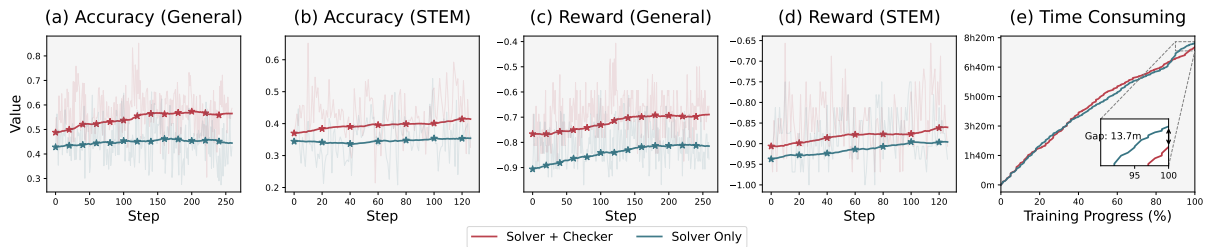


Figure 5: Visualization of training efficiency and dynamics. (a)-(d) illustrate the convergence of accuracy and rewards across general and STEM datasets. (e) Comparison of cumulative training time on the General dataset.

strict "all-or-nothing" penalty is more effective at internalizing logical consistency than a proportional reduction.

**Impact of Reward Scalars** We further compare two scalar assignments within the ZTR framework: a penalty-based setting ( $-1/0$ ) and an incentive-based setting ( $0/1$ ). Experimental results on the STEM dataset show that the  $-1/0$  configuration (59.06%) significantly outperforms the  $0/1$  setting (50.42%). We hypothesize that because factual successes are relatively sparse during early training phases, the  $0/1$  incentive fails to provide sufficient contrast between diverse erroneous paths. In contrast, the  $-1/0$  setting establishes factual accuracy as the default expectation, providing a robust corrective gradient that effectively suppresses hallucinations.

**Generalizability Across Model Families** To verify the model-agnostic nature of MARCH, we extend our evaluation to the Qwen3-8B family. The performance trajectory of Qwen3 mirrors that of Llama3.1, showing consistent gains across all benchmarks. Specifically, MARCH elevates the Qwen3-8B baseline from 56.84% to a maximum of 68.11% average accuracy when trained on STEM data. The specific task-level breakdown (e.g., 74.67% on RAGTruth Summary and 86.00% on HotpotQA) demonstrates that the dual-trajectory optimization consistently enhances both reasoning synthesis and evidentiary precision regardless of the underlying model architecture.

## B.6 Does our method conflict with other boosting methods?

As demonstrated in Figure 3, MARCH is highly orthogonal and can be seamlessly integrated with existing performance-boosting techniques, including RLHF, Few-Shot prompting, and CoT reasoning. When applied to the Llama3.1-8B-Instruct baseline, our method consistently achieves significant performance gains across all benchmarks. For instance, combining MARCH with CoT reasoning further elevates

the average accuracy on STEM tasks from 50.00% to 59.13%, and on General tasks from 50.00% to 57.80%. Similarly, notable improvements are observed when MARCH is integrated with Few-Shot (10-Shots) configurations, reaching an average of 61.27% on STEM. Table 8 indicates that MARCH does not conflict with traditional scaling or prompting methods; rather, it provides a complementary layer of factual rigor through its asymmetric self-play mechanism, leading to a cumulative enhancement in multi-hop question-answering capabilities. We provide the implementation details in Appendix B.3 and more results in Table 8.

### **B.7 Visualization of Training Dynamics and Time Consuming**

To further investigate the practical efficiency of the MARCH framework, we analyze the training dynamics and computational overhead. Figures 5(a)-(d) illustrate the stable convergence of accuracy and rewards across both General and STEM datasets, where the joint Solver + Checker update consistently outperforms the Solver Only configuration. Regarding training efficiency, Figure 5(e) presents the cumulative time consumption on the General dataset (containing 9,000 samples), where we find that simultaneously updating both the *Solver* and the *Checker* trajectories in a single training step does not impose a prohibitive computational burden.

