

Towards Reliable Long-Context Reasoning: Detecting Unanswerable Questions via FactGuard

Qian-Wen Zhang^{*†}, Biao Liu^{*}, Fang Li^{*}, Jie Wang, Lingfeng Qiao,
Yifei Yu, Di Yin, Xing Sun
Tencent Youtu Lab, China
cowenzhang@tencent.com

Abstract

Despite the success of Large Language Models (LLMs) in long-context tasks, their reliability in recognizing unanswerable questions remains a critical challenge. In information-seeking scenarios, models often suffer from overconfidence and hallucinations when evidence is missing or misleading. In this paper, we propose FACTGUARD, a systematic multi-task workflow for generating evidence-grounded answerable and unanswerable questions. We introduce FACTGUARD-BENCH, a benchmark comprising 25,220 examples with context lengths up to 128K, specifically targeting Lack of Evidence and Misleading Evidence scenarios. Evaluations of eleven representative open-source and proprietary LLMs reveal a significant reliability gap: models struggle to reject unanswerable questions as context scales, with the leading model achieving only 71.62% accuracy. By training with FACTGUARD-BENCH, we improve overall accuracy to 81.17% and significantly enhance models' reasoned rejection capabilities. Our work provides a scalable path toward building more faithful and trustworthy long-context LLMs. Our source is publicly available at <https://github.com/biaoliu-kiritsugu/FactGuard>.

1 Introduction

As Large Language Models (LLMs) evolve, their capacity to process long-form documents has expanded, with context windows now routinely supporting up to 128K tokens or more (Comanici et al., 2025; Yang et al., 2025; Liu et al., 2024). While this expansion enables complex downstream applications like multi-document reading comprehension and long-chain reasoning, it introduces a critical, yet under-explored challenge to reliability: the model's capacity to perceive the boundaries of the knowledge contained within its context window. A

model's utility is defined not only by its retrieval accuracy but also by its capacity for reasoned rejection—recognizing when a question is unanswerable based solely on the provided context.

However, modern LLMs frequently suffer from a **reliability gap**. When faced with questions that lack sufficient evidence or contain misleading premises, models often exhibit overconfident behavior, generating plausible but factually untraceable hallucinations (Slobodkin et al., 2023; Deng et al., 2024). This risk is exacerbated in long-context scenarios, where the expansive search space dilutes attentional focus, increasing the probability of models forging spurious connections between segments to force an answer. Achieving a balance between **helpfulness** (answering valid questions) and **faithfulness** (refusing unanswerable ones) is essential for the safe deployment in high-stakes domains like legal or medical analysis.

Existing research into unanswerable questions often relies on limited scenarios or simplified response mechanisms. For instance, SelfAware (Yin et al., 2023) employs a straightforward prompting approach to detect unanswerability, but relies on predefined, static replies like "The answer is unknown", which lacks the explanatory depth required for reasoning. Similarly, KUQ (Amayuelas et al., 2024) investigates "known-unknown" questions in open-ended QA, yet its focus remains on internal model knowledge rather than the rigorous, evidence-based reading comprehension necessitated by long-context inputs. Furthermore, widely used benchmarks such as SQuAD 2.0 (Rajpurkar et al., 2018) are largely constrained by short context lengths (under 4K tokens) and fail to evaluate model behavior across extended text spans. Consequently, there remains a critical need for a systematic framework that addresses Lack of Evidence and Misleading Evidence while providing reasoned rejections in long-context environments.

To address these limitations, we propose FACT-

^{*} Equal contribution.

[†] Corresponding author.




Figure 1: Illustration of FACTGUARD for data synthesis in a collaborative multi-task workflow framework.

GUARD, a systematic, collaborative multi-task workflow designed to automatically synthesize high-quality, long-context datasets for reliability evaluation and enhancement. Our contributions are summarized as follows:

❶ **Systematic Framework for Reliability Enhancement:** We introduce the FACTGUARD workflow, which coordinates multiple LLM agents. Specifically, we design strategies for constructing unanswerable questions through automated evidence removal, entity substitution, and impossible condition insertion, providing a stress test for the faithfulness-helpfulness trade-off.

❷ **Large-Scale Long-Context Benchmark:** We curate FACTGUARD-BENCH, comprising 25,220 examples with context lengths ranging from 4K to 128K. To the best of our knowledge, this represents one of the most extensive benchmarks engineered to quantify and bolster LLM rejection sensitivity in long-form reading comprehension.

❸ **Empirical Insights and Cross-Domain Robustness:** Our evaluation of eleven LLMs reveals a persistent **reliability gap** that exacerbates as context scales. We demonstrate that FACTGUARD-BENCH effectively bridges this gap via **reasoned rejection** without inducing **catastrophic forgetting**, achieving significant reliability gains on out-of-distribution benchmarks such as SQuAD 2.0.

2 Methodology: The FACTGUARD Framework

FACTGUARD is a systematic, multi-agent collaborative workflow designed to synthesize high-quality long-context data focus on reliability. As illustrated in Figure 1, the framework operates in four sequential stages: (1) Contextual Preparation, (2)

Grounded QA Generation, (3) Adversarial Unanswerable Synthesis, and (4) Quality Assurance.

2.1 Contextual Preparation

To handle the computational and semantic complexity of long-form documents, we first decompose the original text $\mathcal{D}$ into a set of coherent fragments $\mathbf{X} = \{x_1, x_2, \dots, x_n\}$. We construct these fragments using paragraph-aware sliding windows, constraining each fragment length to $|x_i| \in [500, 1000]$ tokens to preserve semantic integrity. Each fragment x_i undergoes a two-fold process:

- **Axiomatic Quality Scoring:** For each fragment $x_i \in \mathbf{X}$, a judge LLM assigns a quality score $s_i \in \{1, 2, 3, 4, 5\}$ based on linguistic fluency, logical coherence, and informational density (Allen and Seaman, 2007). Fragments with $s_i < 4$ are pruned to ensure corpus quality.
- **Structured Information Extraction:** For each retained fragment x_i , we apply an extraction LLM to obtain a set of factual primitives, including temporal expressions, entities, events, and topic labels. These primitives serve as verifiable anchors for subsequent question-answer generation.

2.2 Evidence-Grounded Answerable QA Generation

Given a high-quality fragment x_i , the objective is to generate a grounded triplet (q_i, a_i, e_i) , where q_i is the question, a_i is the answer, and e_i is an explicit evidence span extracted from x_i . Unlike standard QA generation, we enforce an *Evidence-First* constraint: the model must first identify a viable evidence span e_i before formulating q_i and a_i . This prevents the model from hallucinating information not present in x_i . A secondary validation agent then cross-checks whether a_i can be

logically derived from e_i without external knowledge, ensuring strict context-faithfulness.

2.3 Adversarial Unanswerable Synthesis

The core of FACTGUARD is the systematic construction of unanswerable instances $(q_i^{\text{un}}, r_i^{\text{un}})$, where q_i^{un} is an unanswerable question and r_i^{un} is a reasoned rejection. We categorize these instances into two distinct failure modes targeting model reliability.

Lack of Evidence In this scenario, we perform Evidence Ablation. Given a grounded triplet (q_i, a_i, e_i) from Stage 2, we remove the evidence span e_i from the fragment x_i , resulting in an evidence-removed context $x'_i = x_i \setminus e_i$. The original question q_i then becomes an unanswerable question q_i^{un} . The model is trained to generate a rejection response r_i^{un} that explicitly identifies the missing factual link and summarizes the remaining relevant context, thereby justifying why the question cannot be answered from the provided context.

Misleading Evidence To test model robustness against false premises, we apply two adversarial transformations to the original question q_i :

- **Entity Substitution:** We replace key entities in q_i with semantically similar but factually incorrect ones (e.g., swapping "Global and Local Mixture Consistency" with "Global and Local Augmentation Consistency").
- **Impossible Condition Insertion:** We augment q_i with constraints or conditions that are absent from the document (e.g., adding "in the year 2025" to a historical text).

For misleading questions, the target response r_i^{un} must follow a *Clarification-then-Correction* template: first, clarify that the specific condition or entity is missing, then provide the correct information related to the original, unmodified context. This ensures the model remains helpful while being strictly honest. As formalized in Table 1, the adversarial questions in FACTGUARD-BENCH are designed to simulate real-world reliability failures. In *Lack of Evidence* cases, the model must identify that while the topic is present, the specific factual sentences have been removed from x_i . In *Misleading Evidence* cases, the model must demonstrate semantic vigilance, recognizing that entity substitutions (e.g., swapping GLMC for GLACL) or impossible constraints (e.g., adding specific dataset names like CIFAR-100-LT), even if the surrounding context remains highly relevant.

Lack of Evidence. The question is related to the article, but the factual basis is deleted.

Fragment: ... There had been a lack of confidence in Murray since Romani, and the two failed Gaza battles increased his unpopularity among both the infantry and the mounted troops. ~~After the war Allenby acknowledged Murray's achievements in a June 1919 despatch in which he summed up his campaigns. ...~~

Question: According to this article, in what year did Allenby recognize Murray's accomplishments in his circular?

Answer: The question cannot be answered. The article mentions Murray's performance in the battle, but does not mention what year Allenby recognized his accomplishments.

Misleading Evidence. The key information of the question is misaligned against the facts of the article.

Fragment: Global and Local Mixture Consistency Cumulative Learning (GLMC) for Long-Tailed Visual Recognition... The paper introduces GLMC, a one-stage training strategy designed to improve long-tailed visual recognition by enhancing the robustness of the feature extractor and reducing the bias of the classifier towards head classes. GLMC uses a global and local mixture consistency loss ...

Raw Question: What are the core ideas behind the Global and Local Mixture Consistency cumulative learning (GLMC) framework and how does it improve long-tailed visual recognition?

