

CONF-RAG: Confidence-Guided Retrieval-Augmented Generation

Yin Huang¹, Yifan Ethan Xu¹, Kai Sun¹, Vera Yan¹, Alicia Sun², Haidar Khan¹,
Jimmy Nguyen¹, Sean Chen¹, Mohammad Kachuee¹, Zhaojiang Lin¹,
Yue Liu¹, Aaron Colak¹, Anuj Kumar¹, Wen-tau Yih², Xin Luna Dong¹

¹Meta Reality Labs, ²Meta FAIR

Correspondence: maggiehuang@meta.com, lunadong@meta.com

Abstract

Production Retrieval-Augmented Generation (RAG) pays retrieval latency and serving cost on every query – even when the model already knows the answer. We present CONF-RAG, a confidence-guided strategy that invokes RAG only when a large language model (LLM) is unlikely to be correct, cutting external retrievals by 5-39% and tail latency by up to 800ms (P90) in a real RAG system while matching the answer quality of always-on RAG. CONF-RAG is enabled by CONFQA, a lightweight fine-tuning strategy (one epoch over 3K knowledge-graph facts) that teaches an LLM to answer only when confident and otherwise say "I am unsure", cutting hallucination rates from 20-40% to 5% or below across multiple factuality benchmarks. Two choices make it work in practice: a dampening prompt that discourages overconfident hallucinations, and training on atomic facts that generalizes confidence calibration to unseen domains and question types. At serving time, LLM-QA and RAG run in parallel with early-stopping, so confident answers return immediately.

1 Introduction

Despite the remarkable capabilities that Large Language Models (LLMs) have demonstrated, they still *hallucinate factual statements* (Zhou et al., 2021; Dziri et al., 2022; Rawte et al., 2023; Mishra et al., 2024; Kalai et al., 2025). It is broadly accepted that factual content shall not be *fabricated* or *generated*, but anchored in *internally parameterized neural knowledge* or *externally recorded symbolic content* (knowledge graphs, webpages, etc). Significant progress has been made on both fronts: knowledge internalization via pre-training (Grattafiori et al., 2024; Kim et al., 2025) and external knowledge use via Retrieval-Augmented Generation (RAG) (surveyed in (Zhou et al., 2024; Gupta et al., 2024; Ni et al., 2025)). However, a critical question remains: *when*

should LLMs rely on parameterized knowledge versus external sources?

Ideally, an LLM would consult external sources only when it detects uncertainty in its own knowledge, but reliable recognition of these knowledge limits remains unsolved. Token-level confidence methods require monitoring hidden-state signals from intermediate activations, often impractical in deployment (Jiang et al., 2023; Su et al., 2024; Li et al., 2025b). Prompting LLMs to explicitly express confidence is undermined by the well-known over-confidence (Wei et al., 2024a; Ni et al., 2024; Xiong et al., 2024). Recent fine-tuning approaches for uncertainty calibration help but remain imperfect, either only modestly reducing hallucinations or substantially suppressing answer correctness (Cheng et al., 2024; Yang et al., 2024b; Grattafiori et al., 2024; Zhang et al., 2024; Kapoor et al., 2025; Li et al., 2025c).

This paper studies a fine-tuning approach that trains LLMs to recognize low-confidence factual knowledge and trigger RAG accordingly. Several key intuitions let us improve over prior work on eliminating hallucinations and avoiding unnecessary RAG invocations, yielding three contributions.




Figure 1: On factual QA benchmarks, fine-tuned model CONFQA reduces hallucination to 5% or below; CONF-RAG increases accuracy by 57.2 points for SimpleQA and 3.6 points for CRAG while cutting latency (Table 3).

First, an effective fine-tuning method for confidence recognition, called CONFQA, forms the core

of our triggering strategy: it checks the LLM’s inherent answer and teaches it to state “I am unsure” when the answer is incorrect. Two design choices distinguish it from existing work (Zhang et al., 2024; Cohen et al., 2024; Cheng et al., 2024; Yang et al., 2024b; Kapoor et al., 2025): (i) a *dampening* prompt—“Answer only if you are confident”—that shapes the LLM’s behavior, and (ii) training data of exclusively *simple* entity attribute questions, whose atomic facts serve as building blocks for complicated statements, that let the confidence behavior generalize to broader domains.