## C Case Study

### C.1 Case 1

#### Input $(x, D)$ of Agent Solver $\nu_{\text{solve}}$

You must respond based strictly on the information in provided passages. Do not incorporate any external knowledge or infer any details beyond what is given in the passages.

##### **Passages:**

##### **Document-1**

TUDY OBJECTIVE: To describe sleep characteristics and rapid eye movement (REM) sleep behavior disorder in patients with Guadeloupean atypical parkinsonism (Gd-PSP)...(omit)... Subjects with idiopathic REM sleep behavior disorder should avoid eating sour sop.

##### **Document-2**

We characterize the clinical features of Parkinson's syndrome on Guadeloupe and describe possible environmental causes...(omit)... The development of atypical parkinsonism in Guadeloupe and probably elsewhere, could result from synergistic toxicity, but acetogenins are probably the most potent neurotoxin, acting as mitochondrial complex I inhibitor.

##### **Document-3**

In Guadeloupe, there is an abnormally high frequency of atypical parkinsonism...(omit)... In conclusion, atypical Guadeloupean parkinsonism comprises two forms of parkinsonism and dementia that differ clinically by the presence of oculomotor signs, but have similar cognitive profiles and neuroimaging features, suggesting that they may constitute a single disease entity, and both were similarly exposed to annonaceous neurotoxins, notably annonacin.

##### **Document-4**

An unusually high frequency of atypical Parkinson syndrome has been delineated over the last 5 years in the French West Indies...(omit)... Guadeloupean parkinsonism may prove to be a tauopathy identical or closely related to PSP.

##### **Document-5**

INTRODUCTION: On Guadeloupe, atypical parkinsonism is abnormally frequent, and represents 75% of progressive parkinsonism while Parkinson's disease (PD) accounts for only 25%, which is an inverted percentage in comparison with Europe...(omit)... This hypothesis is in keeping with epidemiologic data and experimental studies showing neuronal loss after exposure to isoquinolines or acetogenins.

##### **Document-6**

Steele, Richardson and Olszewski in 1964 described a distinctive clinical and pathological entity they called progressive supranuclear palsy (PSP)...(omit)... Future studies will confirm if Guadeloupean parkinsonism is truly a geographic focus of PSP, and if dietary factors account for both.

##### **Document-7**

On the French West Indian island of Guadeloupe...(omit)... Consistent with this view, chronic administration of annonacin to rats through Alzet osmotic minipumps showed that annonacin was able to reproduce the brain lesions characteristic of the human disease.

##### **Document-8**

Atypical parkinsonism is extremely frequent in Guadeloupe and may have an environmental cause...(omit).. This electrophysiological study, added to previous clinical, neuropsychological and neuroradiological studies, has enriched the characterization of Guadeloupean atypical parkinsonism, which thus appears to be a new clinical entity.

**Question:**

Describe clinical presentation of Parkinsonism with dementia of Guadeloupe syndrome.