Question₁ with Entity Substitution: What are the core ideas behind the Global and Local Augmentation Consistency Learning (GLACL) framework and how does it improve long-tailed visual recognition?

Answer₁: The article focuses on GLMC and does not mention GLACL. The core ideas of GLACL cannot be answered, but about GLMC. ...

Question₂ with Impossible Condition Insertion: What are the core ideas behind the Global and Local Mixture Consistency cumulative learning (GLMC) framework and how does it improve long-tailed visual recognition on CIFAR-100-LT?

Answer₂: The article does not mention CIFAR-100-LT. The question of how GLMC improves long-tailed visual recognition on CIFAR-100-LT cannot be answered, but the article mentioned GLMC improve long-tailed visual recognition by enhancing. ...

Table 1: A detailed categorization of unanswerable examples in FACTGUARD-BENCH.

2.4 Quality Assurance

To bridge the gap between automated generation and high-quality benchmarks, FACTGUARD incorporates a rigorous verification protocol:

RAG-based Consistency Check: We formulate the consistency check as a short-context entailment task. For a given unanswerable question q_i^{un} and its associated document $\mathcal{D}$, the agent retrieves the top- k semantically nearest fragments $\mathcal{R}_i = \text{TopK}_{x \in \mathcal{X}} \text{sim}(q_i^{\text{un}}, x)$, where $\text{sim}(\cdot, \cdot)$ is a semantic similarity function. In *Lack of Evidence* cases, $\mathcal{D}$ contains the evidence-removed fragment $x'_i = x_i \setminus e_i$ rather than the original fragment x_i . In *Misleading Evidence* cases, the adversarial premise in q_i^{un} is not entailed by $\mathcal{D}$. The agent must confirm that $\mathcal{R}_i$ provides no sufficient grounds for answering q_i^{un} . Samples where the validation agent's evidence-sufficiency judgment conflicts with the assigned ground-truth are discarded to maintain benchmark fidelity.

Knowledge Decoupling: We use the World Wide Web to identify and remove common-sense questions. If q_i is resolvable via external knowledge alone, it is discarded to ensure FACTGUARD-BENCH strictly evaluates context faithfulness rather than parametric memory.

SPLIT	OVERALL	ENGLISH (EN)				CHINESE (ZH)			
		0-16K	16-32K	32-64K	64-128K	0-16K	16-32K	32-64K	64-128K
TRAIN	19100	2043	2508	3077	3071	4065	3141	826	369
DEV	1920	300	300	270	270	300	120	300	60
TEST	4200	600	600	600	600	600	300	300	600

Table 2: Detailed statistics of FACTGUARD-BENCH.

Dataset	Large-scale (#Num)	Long Context	GT w/ Reason	Low Human Cost	Language
SQuAD 2.0 (Rajpurkar et al., 2018)	✓ (151k)	✗	✗	✗	en
QnotA (Agarwal et al., 2023)	✗ (320)	✗	✗	✗	en
KUQP (Deng et al., 2024)	✗ (320)	✗	✗	✗	en
FACTGUARD-BENCH	✓ (25k)	✓	✓	✓	en&zh

Table 3: Distinguishing features of unanswerable question datasets.




Figure 2: Distributions of FACTGUARD-BENCH in terms of domain, question type and length.

3 Benchmark Constructions

FACTGUARD dynamically generates answerable and unanswerable questions by leveraging a multi-task collaboration process. The LLM underlying the whole process is Qwen2.5-72B-Instruct (Yang et al., 2024b). We collect raw, lengthy texts from the open-source community as the initial input for our process. These texts cover both Chinese and English languages and span domains such as law and books. Specifically, the datasets include legal datasets such as Pile of Law (Henderson et al., 2022), Tiger Law (Chen et al., 2023), the book dataset Gutenberg (Project Gutenberg, 1971), and open-copyright Chinese books.

3.1 Characteristics

We develop a large-scale dataset of long context, FACTGUARD-BENCH, using FACTGUARD framework. FACTGUARD-BENCH includes 25,220 data examples generated from 16,742 texts. Detailed statistical information regarding FACTGUARD-BENCH is presented in Table 2 and distributions of FACTGUARD-BENCH in terms of domain, question type and length illustrate in Figure 2. The dataset

includes English (en) and Chinese (zh) across two domains, law and books, and features two types of questions: answerable and unanswerable. Unanswerable questions are either due to a lack of evidence or misleading evidence. Example lengths range from 4K to 128k tokens. A comparison of relevant existing datasets and FACTGUARD-BENCH is provided in Table 3.

3.2 Manual Review

To verify the quality of the synthetic data, we randomly sampled 480 examples for manual review. Each example was independently assessed by three annotators with human guidelines (Thakur et al., 2025) classifying example as qualified and unqualified. The human guidelines can be found in Appendix B.1, and the annotators’ qualifications, recruitment, and compensation are described in Appendix B.2. The inter-annotator agreement, as measured by Fleiss’s Kappa (Fleiss, 1971), was substantial ($\kappa = 0.64$), indicating a reliable set of human judgments. The overall quality of FACTGUARD-BENCH is 93.96% which indicates that the synthetic data generated by our method maintains high quality and the details can be found in Appendix B.2.

4 Experiments

4.1 Implementation Details

To evaluate the ability of LLMs on FACTGUARD-BENCH, our experiments included several open-source models that have been instruction-tuned using Supervised Fine-Tuning (SFT) (Ouyang et al., 2022) and Reinforcement Learning from Human Feedback (RLHF) (Stiennon et al., 2020; Bai et al., 2022). Specifically, we utilized the following open-source models: Mistral-Large-Instruct-2411 (123B) (Jiang et al., 2024), DeepSeek-V3-0324 (685B) (Liu et al., 2024), Llama3.3-70B-Instruct (Dubey et al., 2024), Qwen2.5 series models (Yang et al., 2024a). We also obtained evaluation results through API calls for several proprietary models,

MODEL	OVERALL	ENGLISH (EN)			CHINESE (ZH)		
		ANS.	LACK.	MIS.	ANS.	LACK.	MIS.
Proprietary & Open-Source Models							
GEMINI-3-PRO-PREVIEW	71.62	90.50	53.00	74.71	89.00	51.33	74.67
GEMINI1.5-PRO	58.20	86.25	54.60	59.61	83.05	45.45	50.81
GPT-4o	45.90	89.89	41.57	40.78	87.89	37.06	30.30
DEEPSEEK-V4-PRO	56.83	87.75	53.50	60.07	84.67	42.22	40.95
DEEPSEEK-V3-0324	46.39	89.57	34.41	40.17	85.55	39.51	36.61
LLAMA-3.3-70B-INSTRUCT	43.81	90.37	46.19	43.18	87.50	27.62	18.70
MISTRAL-LARGE-2411	45.78	89.89	52.41	45.82	86.33	31.25	18.85
Qwen2.5 Series (Base Instruct)							
QWEN2.5-3B-INSTRUCT	45.39	80.75	48.03	43.50	73.83	39.51	27.66
QWEN2.5-7B-INSTRUCT	47.49	85.02	54.96	42.69	80.86	40.91	30.26
QWEN2.5-14B-INSTRUCT	65.17	85.37	68.65	64.80	85.16	59.79	52.50
QWEN2.5-32B-INSTRUCT	67.67	86.36	71.43	67.65	84.76	63.28	55.43
Qwen2.5 Series (After SFT)							
QWEN2.5-3B-INS-SFT	78.74 ^{↑33.35}	82.62 ^{↑1.87}	84.83 ^{↑36.80}	78.88 ^{↑35.38}	80.47 ^{↑6.64}	96.85 ^{↑57.34}	68.34 ^{↑40.68}
QWEN2.5-7B-INS-SFT	79.11 ^{↑31.62}	83.15 ^{↓1.87}	85.06 ^{↑30.10}	81.54 ^{↑38.85}	80.85 ^{↓0.01}	86.36 ^{↑45.45}	68.23 ^{↑37.97}
QWEN2.5-14B-INS-SFT	79.95 ^{↑14.78}	86.33 ^{↑0.96}	85.29 ^{↑16.64}	81.61 ^{↑16.81}	80.07 ^{↓5.09}	89.16 ^{↑29.37}	69.56 ^{↑17.06}
QWEN2.5-32B-INS-SFT	81.17 ^{↑13.50}	89.04 ^{↑2.68}	89.52 ^{↑18.09}	81.86 ^{↑14.21}	84.77 ^{↑0.01}	92.31 ^{↑29.03}	68.86 ^{↑13.43}
Qwen2.5 Series (After LoRA)							
QWEN2.5-7B-INS-LORA	58.19 ^{↑10.70}	79.54 ^{↓5.48}	72.29 ^{↑17.33}	50.00 ^{↑7.31}	71.93 ^{↓8.93}	64.34 ^{↑23.43}	45.24 ^{↑14.98}
QWEN2.5-32B-INS-LORA	79.13 ^{↑11.46}	83.76 ^{↓2.60}	74.15 ^{↑2.72}	86.95 ^{↑19.30}	86.17 ^{↑1.41}	59.14 ^{↓4.14}	74.13 ^{↑18.70}

Table 4: Comprehensive results on the FACTGUARD-BENCH test set. The table compares SOTA baselines with the Qwen2.5 series before and after SFT / LoRA. Subscripts denote the absolute change versus the corresponding Qwen2.5-{3,7,14,32}B-Instruct baseline (↑ red for gains, ↓ green for drops). Ans., Lack., and Mis. denote Answerable questions, Lack of evidence, and Misleading evidence.

including GPT-4o (Hurst et al., 2024), Gemini1.5 Pro & Gemini-3-Pro-Preview-20251118 (GeminiTeam, 2024), and DeepSeek-V4-Pro (DeepSeek-AI, 2026).