Second, a RAG triggering strategy called CONFRAG, which casts triggering as detecting low-confidence answers: LLM-based QA and a RAG pipeline run in *parallel*, with RAG halted early once the LLM gives a confident answer rather than “I am unsure”. Unlike methods inspecting hidden-state signals (Su et al., 2024; Jiang et al., 2023), or prompt-instructed certainty responses (Ni et al., 2024), CONFRAG is lightweight, low latency, more accurate than baselines and broadly applicable.

Third, a comprehensive study across 8 benchmarks spanning short-form, long-form QA, and general knowledge QA. CONFRAG reaches $\approx 95\%$ accuracy with a perfect RAG system, and with a real one matches always-on RAG accuracy while cutting P50 latency by over 600ms on benchmarks like CRAG (Yang et al., 2024a). We also find two standalone uses of CONFQA when RAG is unavailable (such as in Speech-in Speech-out systems (Xie and Wu, 2024; Nguyen et al., 2022)): we recommend CONFQA *without* the dampener in inference to preserve accuracy with fewer hallucinations; we recommend CONFQA *with* the dampener, to drive hallucinations to under 5% (see Figure 1).

2 Related Work

Our work relates to two bodies of work: LLM hallucination suppression and RAG triggering.

LLM hallucination suppression: LLMs are argued to be inherently prone to hallucinations due to the high-entropy next-token pre-training objective (Kalai et al., 2025), prompting many mitigation efforts (surveyed in (Tonmoy et al., 2024)). One direction incentivizes factually correct generation—leveraging internal parameters (Li et al., 2023; Chuang et al., 2024), reciting pre-training passages (Sun et al., 2023), or fine-tuning the model to select correct tokens (Tian et al., 2023; Xie

et al., 2025; Lin et al., 2024); the other detects and suppresses likely-hallucinated content, e.g., verifying output tokens against internal or external corpus (Lei et al., 2023; Dhuliawala et al., 2023; Varshney et al., 2023).

CONFQA belongs to the *confidence-based softening and abstaining* category, whose roots predate LLMs (Kadavath et al., 2022; Mielke et al., 2022; Lin et al., 2022). Confidence has been derived via semantic entropy (Kuhn et al., 2023), token-probability aggregation (Duan et al., 2024), *Local Intrinsic Dimension* (Yin et al., 2024), self-reported confidence in verbal (Zhou et al., 2023) or numerical (Tian et al., 2023; Xiong et al., 2024) form, and consistency across generations (Xiong et al., 2024; Yadkori et al., 2024) or models (Zhang et al., 2023; Feng et al., 2024). Given LLMs’ well-known over-confidence, recent works fine-tune LLMs to abstain on low-confidence facts using three training signals: token probability (Cohen et al., 2024; Chen et al., 2025a), generation consistency (Cheng et al., 2024; Yang et al., 2024b), and answer against external references (Grattafiori et al., 2024; Zhang et al., 2024; Kapoor et al., 2025; Li et al., 2025c).

CONFQA also uses supervised fine-tuning (SFT) based on answer correctness, but our empirical study yields two surprising findings that motivate suppressing hallucinations without materially sacrificing accuracy. (i) a dampener prompt matters in both training and inference, and (ii) training on atomic factual statements best teaches confidence and generalizes to unseen domains. Together these findings reduce hallucinations by an additional 5-11% and improve factuality by up to 25% over correctness-based SFT like (Zhang et al., 2024); we further find that adding a consistency constraint on top of correctness like (Cheng et al., 2024) causes an unnecessary accuracy regression (Section 5).

RAG triggering: RAG is widely studied and deployed to enhance LLMs’ factuality, with numerous surveys (Yu et al., 2022; Gao et al., 2024; Asai et al., 2024; Hu and Lu, 2024; Zhao et al., 2024; Fan et al., 2024; Huang and Huang, 2024; Zhou et al., 2024; Peng et al., 2024; Gupta et al., 2024; Ni et al., 2025). Modular RAG pipelines must optimize triggering (e.g., ETC (Li et al., 2025b)), retrieval (e.g., DeepRetrieval (Jiang et al., 2025)), and noise-robust answer generation (e.g., InstructRAG (Wei et al., 2025)); CONFRAG focuses on **when** to trigger RAG.