**Generated Response  $y$  of Agent Solver  $\nu_{\text{solve}}$  by  $y \sim \nu_{\text{solve}}(\cdot|x, D)$** 

The clinical presentation of Parkinsonism with dementia in Guadeloupe syndrome includes the following features: 1. **Levodopa Resistance**: Patients exhibit levodopa-resistant parkinsonism, distinguishing them from typical Parkinson’s disease. 2. **Fronto-Subcortical Dementia**: Cognitive deficits are primarily fronto-subcortical, characterized by frontolimbic cognitive impairment, hallucinations (52–59%), and dementia with features such as dysautonomia (50%) and tremor (>50%). 3. **Oculomotor Abnormalities**: Some patients display supranuclear oculomotor dysfunction, including reduced vertical saccade velocity and impaired antisaccades. However, oculomotor dysfunction is not universal, as seen in the parkinsonism-dementia complex (PDC) subgroup. 4. **Tremor and Myoclonus**: Jerky tremor and cortical myoclonus (89% of cases) are prominent, with myoclonus originating from cortical regions. 5. **REM Sleep Behavior Disorder (RBD)**: 78% of patients with Guadeloupean atypical parkinsonism (Gd-PSP) experience RBD, often preceding parkinsonism symptoms by years. 6. **Neuroimaging Findings**: Cerebral atrophy is common, with enlargement of the third ventricle and T2-hypointensity in the basal ganglia (particularly substantia nigra). Neuronal loss affects the substantia nigra, striatum, and frontal cortex, alongside tau-containing lesions. 7. **Environmental Factor**: Higher consumption of soursop (*Annona muricata*), a plant containing the mitochondrial complex I inhibitor annonacin, is significantly associated with the syndrome. features collectively define Guadeloupe syndrome as a distinct clinical entity characterized by parkinsonism, dementia, and unique neuroimaging and environmental correlates.

**Atomize  $y$  into verifiable claims through Agent Proposer  $\nu_{\text{propose}}$** 

**You are a data annotation expert.** Your task is to analyze a model’s response to a user query. For every critical number found in the model’s response, you must create a question where the answer is precisely and exactly that number.

**Rule 1. Clarity and Specificity of Questions:** Each question must be clear, complete, and unambiguous. It must contain enough context from the response to ensure the number is the only possible correct answer. For example, if the model’s response contains “In 2025, 50 people will take the bar exam in Beijing,” an effective question would be “How many people will take the bar exam in the Beijing area in 2025?”, and the answer is 50. A bad example would be “How many people will take the bar exam in 2025?”. This is ambiguous because it lacks the specific region (‘Beijing area’).

**Rule 2. Format of the Answer:** The answer must be a pure number, either an integer or a decimal. It must NOT include any other characters or formats, such as percentage signs (%), ‘k’ to denote thousands (e.g., 10k), ranges (e.g., 10-20), or words (e.g., fifty).

Please provide your output as an unordered list using the following format for each question-answer pair:

- Question: xxx [Answer: n]
- Question: xxx [Answer: n]

Now, generate questions for the user query and model response below. Do not output any other content or explanations.

**Model Response:**

The clinical presentation of Parkinsonism with dementia in Guadeloupe syndrome includes the following features: 1. **Levodopa Resistance**: Patients exhibit levodopa-resistant parkinsonism, distinguishing them from typical Parkinson’s disease. 2. **Fronto-Subcortical Dementia**: Cognitive deficits are primarily fronto-subcortical, characterized by frontolimbic cognitive

impairment, hallucinations (52–59%), and dementia with features such as dysautonomia (50%) and tremor (>50%). 3. **Oculomotor Abnormalities**: Some patients display supranuclear oculomotor dysfunction, including reduced vertical saccade velocity and impaired antisaccades. However, oculomotor dysfunction is not universal, as seen in the parkinsonism-dementia complex (PDC) subgroup. 4. **Tremor and Myoclonus**: Jerky tremor and cortical myoclonus (89% of cases) are prominent, with myoclonus originating from cortical regions. 5. **REM Sleep Behavior Disorder (RBD)**: 78% of patients with Guadeloupean atypical parkinsonism (Gd-PSP) experience RBD, often preceding parkinsonism symptoms by years. 6. **Neuroimaging Findings**: Cerebral atrophy is common, with enlargement of the third ventricle and T2-hypointensity in the basal ganglia (particularly substantia nigra). Neuronal loss affects the substantia nigra, striatum, and frontal cortex, alongside tau-containing lesions. 7. **Environmental Factor**: Higher consumption of soursop (*Annona muricata*), a plant containing the mitochondrial complex I inhibitor annonacin, is significantly associated with the syndrome. features collectively define Guadeloupe syndrome as a distinct clinical entity characterized by parkinsonism, dementia, and unique neuroimaging and environmental correlates.