We employ full-parameter SFT training on Qwen2.5 series models (Yang et al., 2024a) to validate the effectiveness of FACTGUARD-BENCH. To demonstrate the model-agnostic generalization of our benchmark, we also adapt our training to two independent model families, Llama-3.1-8B-Instruct and the Mistral-based Ministral series (3B/8B), covering three architectures in total (see Appendix F). We utilized the AdamW optimizer, setting the learning rate to 2×10^{-5} with 2 epoch for full-parameter SFT. We set the warm-up ratio to 0.1 and the weight decay to 0.1. For LoRA configurations, we target the query and value projections (W_q, W_v) in all self-attention layers, utilizing a rank $r = 8$. Appendix A reports the remaining settings, including retrieval, decoding, hardware, and the LoRA rank used in the cross-family runs.

4.2 Evaluation Settings and Metrics

We evaluate the model’s capabilities by assessing the consistency between its predicted answers and




Figure 3: Percentage breakdown of answering unanswerable questions in the FACTGUARD-BENCH test set.

the ground truth. Leveraging the discriminative capabilities of LLM-as-Judge approach (Zheng et al., 2023), our evaluation differentiates between answerable and unanswerable questions, as these two settings require different forms of context-grounded behavior. For answerable questions, a prediction is assigned a score of 1 if it contains the correct information fragments from the ground truth; otherwise, it is scored 0. For unanswerable questions, responses are assigned a score of 1 if they appropriately recognize the unanswerable nature of the question (e.g., through rejection), and a

score of 0 if they generate hallucinatory content.

We selected Qwen2.5-72B-Instruct (Yang et al., 2024b) as the discriminant model for our experiments. Manual validation shows that the LLM-based evaluation achieves about 94% accuracy, with further reliability analysis and cross-judge agreement reported in Appendix C. The full evaluation prompts are provided in Appendix E.

4.3 Experimental Results

4.3.1 The Reliability Gap in Long-Context LLMs

Table 4 reveals a pervasive reliability gap: while frontier models excel at answerable questions, their performance often degrades under adversarial unanswerability. Among the compared frontier systems, Gemini-3-Pro-Preview achieves the strongest overall performance (71.62%) and the best results on answerable questions in both English (90.50%) and Chinese (89.00%), while also showing comparatively robust resistance to misleading evidence. However, its accuracy still drops noticeably on lack-of-evidence cases (53.00% in English and 51.33% in Chinese). DeepSeek-V4-Pro improves over DeepSeek-V3-0324 in overall accuracy (56.83% vs. 46.39%) and yields clearer gains on both lack-of-evidence and misleading-evidence settings, but it still shows a substantial gap between answerable and unanswerable cases. These results reflect a fundamental *Helpfulness-Faithfulness Trade-off*: recent frontier models strengthen long-context retrieval and partial rejection behavior, yet their instruction-tuning objectives still tend to favor plausible answer generation over strict evidence-aware abstention.

Observation: Semantic Entrapment. Several models exhibit substantial accuracy drops on Misleading Evidence tasks across languages; for example, GPT-4o’s accuracy drops to 30.30% in one such setting. This suggests a broader semantic entrapment effect: under long contexts with misleading evidence, models may favor high-confidence lexical matching over stricter factual grounding. A further quantitative analysis is provided in Appendix H.1.

4.3.2 Scaling Laws of Reasoned Rejection

Analysis of the Qwen2.5 series in Table 4 yields two key scaling insights: *SFT as a Reliability Equalizer*: while base models scale traditionally (3B: 45.39% $\rightarrow$ 32B: 67.67%), FACTGUARD serves as a catalyst. The 3B-SFT model (78.74%)

MODEL	TUNING	OVERALL	ANSWERABLE	UNANSWERABLE
<i>Qwen2.5 Series (Base Instruct)</i>				
QWEN2.5-7B-INSTRUCT	-	69.43	94.16	44.77
QWEN2.5-32B-INSTRUCT	-	78.66	94.43	62.93
<i>Qwen2.5 Series (After SFT)</i>				
QWEN2.5-7B-INS-SFT	SFT	83.20 _{+20%}	86.10 _{+1.8%}	80.30 _{+79%}
QWEN2.5-32B-INS-SFT	SFT	79.47 _{+1%}	90.55 _{+4%}	68.41 _{+19%}
<i>Qwen2.5 Series (After LoRA)</i>				
QWEN2.5-7B-INS-LoRA	LoRA	83.68 _{+20.5%}	88.05 _{+6%}	79.33 _{+77%}
QWEN2.5-32B-INS-LoRA	LoRA	82.73 _{+5.2%}	95.95 _{+1%}	68.98 _{+19.6%}

Table 5: Prediction accuracy on the dev set of SQuAD 2.0. We compare the baseline Instruct models with their SFT and LoRA variants. Subscripts denote the percentage change relative to the corresponding Qwen2.5-{7,32}B-Instruct ($\uparrow$ red for gains, $\downarrow$ green for drops).

significantly outperforms the 32B-Base (67.67%), proving that targeted adversarial supervision can unlock high-level rejection capabilities that parameter scaling alone cannot achieve. *Scale Mitigates Out-of-Domain Adaptation Cost*: In-domain, the effect of fine-tuning on answerable accuracy is small and its direction does not follow model size, with both gains and losses appearing at different scales (e.g. Qwen2.5-3B improves in English and Chinese, +1.87 and +6.64, whereas 14B drops 5.09 in Chinese). Out of domain the picture is more systematic: on SQuAD 2.0 the answerable accuracy of the 32B model degrades less than that of the 7B model after SFT (-3.88 versus -8.06 ; Table 5), and larger models retain a higher absolute answerable accuracy post-SFT. We therefore read scale as reducing the out-of-domain cost of adaptation rather than as conferring a general capacity to absorb rejection supervision without trade-off, which the smallest models do not support.

4.3.3 Robustness Across Contextual Scales

Figure 4a and 4b illustrate model reliability as context scales from 4K to 128K tokens. We observe a pronounced reliability attrition in baseline models: as the search space expands, the capacity to verify the absence of information degrades significantly faster than retrieval recall. This suggests that increasing context density often obscures the factual boundaries for standard LLMs.

While all models exhibit an inherent performance decay with increasing context length, those fine-tuned on FACTGUARD-BENCH demonstrate a significantly attenuated decay trajectory (see Figure 4b). By explicitly exposing models to adversarial missing evidence scenarios across diverse context lengths, our framework strengthens the model’s verification logic relative to the input’s positional complexity, maintaining high faithfulness even as the semantic density of the input scales. Appendix D




Figure 4: Model performance under different experimental settings. Panels (a) and (b) report prediction accuracy across input-length intervals for general LLMs and SFT-tuned Qwen models, respectively, while panel (c) shows model accuracy under different data ratios.

provides direct evidence that the improvement is in verification: base models that correctly reject still cite the wrong reason in roughly a third of cases (3B: 66.21% reason accuracy), a gap that fine-tuning largely closes (87.91%).

4.3.4 From Naive Refusal to Reasoned Rejection

A critical requirement for trustworthy LLMs is the transition from categorical refusal to Reasoned Rejection. We evaluate this transition by fine-tuning Qwen2.5-7B-Instruct on FACTGUARD-BENCH and categorizing its unanswerable responses.

As illustrated in Figure 3, baseline models exhibit a tendency toward hallucinatory completion when evidence is missing. Conversely, our SFT variants demonstrate a qualitative shift: instead of generating generic refusals, FACTGUARD-enhanced models can explicitly identify the absent factual information. This capability indicates that the model has transitioned from surface-level keyword matching to a deeper verification of the logical alignment between query constraints and contextual availability.

To quantify this shift rather than infer it from response categories alone, we add a stricter metric that checks whether a correct rejection is accompanied by a correctly grounded reason (Appendix D). Conditioned on a correct rejection, reason accuracy rises from 66.21% to 87.91% for Mistral-3B and from 73.94% to 92.55% for Mistral-8B after SFT, confirming that fine-tuned models do not merely reject more often but increasingly identify the specific missing or substituted evidence.

The model’s reasoning capability is highly sensitive to the mixture ratio between answerable and unanswerable training examples. As shown in Figure 4c, even a sparse inclusion of unanswerable ex-

amples (e.g., an 8:1 ratio) significantly recalibrates the model’s faithfulness threshold. This recalibration is not a loss of intelligence but a strategic realignment toward context-faithfulness, effectively suppressing the sycophancy bias identified in Section 4.3.1.

4.4 Analysis of Performance Trade-offs

4.4.1 Cross-BENCHMARK Generalization

To verify that FACTGUARD-BENCH fosters universal grounding principles rather than distribution-specific patterns, we evaluated models on the human-annotated SQuAD 2.0 benchmark.

Zero-Shot Reliability Transfer: Models trained on our synthetic legal and literary data show substantial reliability gains on the Wikipedia-based SQuAD 2.0 task. As reported in Table 5, the Qwen2.5-7B-Instruct model, after fine-tuning on FACTGUARD-BENCH, achieves a nearly twofold increase in unanswerable question accuracy, rising from 44.77% (baseline) to 80.30% (SFT).

The Helpfulness-Faithfulness Frontier: The significant gain in faithfulness necessitates a tactical trade-off with answerable accuracy. For the Qwen2.5-7B-Instruct model on SQuAD 2.0, unanswerable accuracy rises from 44.77% to 80.30% while answerable accuracy shifts from 94.16% to 86.10% (see Table 5). We characterize this phenomenon as an *Elevation of the Evidentiary Threshold*: the model adopts a more stringent verification criterion, prioritizing factual grounding over speculative helpfulness. This joint pattern, a large gain in abstention accompanied by a smaller loss on answerable items, is the signature the characterization predicts, and distinguishes it from a general degradation in comprehension, which would lower both. This shift represents a deliberate recalibration of the model’s epistemic uncertainty, which is

essential for decision-critical domains where the categorical cost of a hallucination far outweighs that of a cautious, yet verifiable, refusal.