RAG triggering shares its core intuitions




Figure 2: Training data generation for CONFQA SFT.

with confidence-based abstention—withholding low-confidence facts in favor of external references—so related strategies overlap. *Token-level* approaches need internal generation confidence, triggering RAG when next token has low certainty (He et al., 2021; Jiang et al., 2023) or high entropy (Su et al., 2024; Li et al., 2025b); *fact-level* approaches trigger RAG on low fact-level confidence (Ni et al., 2024; Li et al., 2025a). CONF-RAG is a fact-level approach, but differs in two aspects: it uses a fine-tuned model (CONFQA) with low hallucination and only mild correction regression to best trade off quality and cost, and it parallelizes to hide the latency of model-based confidence testing.

Two other triggering lines: (1) iterative methods that interleave retrieval and generation until the retrieved content suffices (Peng et al., 2023; Trivedi et al., 2023; Jeong et al., 2024; Zhu et al., 2025); (2) R1-style generation that emits a <search> token to trigger retrieval, then continues generation over the appended results (Asai et al., 2023; Jin et al., 2025; Li et al., 2025d; Chen et al., 2025b). Neither targets the quality-latency trade-off of triggering. Our method is complementary and can improve their performance, either through training-data generation, or a curriculum-learning mechanism (Guo et al., 2018).

3 Methodology

We first describe CONFQA, a fine-tuning approach for calibrating model confidence, and then show how it supports the triggering strategy.

3.1 CONFQA: When to Say Unsure?

CONFQA fine-tunes an LLM to answer only questions that it is confident about and to admit “*I am unsure*” otherwise, guided by three intuitions: (1) calibrate the LLM’s confidence by showing the ground truth; (2) add a *dampener prompt*, “*Answer only if you are confident*”, to steer behavior explicitly; and (3) teach on *atomic facts* (entity attributes) so regularization targets factual statements.




Figure 3: CONF-RAG invocation architecture.

Training data are question–label pairs (see Figure 2). Questions ask for atomic facts and are generated from *DBPedia*, which spans diverse domains (Intuition #3), using open-source script of (Sun et al., 2024), evenly sampled across entity popularity for representativeness. For labels, we prompt the Llama-3.1-70B model to answer each question (Prompt 1 Appendix C) and Llama-3.1-405B to judge the answer’s consistency with the ground truth (Prompt 2 Appendix C): correct answers are labelled with the ground truth, otherwise with “*I am unsure about the answer*” (Intuition #1).

The dampener prompt is supplied as the system prompt in both training and inference to suppress hallucinations (Intuition #2). Appendix D.6 shows rephrasing it has little impact.

3.2 CONF-RAG: When to Trigger RAG?

Since CONFQA answers only when confident, we invoke RAG only when it says “unsure”.

Figure 3 depicts the CONF-RAG architecture, which extends *Factual Question Answering* to handle *dynamic* questions, those needing up-to-date information, best answered by RAG. These are identified either by fine-tuning the same model or by a smaller dedicated model (Appendix E). For *static* questions, CONFQA and RAG run in parallel: if the LLM produces a valid answer, RAG is early-stopped and that answer is returned; otherwise, the system waits for and returns the RAG output.

4 Experiment Setup

4.1 Benchmarks and Metrics

Data sets: We experimented with three *short-form factuality benchmarks*, where answers are mostly short. Questions range from *simple* questions about an entity’s attribute to *complex* ones requiring comparison, aggregation, reasoning, and post-processing. Table 1 summarizes the benchmarks and Appendix A gives details. Appendix B shows that on all three, well-known LLMs are consistently over-confident in their answers, while answer consistency correlates better with accuracy—

Benchmark	Category	Question types	# Domain	Size
Head-to-Tail (Sun et al., 2024)	short-form	simple questions (attribute of entities)	dbpedia, imdb	1,200
SimpleQA (Wei et al., 2024a)	short-form	general fact-seeking questions	multiple domains	4,326
CRAG (Yang et al., 2024a)	short-form	simple questions, reasoning questions	5 domains	642
LongFact (Wei et al., 2024b)	long-form	general questions	38 domains	250
AlpacaFact (Dubois et al., 2023)	long-form	fact-seeking instruction-following	multiple domains	241
Biography (Min et al., 2023)	long-form	biography questions	celebrity	183
MMLU 5-shot (Hendrycks et al., 2021)	general knowl.	multi-choice questions	57 domains	14,042
MMLU pro (Wang et al., 2024)	general knowl.	multi-choice questions	multiple domains	12,032

Table 1: The overall statistics of evaluation datasets.

matching prior observations (Mielke et al., 2022; Zhou et al., 2023; Xiong et al., 2024).