### Generated Question-Answers pairs of the *Proposer*.

- Question: What is the lower percentage range of patients with Guadeloupe syndrome experiencing hallucinations? [Answer: 52]
- Question: What is the upper percentage range of patients with Guadeloupe syndrome experiencing hallucinations? [Answer: 59]
- Question: What percentage of patients with Guadeloupe syndrome have dysautonomia? [Answer: 50]
- Question: What percentage of patients with Guadeloupe syndrome have cortical myoclonus? [Answer: 89]
- Question: What percentage of patients with Guadeloupe syndrome experience REM Sleep Behavior Disorder (RBD)? [Answer: 78]

### Input ( $\{q_i\}_{i=1}^n, D$ ) of Agent Checker $\nu_{\text{check}}$

**You are a data validation expert.** You will be given a set of reference materials and a list of user questions. Your task is to retrieve information from the reference materials to answer all questions. Before providing each answer, you must state the evidence for it. You must adhere strictly to the content of the reference materials and must not fabricate or infer any information that is not explicitly mentioned. If the relevant information cannot be found in the materials, you must state “Cannot answer”. It is important to note that all answers are pure numbers. Therefore, your answer must also be a pure number and must not contain any other content, such as percentage signs (%), 'k' to denote thousands, ranges, or numbers written as words.

*Here is an example:*

#### **Reference Materials:**

Document 1: ...in 2024, 50 people will take the bar exam in Beijing.

Document 2: ...(omitted)

- Question 1: How many people will take the bar exam in Beijing in 2024?
- Question 2: How many people will take the bar exam in Beijing in 2025?

#### **Example Response:**

1. Evidence: Document 1 states that 50 people will take the bar exam in Beijing in 2024.  
[Answer: 50]
  2. Evidence: The materials do not contain information about the bar exam in Beijing for 2025, therefore I cannot answer.  
[Answer: Cannot answer]
- Now, based on the reference materials below, answer all the following questions. Strictly adhere to the example format provided.

**Reference Materials:**

**Document-1** (omit)

**Document-2** (omit)

**Document-3** (omit)

**Document-4** (omit)

**Document-5** (omit)

**Document-6** (omit)

**Document-7** (omit)

**Document-8** (omit)

**Questions:**

- Question: What is the lower percentage range of patients with Guadeloupe syndrome experiencing hallucinations?
- Question: What is the upper percentage range of patients with Guadeloupe syndrome experiencing hallucinations?
- Question: What percentage of patients with Guadeloupe syndrome have dysautonomia?
- Question: What percentage of patients with Guadeloupe syndrome have cortical myoclonus?
- Question: What percentage of patients with Guadeloupe syndrome experience REM Sleep Behavior Disorder (RBD)?

**Generated Response  $\{\hat{a}_i\}_{i=1}^n$  of Agent Checker  $\nu_{\text{check}}$  by  $\{\hat{a}_i\}_{i=1}^n \sim \nu_{\text{check}}(\cdot | \{\mathbf{q}_i\}_{i=1}^n, \mathbf{D})$**

1. Evidence: Document-5 states that 52% of PDC patients had hallucinations, which is the lower percentage range mentioned for Guadeloupe syndrome.  
Answer: 52
2. Evidence: Document-3 states that 59% of patients with PSP-like syndrome experienced hallucinations, which is the upper percentage range mentioned for Guadeloupe syndrome.  
Answer: 59
3. Evidence: Document-3 states that dysautonomia was present in 50% of patients with Guadeloupe syndrome.  
Answer: 50

4. Evidence: Document-8 states that 89% of Gd-PSP patients had cortical myoclonus.  
Answer: 89

5. Evidence: Document-1 states that 78% of patients with Gd-PSP experienced REM sleep behavior disorder.  
Answer: 78

**Calculate Zero-Tolerance Reward (ZTR) according to  $\{a_i\}_{i=1}^n$  contained in  $y$  from Agent Solver  $\nu_{\text{solve}}$  and  $\{\hat{a}_i\}_{i=1}^n$  generated by Agent Checker  $\nu_{\text{check}}$ .**