Transfer to Long-Context, Human-Written

Data: SQuAD 2.0 contexts are short (mean 167 tokens), so we further evaluate on PeerQA (Baumgärtner et al., 2025), whose contexts are full scientific papers and whose questions are drawn from peer reviews (Appendix G). The rejection behavior carries over to this long-context regime: for Ministral-8B, unanswerable recall rises from 0.09 to 0.62 and unanswerable accuracy from 40.18% to 79.46%. Under the binary-classification protocol of Baumgärtner et al. (2025), Macro-F1 improves for LoRA (0.52 $\rightarrow$ 0.59) but remains flat for full SFT (0.52 $\rightarrow$ 0.53), since this metric penalizes the drop in answerable recall (0.99 $\rightarrow$ 0.57) that stricter abstention entails on an answerable-heavy test set.

4.4.2 Mitigation via Parameter-Efficient Tuning (LoRA)

To mitigate the observed trade-off on answerable questions, we investigate LoRA as a modular regularization strategy.

In-Domain Adversarial Robustness: As reported in Table 4, while the Qwen2.5-32B-Ins-SFT model achieves the highest overall accuracy (81.17%) on FACTGUARD-BENCH, the LoRA variant demonstrates exceptional robustness in misleading scenarios. Notably, Qwen2.5-32B-Ins-LoRA achieves peak accuracy in the Misleading Evidence category across both English (86.95%) and Chinese (74.13%), significantly outperforming all other fine-tuning configurations. This suggests that parameter-efficient updates are particularly effective at suppressing the sycophancy bias in high-conflict adversarial tasks.

Preserving Retrieval Proficiency: The advantages of LoRA are even more pronounced in cross-domain generalization. On the SQuAD 2.0 benchmark (Table 5), while full SFT leads to an 8% and 4% drop in answerable accuracy for the 7B and 32B models respectively, LoRA remarkably yields a 1% improvement (rising from 94.43% to 95.95%) for the 32B variant.

These results indicate that for high-capacity models, parameter-efficient updates act as a Modular Extension of reliability logic. By constraining the update space, LoRA integrates superior rejection capabilities for misleading premises while preserving the pre-existing knowledge density and retrieval proficiency of the pre-trained model. We note, how-

ever, that which method is preferable is not universal but depends on the distribution and the metric: full SFT attains the highest overall accuracy in-domain and the largest gains in abstention, whereas LoRA is favored where answerable accuracy must be preserved, a contrast that is most visible on PeerQA (Appendix G).

5 Related Work

5.1 Machine Reading Comprehension

Machine reading comprehension (MRC) is a hot research topic in the field of NLP, which focuses on reading documents and answering related questions (Liu et al., 2019; Baradaran et al., 2022). Over the years, machine reading comprehension has garnered significant attention from both academia and industry (Hermann et al., 2015; Liu et al., 2019). With the rapid advancements of large language models (LLMs) (Zhao et al., 2023; Liu et al., 2023), retrieval-augmented generation (RAG) has emerged as a promising framework for tackling reading comprehension tasks across diverse specialized domains (Zhao et al., 2024; Lewis et al., 2020). Nevertheless, even state-of-the-art RAG frameworks are susceptible to retrieval accuracy limitations (Hu et al., 2019; Wang et al., 2025), which emphasizes the critical importance of facticity (Jacovi et al., 2025; Bi et al., 2025), i.e., the ability of a model to generate factually consistent and verifiable responses in information-seeking scenarios. In this work, we emphasize scalable and robust evaluation of answerable and unanswerable questions in reading comprehension.

5.2 Long Context LLMs and Benchmarks

Recent studies have emphasized the importance of extending positional embeddings to improve the ability of LLMs to handle long contexts effectively (Su et al., 2024; Press et al., 2021; Chi et al., 2022). Closed-source LLMs, in particular, have emerged as leaders in long-context modeling, benefiting from progressively larger context windows. For instance, models such as GPT-4 (Achiam et al., 2023) and Gemini Pro 1.5-1000k (Team et al., 2024) are capable of processing increasingly longer documents, with context lengths ranging from 128k to 1000k tokens. Similarly, open-source LLMs, including Qwen 2.5 (Yang et al., 2024a) and DeepSeek (DeepSeek-AI, 2024), also support context lengths of at least 128k tokens. Key benchmarks for assessing long-context capa-

bilities include NIAH (gkamradt, 2023; Yu et al., 2025), Longbench Series (Bai et al., 2024, 2025), LooGLE (Li et al., 2024), L-Eval (An et al., 2024), and, in the financial domain, DocFinQA (Reddy et al., 2024). These benchmarks measure comprehension and reasoning over long inputs, but their questions are answerable by construction, so they do not test whether a model recognizes that the context fails to support a question. In FACTGUARD-BENCH, we utilize a wider range of context lengths to evaluate the LLM’s ability to understand, learn, and reason about information in text.

5.3 Detection of unanswerable questions

Early paradigms (Dasigi et al., 2021; Pang et al., 2022) primarily leveraged constrained-response formats, such as dichotomous or multiple-choice queries. In recent years, studies have increasingly focused on enhancing the ability of reading comprehension models to detect unanswerable questions. Approaches such as those by (Yin et al., 2023) and (Slobodkin et al., 2023) employ prompt engineering—for instance, by incorporating hints such as, “If the question cannot be answered based on the passage, reply ‘unanswerable’”—to improve the model’s ability to detect unanswerable questions. On the other hand, some methods (Agarwal et al., 2023; Deng et al., 2024; Rajpurkar et al., 2018) construct datasets related to unanswer questions to evaluate the model’s ability of detection of unanswerable questions. For example, Agarwal et al. (2023) categorized unanswerable questions into five distinct types: Incomplete Information, Future Questions, Incorrect Information, Ambiguous, and Unmeasurable. They also introduced QnotA—a dataset consisting of 400 samples designed to support this taxonomy. Closest to our setting is PeerQA (Baumgärtner et al., 2025), which pairs full scientific papers with questions taken from peer reviews and includes an unanswerable-question classification task. However, these datasets are often small in scale, require expensive annotation manpower, and have a short context information.

6 Conclusion

This paper addresses the critical challenge of model reliability in long-form document understanding by introducing a multi-task workflow to synthesize high-fidelity datasets, revealing through its benchmark that current LLMs suffer from overconfident hallucinations as context scales. The empirical find-

ings demonstrate that targeted adversarial supervision and precise internal calibration can unlock robust reasoning and rejection capabilities. Ultimately, this work shifts the focus of long-context research from mere retrieval capacity to the critical balance between helpfulness and faithfulness, paving the way for safer, more reliable AI systems.

Limitations

Our work has two primary limitations. First, while our current dataset relies on naturally occurring long-form text, it is limited to the domains of literature (books) and legal texts (law). Consequently, the learned rejection patterns may require specific calibration or expansion when applied to other highly specialized fields. Second, FACTGUARD primarily targets factual grounding; future work must further address complex logical hallucinations where evidence is present but the model executes flawed multi-step reasoning.

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, and 1 others. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Ayushi Agarwal, Nisarg Patel, Neeraj Varshney, Mihir Parmar, Pavan Mallina, Aryan Shah, Srihari Raju Sangaraju, Tirth Patel, Nihar Thakkar, and Chitta Baral. 2023. Can nlp models’ identify’, ‘distinguish’, and ‘justify’ questions that don’t have a definitive answer? In *The 61st Annual Meeting Of The Association For Computational Linguistics*.
- I Elaine Allen and Christopher A Seaman. 2007. Likert scales and data analyses. *Quality progress*, 40(7):64–65.
- Alfonso Amayuelas, Kyle Wong, Liangming Pan, Wenhu Chen, and William Yang Wang. 2024. *Knowledge of knowledge: Exploring known-unknowns uncertainty with large language models*. pages 6416–6432. Association for Computational Linguistics.
- Chenxin An, Shansan Gong, Ming Zhong, Xingjian Zhao, Mukai Li, Jun Zhang, Lingpeng Kong, and Xipeng Qiu. 2024. *L-eval: Instituting standardized evaluation for long context language models*. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14388–14411, Bangkok, Thailand. Association for Computational Linguistics.
- Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, and 1