Metrics and evaluation: For *QA quality*, following CRAG (Yang et al., 2024a), we compute the percentage of *correct*, *incorrect* (hallucinations), and *missing* (not attempted) answers, and take *Factuality* = correct% - incorrect% as our main metric. Ranging from -1 to 1, Factuality penalizes hallucinations more than missing answers, letting us quantify hallucination mitigation as advocated in (Kalai et al., 2025); Appendix A further shows it penalizes hallucinations better than the usual *F-measure*. We judge answer correctness via prompt-based LLM-as-a-judge, which was shown to achieve 99% accuracy in (Yang et al., 2024a).

For *triggering*, we report precision and recall against an *oracle* that triggers RAG whenever the LLM’s answer from internal knowledge is not mostly correct (correct <2 out of 3 times). *Precision* is the fraction of triggered queries that the oracle would also trigger (i.e., how many of our RAG calls were necessary); *Recall* is the fraction of oracle-trigger queries that we actually triggered (i.e., how many necessary calls we caught). *F-measure* computes their harmonic mean: $F_{msr} = \frac{2 \cdot \text{prec-rec}}{\text{prec} + \text{rec}}$.

4.2 Models and Implementations

Implementations: We fine-tuned Llama-3.1-70B (Grattafiori et al., 2024). Guided by a scaling-law study (see Appendix D.1), we ran one epoch on 3K training samples, with learning rate 1e-6, batch size 1, and gradient accumulation step 1. Experiments used Nvidia H100 96GB HBM2e GPUs: 32 GPUs for fine-tuning, 8 for inference. We observed similar results on Qwen2.5-7B-Instruct (Qwen et al., 2025) and Gemma-3-4B-IT (Team et al., 2025), see Appendix D.5.

Alternative Solutions: We compared against two baselines: the oracle from Section 4.1, and Llama-3.1-70B with an inference-time dampener–

plus three state-of-the-art solutions:

Prompt-based (Ni et al., 2024; Li et al., 2025a) triggers RAG when the LLM signals low confidence. We implemented both the vanilla version, and an enhanced one (Prompt+) urging prudence and improving correctness (punish+explain in (Ni et al., 2024); Prompt 4 and 5 in Appendix C).

R-Tuning (Zhang et al., 2024) builds training data by appending “Are you sure you accurately answered the question based on your internal knowledge?” in the question, and padding the ground truth with “I am sure” or “I am unsure” by correctness. We trigger RAG on an “I am unsure” suffix. We trained on both MMLU (Hendrycks et al., 2021) and DBPedia.

IDK (Cheng et al., 2024) requires answer consistency (here at least four out of five times) as well as correctness when labeling ground truth. We fine-tuned on DBPedia for a more direct comparison.

5 Experimental Results

In this section, we study the effectiveness of CONFQA and CONFRAQ.

5.1 Effectiveness of CONFQA

The left part of Table 2 reports CONFQA’s answer quality. First, baseline methods are hallucination prone. The dampener cut incorrect answers by only 2.5-7.8 points across benchmarks.

Second, prompt-based methods suppress hallucinations, but hurt correctness, sometimes lowering factuality. The vanilla version still incurs high incorrectness (DBPedia, IMDb); while the stricter enhanced version suppresses hallucinations effectively, but sharply reduced correctness.

Third, R-tuning methods suppress hallucinations but hurt correctness. R-tuning on DBPedia suppresses hallucinations more effectively than on MMLU, but hurt correctness more. R-tuning on MMLU has high incorrectness rate—sometimes above baseline—for two reasons: 1) feeding ground