- others. 2022. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*.
- Yushi Bai, Xin Lv, Jiajie Zhang, Hongchang Lyu, Jiankai Tang, Zhidian Huang, Zhengxiao Du, Xiao Liu, Aohan Zeng, Lei Hou, Yuxiao Dong, Jie Tang, and Juanzi Li. 2024. [LongBench: A bilingual, multitask benchmark for long context understanding](#). pages 3119–3137, Bangkok, Thailand. Association for Computational Linguistics.
- Yushi Bai, Shangqing Tu, Jiajie Zhang, Hao Peng, Xiaozhi Wang, Xin Lv, Shulin Cao, Jiazheng Xu, Lei Hou, Yuxiao Dong, Jie Tang, and Juanzi Li. 2025. [LongBench v2: Towards deeper understanding and reasoning on realistic long-context multitasks](#). pages 3639–3664, Vienna, Austria. Association for Computational Linguistics.
- Razieh Baradaran, Razieh Ghiasi, and Hossein Amirkhani. 2022. A survey on machine reading comprehension systems. *Natural Language Engineering*, 28(6):683–732.
- Tim Baumgärtner, Ted Briscoe, and Iryna Gurevych. 2025. PeerQA: A scientific question answering dataset from peer reviews. In *Proceedings of the 2025 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL)*.
- Baolong Bi, Shenghua Liu, Yiwei Wang, Lingrui Mei, Junfeng Fang, Hongcheng Gao, Shiyu Ni, and Xueqi Cheng. 2025. Is factuality enhancement a free lunch for llms? better factuality can lead to worse context-faithfulness. *The Thirteenth International Conference on Learning Representations*.
- Ye Chen, Wei Cai, Liangmin Wu, Xiaowei Li, Zhanxuan Xin, and Cong Fu. 2023. [Tigerbot: An open multilingual multitask llm](#). *Preprint*, arXiv:2312.08688.
- Ta-Chung Chi, Ting-Han Fan, Peter J Ramadge, and Alexander Rudnicky. 2022. Kerple: Kernelized relative positional embedding for length extrapolation. *Advances in Neural Information Processing Systems*, 35:8386–8399.
- Gheorghe Comanici, Eric Bieber, Mike Schaekermann, Ice Pasupat, Naveen Sachdeva, Inderjit Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, and 1 others. 2025. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*.
- Pradeep Dasigi, Kyle Lo, Iz Beltagy, Arman Cohan, Noah A Smith, and Matt Gardner. 2021. A dataset of information-seeking questions and answers anchored in research papers. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4599–4610.
- DeepSeek-AI. 2024. [Deepseek-v2: A strong, economical, and efficient mixture-of-experts language model](#). *Preprint*, arXiv:2405.04434.
- DeepSeek-AI. 2026. Deepseek-v4: Towards highly efficient million-token context intelligence.
- Yang Deng, Yong Zhao, Moxin Li, See-Kiong Ng, and Tat-Seng Chua. 2024. Don’t just say “i don’t know”! self-aligning large language models for responding to unknown questions with explanations. Association for Computational Linguistics.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, and 1 others. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Joseph L Fleiss. 1971. Measuring nominal scale agreement among many raters. *Psychological bulletin*, 76(5):378.
- GeminiTeam. 2024. [Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context](#). *Preprint*, arXiv:2403.05530.
- gkamradt. 2023. [Needle in a haystack - pressure testing llms](#).
- Peter Henderson, Mark S Krass, Lucia Zheng, Neel Guha, Christopher D Manning, Dan Jurafsky, and Daniel E Ho. 2022. Pile of law: learning responsible data filtering from the law and a 256gb open-source legal dataset. In *Proceedings of the 36th International Conference on Neural Information Processing Systems*, pages 29217–29234.
- Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. Teaching machines to read and comprehend. *Advances in neural information processing systems*, 28.
- Minghao Hu, Furu Wei, Yuxing Peng, Zhen Huang, Nan Yang, and Dongsheng Li. 2019. Read+ verify: Machine reading comprehension with unanswerable questions. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 6529–6537.
- Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, and 1 others. 2024. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*.
- Alon Jacovi, Andrew Wang, Chris Alberti, Connie Tao, Jon Lipovetz, Kate Olszewska, Lukas Haas, Michelle Liu, Nate Keating, Adam Bloniarz, Carl Saroufim, Corey Fry, Dror Marcus, Doron Kukliansky, Gaurav Singh Tomar, James Swirhun, Jinwei Xing, Lily Wang, Madhu Gurumurthy, and 7 others. 2025. [The facts grounding leaderboard: Benchmarking llms’ ability to ground responses to long-form input](#). *Preprint*, arXiv:2501.03200.

- Albert Q Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, and 1 others. 2024. Mixtral of experts. *arXiv preprint arXiv:2401.04088*.
- 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.
- Jiaqi Li, Mengmeng Wang, Zilong Zheng, and Muhan Zhang. 2024. [LooGLE: Can long-context language models understand long contexts?](#) pages 16304–16333. Association for Computational Linguistics.
- Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, and 1 others. 2024. Deepseek-v3 technical report. *arXiv preprint arXiv:2412.19437*.
- Shanshan Liu, Xin Zhang, Sheng Zhang, Hui Wang, and Weiming Zhang. 2019. Neural machine reading comprehension: Methods and trends. *Applied Sciences*, 9(18):3698.
- Yiheng Liu, Tianle Han, Siyuan Ma, Jiayue Zhang, Yuanyuan Yang, Jiaming Tian, Hao He, Antong Li, Mengshen He, Zhengliang Liu, and 1 others. 2023. Summary of chatgpt-related research and perspective towards the future of large language models. *Meta-Radiology*, page 100017.
- 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.
- Richard Yuanzhe Pang, Alicia Parrish, Nitish Joshi, Nikita Nangia, Jason Phang, Angelica Chen, Vishakh Padmakumar, Johnny Ma, Jana Thompson, He He, and 1 others. 2022. Quality: Question answering with long input texts, yes! In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5336–5358.
- Ofir Press, Noah A Smith, and Mike Lewis. 2021. Train short, test long: Attention with linear biases enables input length extrapolation. *arXiv preprint arXiv:2108.12409*.
- Project Gutenberg. 1971. [Project Gutenberg](#).
- Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018. Know what you don’t know: Unanswerable questions for squad. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 784–789.
- Varshini Reddy, Rik Koncel-Kedziorski, Viet Dac Lai, Michael Krumbick, Charles Lovering, and Chris Tanner. 2024. DocFinQA: A long-context financial reasoning dataset. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Aviv Slobodkin, Omer Goldman, Avi Caciularu, Ido Dagan, and Shauli Ravfogel. 2023. [The curious case of hallucinatory \(un\)answerability: Finding truths in the hidden states of over-confident large language models](#). pages 3607–3625. Association for Computational Linguistics.
- Nisan Stiennon, Long Ouyang, Jeffrey Wu, Daniel Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul F Christiano. 2020. Learning to summarize with human feedback. *Advances in Neural Information Processing Systems*, 33:3008–3021.
- Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. 2024. [Roformer: Enhanced transformer with rotary position embedding](#). *Neurocomputing*.
- Gemini Team, Petko Georgiev, Ving Ian Lei, Ryan Burnell, Libin Bai, Anmol Gulati, Garrett Tanzer, Damien Vincent, Zhufeng Pan, Shibo Wang, and 1 others. 2024. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *arXiv preprint arXiv:2403.05530*.
- Aman Singh Thakur, Kartik Choudhary, Venkat Srinik Ramayapally, Sankaran Vaidyanathan, and Dieuwke Hupkes. 2025. [Judging the judges: Evaluating alignment and vulnerabilities in LLMs-as-judges](#). Association for Computational Linguistics.
- Fei Wang, Xingchen Wan, Ruoxi Sun, Jiefeng Chen, and Serkan O Arik. 2025. [Astute RAG: Overcoming imperfect retrieval augmentation and knowledge conflicts for large language models](#). Association for Computational Linguistics.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, and 1 others. 2025. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.
- An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, and 1 others. 2024a. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*.
- An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, and 1 others. 2024b. Qwen2.5 technical report. *arXiv preprint arXiv:2412.15115*.
- Zhangyue Yin, Qiushi Sun, Qipeng Guo, Jiawen Wu, Xipeng Qiu, and Xuanjing Huang. 2023. [Do large language models know what they don’t know?](#) Association for Computational Linguistics.

Yifei Yu, Qian-Wen Zhang, Lingfeng Qiao, Di Yin, Fang Li, Jie Wang, Chen Zeng Xi, Suncong Zheng, Xiaolong Liang, and Xing Sun. 2025. [Sequential-NIAH: A needle-in-a-haystack benchmark for extracting sequential needles from long contexts](#). Association for Computational Linguistics.

Siyun Zhao, Yuqing Yang, Zilong Wang, Zhiyuan He, Luna K Qiu, and Lili Qiu. 2024. Retrieval augmented generation (rag) and beyond: A comprehensive survey on how to make your llms use external data more wisely. *arXiv preprint arXiv:2409.14924*.

Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, and 1 others. 2023. A survey of large language models. *arXiv preprint arXiv:2303.18223*.

Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, and 1 others. 2023. Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in neural information processing systems*, 36:46595–46623.

A Reproducibility

This section consolidates the settings needed to reproduce our results. All prompts, configuration files, and launch scripts are released in the repository linked in the abstract.

Data synthesis. The workflow is driven by Qwen2.5-72B-Instruct. Raw documents are kept if they contain between 2K and 128K tokens. Each document is split at line boundaries into fragments of 500–1000 characters, and fragments longer than 5,000 characters are discarded. For the RAG-based consistency check of Section 2.4, fragments are rechunked with a recursive character splitter (chunk size 1,024, overlap 512), embedded with an in-house Hunyuan retrieval embedding model, indexed with FAISS, and queried with cosine similarity; we retrieve the top $k = 5$ fragments and apply no score threshold, so filtering is decided by the entailment judgment rather than by the retrieval score.

Fine-tuning. We train with AdamW, learning rate 2×10^{-5} , 2 epochs, cosine schedule, warmup ratio 0.1, weight decay 0.1, bf16 precision, global batch size 128, and micro batch size 1 per data-parallel rank, with a maximum sequence length of 131,072 tokens. LoRA adapters target the query and value projections (W_q, W_v) of every self-attention layer. Our original Qwen2.5 experiments use rank $r = 8$; the cross-family runs in Appendix F use rank $r = 128$ with $\alpha = 256$ and dropout 0.05, so LoRA rank is not held constant between the two settings and the Qwen and Ministral LoRA columns should not be read as a controlled comparison of adapter capacity. We did not perform a hyperparameter search: the values above were fixed a priori from common practice and reused for every model and both methods, so no configuration was selected on test performance.

Inference and evaluation. Open-weight models are served with vLLM using greedy decoding (temperature 0), at most 1,024 new tokens, and a maximum model length of 131,072 tokens. Proprietary models are queried through their APIs with the same decoding settings where exposed. The judge is Qwen2.5-72B-Instruct, also at temperature 0; the cross-judge check of Appendix C uses DeepSeek-V3-0324. The evaluated checkpoints are the official -BF16 releases for Ministral-3, and the corresponding official Instruct releases for the Qwen2.5 and Llama-3.1 families.

Data splits. FACTGUARD-BENCH is divided into 19,100 training, 1,920 development, and 4,200 test examples (Table 2), with no document appearing in more than one split, and the split is fixed and released with the data. Training and inference were run on a single eight-GPU node, using ZeRO-3 with optimizer and parameter offload for full-parameter SFT of the larger models.

B Manual Review of Details

B.1 Human Annotation Guidelines

Annotation Guidelines for Human Reviewers

You will be given a document, a question, and an answer produced by an LLM. Your task is to judge whether the LLM’s answer is correct, as if you were the LLM’s teacher grading their exam. An answer should be counted as correct if it correctly answers the question based on the content of the document. In doing so, please follow the guidelines below:

- **Answerable questions.** Answers that are generated strictly based on the provided document content, ensuring they remain accurate and free from hallucinations, should be marked *correct*.
- **Unanswerable questions.** Answers that correctly indicate unanswerability and provide appropriate justification based on the text content should be considered *correct*.

If you have trouble judging whether the answer is correct—for instance, because you feel you are lacking the knowledge required to judge so—please indicate this by marking the answer “maybe correct” or “maybe incorrect”, so that we can further review it.

B.2 Synthetic Data Quality

Annotators, recruitment, and compensation.

The annotation was performed by three annotators recruited through an internal crowdsourcing platform operated within our own organization. All three are full-time employees of our organization, so they are both colleagues and participants reached via a crowdsourcing platform. Because the annotation was carried out as part of their regular salaried duties, they received their normal salary and no separate per-task payment. All three are university graduates, are native speakers of Chinese and certified at CET-6 in English, and are experienced with NLP annotation tasks and with reading model-generated text, which qualifies them to assess both the English and the Chinese examples in FACTGUARD-BENCH. No personal or sensitive data was collected from them, and the material they reviewed consisted only of documents drawn from

Category	#Samples	Quality (%)
Answerable	120	94.17
<i>Unanswerable</i>		
Lack of evidence	120	96.67
Misleading evidence	240	92.50
<i>Sub-total</i>	360	93.89
Overall	480	93.96

Table 6: Manual review of synthetic data quality.

Category	#Samples	Quality (%)
Task 1: Answer Consistency Evaluation		
Answerable	80	95.00
Lack of evidence	60	93.33
Misleading evidence	160	93.75
<i>Overall</i>	<i>300</i>	<i>94.00</i>
Task 2: Reasoning Ability for Unanswerable Questions		
Incorrect answers	60	93.33
Direct refusals	111	91.89
Reasoned answers	49	97.96
<i>Overall</i>	<i>220</i>	<i>93.64</i>

Table 7: Manual review results of judgment quality by the discriminative model on Qwen2.5-72B response answers.

Question Type	Agreement Rate	Cohen’s κ
Lack of Evidence	96.3%	0.925
Misleading Evidence	93.0%	0.738
Answerable	93.5%	0.648

Table 8: Agreement and Cohen’s κ between Qwen2.5-72B-Instruct and DeepSeek-V3-0324

public corpora together with model-generated questions and answers.

Procedure and results. Each annotator independently labelled every example as qualified or unqualified following the guidelines in Appendix B.1, which state that the text under review is model-generated, and the final label was decided by majority vote. Annotators were asked to spend at least 10 minutes reading the document and evaluating each example. The results are shown in Table 6. The lower quality in the misleading evidence category was due to the omission of clarifications during the synthesis of answers, as the relevant instructions were not followed. However, the overall quality of 93.96% indicates the high value of our method.

C LLM-Based Evaluation

To ascertain the reliability of the discriminative model employed in our evaluation, we randomly selected 300 samples for manual review based on the discriminant model’s results of discriminating Qwen2.5-72B answers from standardized answers. Consistent with our approach to validating synthetic data quality, we employed a three-person voting mechanism. The outcome of this manual review is detailed in Table 7.

In **Task 1: Answer Consistency Evaluation**, human annotators evaluated whether the discriminative model accurately identified the consistency between its predictions and the ground truth for answerable and unanswerable questions. The results demonstrate that the discriminative model achieved a commendable accuracy of **94.00%** in Task 1.

In **Task 2: Reasoning Ability for Unanswerable Questions**, the manual review focused on whether the discriminative model could accurately classify responses into three distinct categories: *incorrect answers*, *direct refusals*, and *reasoned answers*. The evaluation revealed that the model achieved an overall classification accuracy of **93.64%**. However, due to subtle or ambiguous rejection/clarification phrasing, the model produced more false negatives than false positives. Although slightly outperformed by human benchmarks, the automated system excels in efficiency, consistency, and scalability, enabling robust iterative refinement.

To explore the dynamics between **single-judge and multi-judge paradigms** and validate the robustness of our primary evaluator, we conducted an agreement experiment. Specifically, we measured the cross-model agreement between our primary judge, **Qwen2.5-72B-Instruct**, and an independent secondary judge, **DeepSeek-V3-0324**. The alignment across different question types is summarized in Table 8. The high Cohen’s κ values validate that our evaluation claims are objective. ($\kappa > 0.6$ indicates ‘substantial’ agreement, and $\kappa > 0.8$ denotes ‘almost perfect’ agreement, reinforcing the objectivity of our LLM-as-a-Judge approach.) Given this high level of consensus, adopting a single-judge paradigm (Qwen2.5-72B-Instruct) not only yields substantial savings in computational costs and inference latency, but also ensures a more cohesive and consistent scoring rationale across iterations. This approach maintains rigorous evaluation standards while effectively avoiding the conflicting heuristics

Model	Rej.	Rej.+Reason	Reason Acc.
<i>Ministral-3B-Instruct</i>			
Base	1453	962	66.21
+ LoRA	2131	1469	68.93
+ SFT	2704	2377	87.91
<i>Ministral-8B-Instruct</i>			
Base	1880	1390	73.94
+ LoRA	2072	1704	82.24
+ SFT	2804	2595	92.55

Table 9: Reason-grounded rejection on the 3,500 unanswerable test questions of FACTGUARD-BENCH. **Rej.** counts correct rejections (the surface metric used in the main results); **Rej.+Reason** counts rejections whose stated reason is also grounded in the document; **Reason Acc.** (%) is the fraction of correct rejections that carry a correct reason.

or semantic noise that often arise in multi-judge assemblies.

Because both judges belong to different model families and are trained by different organizations, this agreement also speaks to a potential circularity concern: Qwen2.5-72B-Instruct is used both to synthesize FACTGUARD-BENCH and to score model outputs. If the reported gains merely reflected alignment with one evaluator’s preferred refusal style, an independent judge from another family would not agree at this rate. Combined with the cross-family training results in Appendix F, this indicates that the measured improvements reflect genuine changes in rejection behavior rather than evaluator-specific stylistic preferences.

D Reason-Grounded Rejection Evaluation

The metric used in our main results credits a model whenever it recognizes that a question is unanswerable, but it does not check whether the model rejected for the *right* reason. A model could reject because it failed to locate any relevant passage, which happens to be correct on a lack-of-evidence item but reflects retrieval failure rather than verification. To separate these cases, we add a stricter second-stage evaluation on the 3,500 unanswerable test questions.

Given a response that the first stage already judged to be a correct rejection, the judge (Qwen2.5-72B-Instruct) additionally receives the document and the ground-truth explanation of *why* the question is unanswerable, and decides whether the reason stated by the model matches it, i.e. whether the model identified the specific missing

evidence span or the specific substituted or unsupported constraint. Generic refusals such as “the document does not contain enough information” are scored as correct rejections but incorrect reasons. The prompt is given in Appendix E.

Table 9 reports both metrics. Fine-tuning improves the number of correct rejections, but the larger effect is on reason accuracy conditioned on a correct rejection, which rises from 66.21% to 87.91% for Ministral-3B and from 73.94% to 92.55% for Ministral-8B under SFT. That is, base models that do reject often reject for the wrong reason, whereas fine-tuned models increasingly identify the specific evidentiary defect. This supports the claim in Section 4.3.4 that the gains reflect a shift toward reasoned rejection rather than an increased propensity to refuse, and it locates the base-model failure in verification rather than retrieval.

E LLM-as-Judge Evaluation Prompts

We use the following prompts for the Qwen2.5-72B-Instruct discriminative model in our evaluation experiments. The first stage detects whether a response appropriately refuses or clarifies an unanswerable question. For answerable questions that are not incorrectly rejected, the second stage evaluates answer consistency against the reference answer. The third stage applies only to correctly rejected unanswerable questions and checks whether the stated rejection reason is grounded in the document.

E.1 Stage 1: Refusal and Clarification Detection

Single-point unanswerability detection (LACK.)

Below is a question and an answer about a document.
Question: {question}
Answer: {completion}
 Does the answer indicate that the question cannot be answered or that the relevant information does not exist?
 Please provide the corresponding span from the answer.

Entity-mismatch clarification detection (MIS.)

Below is a question and an answer about a document.
Question: {question}
Answer: {completion}
 It is known that the document does not mention {replaced entity}; instead, it mentions {original entity}. Therefore, the answer needs to clarify this mismatch. Please judge whether the answer contains a clarification similar to the examples below. You only need to analyze whether such a clarification exists and make a judgment; do not revise the answer.
Clarification examples:
 1. The document does not mention {replaced

entity}.

2. The question about {replaced entity} cannot be answered.

3. {replaced entity}({original entity}).

Note that merely mentioning {original entity} in the answer does not count as clarification.

Unsupported-condition clarification detection (MIS.)

Below is a question and an answer about a document.