Model	Corr	Miss	Incor	Fac.	Ceiling Corr	Ceiling Fac.	Tri. Prec	Tri. Rec	F _{msr}
DBpedia (in-domain)									
Llama-3.1 Oracle	47.2	21.5	31.3	15.9	68.7	37.4	-	-	-
Llama-3.1 Dampen	47.0	26.8	26.2	20.8	73.8	47.6	<i>94.4</i>	47.9	63.3
Prompt	28.7	56.3	15.0	13.7	85.0	70.0	68.0	72.6	70.2
Prompt+	12.8	84.0	3.2	9.6	96.8	93.6	60.1	<i>95.6</i>	73.8
R-tuning (MMLU)	24.3	67.3	8.3	16.0	91.6	83.3	66.6	84.9	74.6
R-tuning (DBPedia)	24.5	67.8	7.7	16.8	92.3	84.6	67.3	86.4	75.7
IDK (DBPedia)	17.0	81.5	1.5	15.5	98.5	97.0	63.4	97.8	76.9
Llama-3.1 20-consis	-	-	-	-	-	-	74.6	90.9	81.9
CONFQA (no dampener)	49.0	33.5	17.5	31.5	82.5	65.0	95.5	60.6	74.1
CONFQA	31.5	63.3	5.2	26.3	94.8	89.6	73.9	88.6	<i>80.6</i>
IMDb (out-of-domain)									
Llama-3.1 Oracle	<i>41.0</i>	33.3	25.7	15.3	74.3	48.6	-	-	-
Llama-3.1 Dampen	40.7	36.2	23.2	17.5	76.9	53.7	97.7	59.9	74.3
Prompt	31.8	54.3	13.8	18.0	86.1	72.3	85.0	78.2	81.5
Prompt+	19.8	78.8	<i>1.3</i>	18.5	98.6	97.3	73.4	98.0	83.9
R-tuning (MMLU)	28.2	60.5	11.3	16.9	88.7	77.4	80.2	82.2	81.2
R-tuning (DBPedia)	25.3	70.2	4.5	20.8	95.5	91.0	80.3	95.5	87.2
IDK (DBPedia)	22.0	77.0	1.0	21.0	99.0	98.0	76.2	99.4	86.3
Llama-3.1 20-consis	-	-	-	-	-	-	87.8	95.2	<i>91.3</i>
CONFQA (no dampener)	43.0	42.0	16.0	27.0	85.0	69.0	<i>96.4</i>	68.1	79.8
CONFQA	32.5	63.3	4.2	28.3	95.8	91.6	89.2	95.8	92.4
SimpleQA (out-of-domain)									
Llama-3.1 Oracle	16.0	42.7	41.3	-25.3	58.7	17.4	-	-	-
Llama-3.1 Dampen	16.8	48.0	35.2	-18.4	64.8	29.6	97.7	55.8	71.0
Prompt	0.9	97.6	<i>1.5</i>	-0.6	98.5	97.0	84.3	97.9	90.6
Prompt+	1.5	96.8	1.7	-0.2	98.3	96.6	85.1	<i>98.0</i>	91.1
R-tuning (MMLU)	20.3	38.0	41.7	-21.4	58.3	16.6	87.2	88.4	87.8
R-tuning (DBPedia)	3.7	83.3	13.0	-9.3	87.0	74.0	85.4	84.7	85.1
IDK (DBPedia)	0.6	99.1	0.2	<i>0.4</i>	99.7	99.5	84.4	99.7	91.4
Llama-3.1 20-consis	-	-	-	-	-	-	88.9	95.2	<i>92.0</i>
CONFQA (no dampener)	<i>17.3</i>	55.8	26.8	-9.5	73.1	46.3	97.3	64.7	77.7
CONFQA	4.9	93.1	2.1	2.8	98.0	95.9	87.6	97.1	92.1
CRAG (out-of-domain)									
Llama-3.1 Oracle	60.3	11.7	28.0	32.3	72.0	44.0	-	-	-
Llama-3.1 Dampen	57.5	22.3	20.2	37.2	79.8	59.6	83.9	47.1	60.3
Prompt	23.7	69.3	7.0	16.7	93.0	86.0	48.1	83.9	61.1
Prompt+	26.5	71.2	2.3	24.2	97.7	95.4	53.2	95.3	68.3
R-tuning (MMLU)	57.8	<i>17.1</i>	25.1	32.7	74.9	49.8	51.2	72.9	60.2
R-tuning (DBPedia)	31.6	55.0	13.4	18.2	86.6	73.2	50.7	70.2	58.9
IDK (DBPedia)	20.7	78.2	1.1	19.6	98.9	97.8	49.8	98.0	66.1
Llama-3.1 20-consis	-	-	-	-	-	-	65.9	80.4	72.4
CONFQA (no dampener)	57.0	19.6	23.4	33.6	76.6	53.2	<i>83.3</i>	41.2	55.1
CONFQA	39.4	56.2	4.4	35.0	95.6	91.2	63.7	90.2	74.7

Table 2: Overall factuality and RAG triggering F_{msr} (all numbers in percentage %; bold indicates the best performance and italic the second best). CONFQA improves factuality and reduces incorrectness to 5% and below, achieving highest triggering F_{msr} on 3 benchmarks and close to highest triggering F_{msr} on DBPedia (Llama-3.1 20-consis). Note: Llama-3.1 20-consis calls LLM 21 times, with 20 response generations and 1 consistency judge. Not practical in real deployment.

truths as extra knowledge can encourage hallucinating what was not learned in pre-training; 2) MMLU mixes facts with reasoning, introducing ambiguity. This supports our hypothesis that atomic facts yield better training examples (Intuition #3).