Question: {question with added condition}

Answer: {completion}

Known condition: {explanation}

Therefore, {added condition} in the question does not exist. The answer needs to discuss this issue, including pointing out that the part of the question related to {added condition} cannot be answered or that the relevant information does not exist. Please judge whether the answer contains such a statement. You only need to analyze and judge whether the answer contains this statement; do not revise the answer.

False-refusal detection for answerable questions

Below is a question and an answer about a document.

Question: {question}

Answer: {completion}

Does the answer indicate that the question cannot be answered or that the relevant information does not exist? Please provide the corresponding span from the answer.

E.2 Stage 2: Answer Consistency Evaluation

For answerable examples, answer consistency is evaluated only when the JSON extracted in Stage 1 indicates that the model did not incorrectly refuse to answer.

Answer consistency prompt

Below is a question and two answers about a document.

Question: {question}

Answer 1: {completion}

Answer 2: {ref_answer}

Ignoring differences in wording, please judge whether the two answers have the same main conclusion for the question. Return the result in the following JSON format:

```
{
  "analysis of whether the conclusions are the same": str,
  "whether the conclusions are the same":
    "yes" | "no"
}
```

E.3 Stage 3: Rejection Reason Grounding

This stage is applied only to unanswerable examples that Stage 1 already scored as a correct rejection, and produces the reason-grounded metric reported in Appendix D.

Rejection reason grounding prompt

A question cannot be answered from the given document. Below are the true reason why it is unanswerable and a model response that correctly declined to answer.

Question: {question}

True reason: {ref_reason}

Model response: {completion}

Judge whether the reason given in the model response matches the true reason. The response must identify the specific defect, i.e. the particular missing piece of evidence, or the particular condition in the question that is substituted or unsupported by the document. A generic statement that the document lacks sufficient information, without identifying which specific fact is missing or incorrect, does not count as a match. Return the result in the following JSON format:

```
{
  "analysis of the stated reason": str,
  "reason matches": "yes" | "no"
}
```

F Cross-Family Generalization

To examine whether the training signal from FACTGUARD-BENCH transfers beyond the Qwen2.5 model family, we apply the same fine-tuning recipe used in our main experiments to two additional families: Llama-3.1-8B-Instruct and the Mistral-based Ministral series (3B and 8B). Table 10 reports results on both FACTGUARD-BENCH and SQuAD 2.0.

Llama-3.1-8B-Instruct. After SFT, the overall accuracy on FACTGUARD-BENCH increases from 41.21% to 77.91%, with especially large gains on unanswerable cases: English misleading-evidence accuracy improves from 41.20% to 83.32%, while Chinese lack-of-evidence accuracy improves from 31.28% to 86.31%. The model also shows improvement on SQuAD 2.0, where overall accuracy rises from 73.22% to 80.34% and unanswerable accuracy increases from 52.95% to 79.85%. The decrease in SQuAD answerable accuracy suggests a trade-off between more conservative evidence verification and answer extraction, but the consistent gains on unanswerable categories indicate that FACTGUARD-BENCH provides a transferable supervision signal for recognizing unsupported or misleading questions.

Ministral series. The same pattern holds on a third architecture. Full SFT raises overall FACTGUARD-BENCH accuracy from 46.29% to 76.79% for Ministral-3B and from 59.02% to 80.81% for Ministral-8B, with gains concentrated in the unanswerable categories in both languages, and it transfers to SQuAD 2.0 (3B: 58.26% → 75.43%; 8B: 59.94% → 80.89%, with unanswerable accuracy rising from 24.46% to 72.03% for 8B). Unlike the Qwen and Llama results, Ministral-3B improves on answerable questions as well (En: 75.25% →

Question Who is described as untameable, irresistible, and peremptory? Answer The document does not describe any character or individual as untameable, irresistible, and peremptory. However, the character of Kenyon in the "Westminster Play" section displays traits that could be interpreted as matching this description, given his determined and unyielding opposition to the closure of the debate, despite the efforts of others to persuade him otherwise. He is portrayed as immovable in his stance and unwilling to compromise, which aligns with the traits you mentioned. Yet, the exact adjectives "untameable, irresistible, and peremptory" are not used to describe him or anyone else. GPT-4o <u>Mr. Inderwick, O.C., is described as "untameable, irresistible, and peremptory"</u> in the context of his examination in a divorce case. 【False】 Llama-3.3-70B-Instruct <i>According to the provided document, there is no specific mention of any individual or entity being described as untameable, irresistible, and peremptory. The document primarily focuses on the technological advancements and market trends in cloud computing, detailing the innovative solutions and services offered by Alibaba Cloud. It discusses various aspects such as the growth of cloud infrastructure, the impact of artificial intelligence, and the strategic directions for future development, but does not include character descriptions or attributes like those mentioned in your question.</i> 【True】 Qwen2.5-7B-Instruct <u>The Bishop of St. Asaph is described as untameable, irresistible, and peremptory</u>, but more specifically, it is Kenyon who is described as untameable and irresistible, and also as having an untameable and irresistible "Yoic" (which he later corrects to "No"). 【False】 Qwen2.5-7B-Instruct-SFT <i>The document does not provide a specific description of someone being untameable, irresistible, and peremptory. The text primarily focuses on other aspects and does not contain this particular combination of words. Therefore, I cannot accurately answer your question based on the information given in the document.</i> 【True】	Question 剑楠对切青的感情发生了怎样的变化? Ho w has Jian Nan's feelings for Ren Jing changed? Answer 根据原文，文中描述的是剑尘对切青的感情变化，而不是剑楠。剑尘对切青的感情经历了从同情到深深的爱慕，再到因丧失信任而转变的过程。... According to the original text, the text describes the change in Jian Chen's feelings for Ren Jing, not Jian Nan. Jian Chen's feelings for Ren Jing went through a process of change from sympathy to deep admiration, and then to loss of trust. ... GPT-4o 通过文档内容分析，<u>剑尘（文中提到的剑楠即是剑尘）</u>对切青的感情经历了以下几个明显的变化... Through the analysis of the document content, <u>Jian Chen (the Jian Nan mentioned in the text is Jian Chen)</u>'s feelings towards Ren Jing have undergone the following obvious changes... 【False】 Llama-3.3-70B-Instruct <u>剑尘对切青的感情经历了从热烈到失望，再到试图超脱的过程。</u>... <u>Jian Chen's feelings</u> for Ren Jing went from passion to disappointment, and then to trying to detach himself. ... 【False】 Qwen2.5-7B-Instruct <u>剑楠对切青的感情从最初的同情和爱慕逐渐转变为对切青的理解和宽容。</u>... <u>Jian Nan's feelings</u> towards Ren Jing gradually changed from sympathy and love to understanding and tolerance of Renjing. ... 【False】 Qwen2.5-7B-Instruct-SFT <i>根据原文，文中描述的是**剑尘**对切青感情的变化，而非**剑楠**。剑尘对切青的感情经历了显著的变化。最初，...</i> <i>According to the original text, the text describes the change in **Jian Chen**'s feelings towards Ren Jing, not **Jian Nan**.</i> Jian Chen's feelings towards Ren Jing have undergone significant changes. Initially,... 【True】
--	---

Figure 5: Case study. An example of lack of evidence in English on the left, and an example of misleading evidence in Chinese on the right (translated below). Red underlined text indicates hallucinatory content and green italicized text indicates useful explanations.

FACTGUARD-BENCH								SQuAD 2.0		
Model	Overall	English (En)			Chinese (Zh)			Overall	Ans.	Unans.
		Ans.	Lack.	Mis.	Ans.	Lack.	Mis.			
Llama-3.1-8B-Instruct										
Base	41.21	82.00	58.35	41.20	82.67	31.28	13.14	73.22	93.55	52.95
+ SFT	77.91	83.25	72.08	83.32	69.67	86.31	74.19	80.34	80.83	79.85
Ministral-3B-Instruct										
Base	46.29	75.25	46.50	51.50	63.33	32.44	29.24	58.26	94.74	21.88
+ LoRA	59.43	59.00	86.00	58.93	43.00	66.89	46.57	65.60	89.96	41.31
+ SFT	76.79	78.00	77.50	82.00	69.67	78.67	70.19	75.43	82.00	68.88
Ministral-8B-Instruct										
Base	59.02	87.75	45.50	67.79	82.67	36.00	47.24	59.94	95.53	24.46
+ LoRA	62.67	84.75	59.50	53.93	73.67	67.11	62.67	72.10	95.04	49.22
+ SFT	80.81	86.50	77.17	86.57	81.33	76.89	74.57	80.89	89.78	72.03

Table 10: Cross-family generalization. Performance of Llama-3.1-8B-Instruct and the Ministral series on FACTGUARD-BENCH (test set) and SQuAD 2.0 (dev set) before and after fine-tuning with FACTGUARD-BENCH. Ans., Lack., Mis., and Unans. denote Answerable, Lack of evidence, Misleading evidence, and Unanswerable, respectively. Bold marks the best overall accuracy within each model family.

78.00%; Zh: 63.33% → 69.67%), indicating that the answerable-accuracy trade-off discussed in Section 4.4.1 is model-dependent rather than an intrinsic cost of our supervision. Consistent with the LoRA analysis in Section 4.4.2, LoRA is more

conservative in-domain but better preserves answerable accuracy out-of-domain: Ministral-8B + LoRA reaches 95.04% SQuAD answerable accuracy, essentially matching its base model (95.53%), whereas SFT drops to 89.78%.

Dataset	Samples	Mean	Min / Max
SQuAD 2.0	11,873	167	30 / 780
PeerQA	579	11,645	2,670 / 52,360
FACTGUARD-BENCH	4,200	15,554	935 / 127,665

Table 11: Context length of the evaluation datasets, measured in Llama-3 tokenizer tokens. Sample counts refer to the split used for evaluation (FACTGUARD-BENCH test set).

Model	Binary classification			Gen.
	Unans. P/R	Ans. P/R	M-F1	R-L
<i>Ministral-3B-Instruct</i>				
Base	0.54 / 0.12	0.80 / 0.97	0.54	0.2825
+ LoRA	0.28 / 0.52	0.83 / 0.64	0.54	0.3073
+ SFT	0.25 / 0.59	0.83 / 0.53	0.50	0.2803
<i>Ministral-8B-Instruct</i>				
Base	0.62 / 0.09	0.80 / 0.99	0.52	0.2815
+ LoRA	0.37 / 0.32	0.82 / 0.85	0.59	0.3530
+ SFT	0.28 / 0.62	0.85 / 0.57	0.53	0.2945

Table 12: PeerQA under the **binary-classification** protocol of Baumgärtner et al. (2025) (414 answerable / 112 unanswerable). P/R are per-class precision and recall, M-F1 is Macro-F1 over the two classes, and R-L is ROUGE-L against the GPT-4-augmented reference answer on the answerable subset.

Model	Ans.	Unans.	Overall
<i>Ministral-3B-Instruct</i>			
Base	26.81	28.57	27.19
+ LoRA	38.41	66.07	44.30
+ SFT	46.86	72.32	52.28
<i>Ministral-8B-Instruct</i>			
Base	35.02	40.18	36.12
+ LoRA	34.54	47.32	37.26
+ SFT	49.28	79.46	55.70

Table 13: PeerQA under the **LLM-as-judge** protocol used throughout this paper (Qwen2.5-72B-Instruct), which scores the generated response rather than an explicit label. Accuracy (%) on the same 526 samples.

Taken together, the consistent overall gains across Qwen2.5, Llama-3.1, and Ministral indicate that the improvements come from the supervision signal in FACTGUARD-BENCH rather than from adapting to one particular architecture or to the stylistic preferences of the Qwen-based generator used to build the data.

G Out-of-Domain Evaluation on PeerQA

SQuAD 2.0 provides human-annotated out-of-domain validation, but its contexts are short (Table 11), so it cannot show whether the learned rejection behavior survives in the long-context regime it

was trained for. We therefore additionally evaluate on PeerQA (Baumgärtner et al., 2025), a human-written scientific QA dataset in which the contexts are full papers and the questions, drawn from peer reviews, include both answerable and unanswerable cases. This is out of domain along three axes at once: scientific rather than legal or literary text, human-written rather than synthesized questions, and English-only.

We use the 526 samples that carry a non-null answerability label, excluding 53 entries, yielding 414 answerable and 112 unanswerable questions. We report two protocols, which measure different things and, as discussed below, do not agree.

Binary-classification protocol. Following Baumgärtner et al. (2025), we parse the model’s explicit answerable/unanswerable decision and score it against the gold label, reporting per-class precision and recall, Macro-F1, and ROUGE-L on the answerable subset (Table 12). Under this protocol the transfer is *not* uniformly positive. LoRA improves Macro-F1 on the 8B model (0.52 $\rightarrow$ 0.59) and leaves it unchanged on 3B (0.54), whereas full SFT leaves it flat or lower (8B 0.52 $\rightarrow$ 0.53; 3B 0.54 $\rightarrow$ 0.50). Generation quality follows the same ordering, with LoRA the only variant that consistently improves over the base (8B: 0.2815 $\rightarrow$ 0.3530).

The per-class numbers explain why the base models score as well as they do. Both base models are strongly biased toward the majority answerable class (Ans. recall 0.97–0.99) and almost never abstain (Unans. recall 0.09–0.12). Because the test set is answerable-heavy, this bias inflates Macro-F1 to 0.52–0.54 while the model in practice has essentially no capacity to abstain. Fine-tuning moves the operating point in the opposite direction: SFT raises unanswerable recall to 0.59–0.62 at the cost of answerable recall (8B: 0.99 $\rightarrow$ 0.57), which Macro-F1 penalizes on this label distribution. LoRA lands in between, retaining high answerable recall (0.85) while introducing moderate rejection ability (0.09 $\rightarrow$ 0.32), which is why it attains the highest Macro-F1.

LLM-as-judge protocol. Under the protocol used elsewhere in this paper, which scores the generated response rather than a parsed label, the ordering reverses: SFT is the strongest variant, raising unanswerable accuracy from 40.18% to 79.46% and answerable accuracy from 35.02% to 49.28% on the 8B model (Table 13).

Reconciling the two protocols. The divergence is a property of the protocols, not of the models. The binary protocol requires an explicit label, whereas the judge also credits *implicit* rejections, where a model does not emit an “unanswerable” verdict but its text declines to answer (“the document does not specify...”, “it is not possible to determine...”). Fine-tuned models frequently reject in this implicit style, so they are rewarded by the judge and penalized by the label parser. The two protocols also score different objects: the binary protocol scores a decision, while the judge scores answer content, which is why SFT can simultaneously lower answerable recall and raise answerable accuracy, i.e. answer less often but more correctly.

We therefore do not claim uniform improvement on PeerQA. What both protocols do support is that fine-tuning on FACTGUARD-BENCH substantially increases the willingness and ability to abstain on a long-context, human-written, out-of-domain benchmark (unanswerable recall $0.09 \rightarrow 0.62$; unanswerable accuracy $40.18\% \rightarrow 79.46\%$ for 8B). Which variant is preferable depends on the target operating point: LoRA if the deployment is answerable-heavy and a balanced label distribution matters, full SFT if unrecognized unanswerable questions are the more costly error.

H Case Study

To facilitate a clear and intuitive comparison of various models for generating reasoning-based answers to unanswerable questions, we present two distinct scenarios in Figure 5. In the lack of evidence scenario, GPT4o and Qwen2.5-7B-Instruct display significant hallucination in their responses, frequently generating factually incorrect answers. Llama-3.3-70B-Instruct had both rejection tendencies and reasoning, making it a highly desirable response. In the misleading evidence scenario, all baseline models are misled by the question, resulting in incorrect answers. However, after fine-tuning with SFT, this issue is mitigated, enabling the models to provide accurate responses that align with the given text.

H.1 Semantic Entrapment Cases

To further illustrate the failure mode discussed in Section 4.3.1, we first quantify it and then present two representative misleading-evidence cases from Gemini-3-Pro-Preview.

Quantifying semantic entrapment. If models fail on misleading evidence because they rely on lexical familiarity instead of verifying the decisive constraint, then the failure rate should depend on how closely the substituted entity resembles the original: a near-identical substitution offers the strongest illusion of a match. We test this on the Gemini-3-Pro-Preview results by grouping the entity-substitution items by the Normalized Levenshtein Similarity (NLS) between the substituted entity and the original one, where a higher NLS denotes a more subtle, higher-overlap substitution.

NLS interval	Accuracy (%)
< 0.4 (significant change)	71.2
0.4 – 0.6	63.5
0.6 – 0.8	59.9
> 0.8 (subtle modification)	56.2

Accuracy decreases monotonically as lexical similarity increases ($71.2\% \rightarrow 56.2\%$). The difficulty of an item therefore tracks its surface similarity to a legitimate match rather than any property of the answer itself, which is the signature of a verification failure: the more a misleading entity resembles the correct one, the more readily the model accepts high-level lexical familiarity in place of an explicit factual check. This also argues against the concern that our substitutions introduce an easy lexical shortcut, since the hardest items are precisely those where the surface form is least distinguishable.

In the two cases below, the retrieved passage has high lexical overlap with the question, but the decisive constraint is either substituted or unsupported. The model therefore produces confident answers instead of recognizing that the question is unanswerable from the given document.

Case 1: Entity substitution in a legal settlement question

Misleading question. What was the total settlement amount and requested fee in *In re American Realty Capital Management, Inc. Litigation*?

Overlapping evidence. The document states that, in *In re American Realty Capital Properties, Inc. Litigation*, the proposed settlement was estimated at \$1,025,000,000 and Class Counsel requested a fee equal to 12.4% of the settlement amount.

Model response. “Based on the provided document, the total settlement amount is \$1.025 billion (\$1,025,000,000). Class Counsel is requesting 12.4% of this settlement amount as their fee.”

Model failure. The model answered that the total settlement was \$1.025 billion and that the requested fee was approximately \$127.19 million.

Why this is misleading. The question asks about *American Realty Capital Management*, whereas the evidence

concerns *American Realty Capital Properties*. Because the monetary amount, litigation name, and fee percentage are otherwise highly similar, the model treats the passage as supporting evidence and misses the substituted entity constraint.

Case 2: Unsupported temporal constraint in an over-time question

Misleading question. Did the Defendant United States Department of Veterans Affairs compensate Plaintiff Cindie Hamm *for the month of May 2016* at a rate of not less than one and one-half times the regular rate for hours worked over 40 hours in a workweek?

Overlapping evidence. The document states that the Defendant employed Cindie Hamm for more than forty hours in a workweek without compensating her at one and one-half times the regular rate for overtime hours.

Model response. “Based on the provided document, no. According to Paragraph 29 of the complaint, Defendant United States Department of Veterans Affairs employed Plaintiff Cindie Hamm for more than forty hours in a workweek without compensating her at a rate of not less than one and one-half times her regular rate for all hours worked in excess of 40 hours. Furthermore, Paragraphs 13 and 14 establish that this failure to properly compensate her occurred throughout the period from April 1, 2015, until 2017, which would include the month of May 2016.”

Model failure. The model answered “no” and justified the answer by matching the overtime-compensation language in the passage.

Why this is misleading. The central temporal condition, “for the month of May 2016,” is not directly supported by the cited evidence. The model focuses on the highly overlapping overtime clause and fails to verify whether the specific month required by the question is grounded in the document.