Fourth, IDK achieves the lowest incorrectness rate (below 1.5% on all benchmarks) but drops cor-

rectness significantly, reducing overall factuality—showing that additionally requiring the consistency signal can be unnecessarily strict.

Fifth, fine-tuned CONFQA improves factuality by up to 25%, with incorrectness dropping to 5% or below on all benchmarks. Correctness drops too, but far less than hallucinations drop, yielding much

higher factuality. Although trained only on DBPedia, CONFQA generalizes to other benchmarks (Intuition #3).

5.2 CONFRAg Triggering and Serving

The right part of Table 2 compares triggering metrics across solutions. As a strong but impractical reference point, we include Llama-3.1 20-consis, which generates an answer 20 times per question and then calls Llama-3.1 once more to judge whether the 20 responses are mutually consistent; it triggers RAG unless all 20 responses agree and none is a refusal. Because this consistency signal is not an answer-quality decision, Llama-3.1 20-consis has no entries in the left (answer-quality) part of Table 2. It is also far costlier than CONFRAg— $\approx 20\times$ the generation tokens and at least twice the latency—which makes it unsuitable for deployment.

CONFRAg lifts potential correctness to $\approx 95\%$, while cutting unnecessary external retrievals by 5-39%. It attains the highest triggering F-measure on every benchmark except DBPedia, where Llama-3.1 20-consis is highest; even there, CONFRAg is within a small margin—so CONFRAg matches or exceeds Llama-3.1 20-consis on triggering quality at a fraction of its token and latency cost. The remaining methods fall short on one of two axes: Prompt and R-tuning have low ceiling correctness, while Prompt+ and IDK reach a higher ceiling but at much lower triggering precision, incurring unnecessary retrieval cost.

Deployment architecture: at serving time (Figure 3), LLM-QA and the RAG pipeline run in parallel; once CONFQA returns a confident answer, retrieval is early-stopped and the answer returned, otherwise the system waits for RAG. Because CONFQA abstains only on genuinely uncertain queries, CONFRAg finishes retrievals on a fraction of traffic, cutting unnecessary external retrievals by 5-39% (Table 2) and end-to-end latency by 600ms (P50) and 800ms (P90) on CRAG (Table 3). SimpleQA requires triggering RAG for most questions to reach high quality, so its latency gain is smaller. For a service handling N queries/day at retrieval unit cost c , this avoids $\approx (0.05 \text{ to } 0.39) \cdot N \cdot c$ in retrieval spend plus the corresponding generator GPU-seconds and tokens.

Table 3 shows results of production implementation of CONFRAg and RAG-everywhere for SimpleQA and CRAG benchmarks. CONFRAg shows

slight under-performance on SimpleQA factuality and Correctness and slight improvement on CRAG compared with always-on RAG. These numbers came from production implementation of a RAG system, so accounting for small variations, CONFRAg is more or less on-par with always-on RAG for quality. On our production-representative pipeline (Bing + Knowledge Graph APIs feeding Llama-3.1-70B), retrieval and summarization are the dominant per-query cost.

RAG cost saving: our production RAG pipeline runs four sequential stages: a search API call (500ms), crawling 10 pages (1300ms), content extraction and post-processing from HTML (400ms), and an LLM call over the extracted content to generate an answer (3000ms) – 5200ms end to end. When CONFRAg decides to answer directly, the search API is still triggered, but the pipeline halts during crawling, before post-processing and summarization. Because search runs in parallel with direct QA, its 500ms is fully overlapped and off the critical path; the three downstream stages are eliminated along with their cost: 1). Crawling (10 pages) – interrupted mid-process, avoiding most of the 1300ms latency and the associated bandwidth/crawl cost. 2). Extraction, chunking, and re-ranking – the full 400ms of post-processing and any paid API cost therein are saved. 3). LLM summarization – with no page content produced, both the summarization call and the tokens it would consume are avoided.

The token savings are substantial: summarizing 10 webpages to answer a question consumes roughly 15,000-50,000+ tokens, making the retrieval path 30-100 $\times$ more expensive per query than direct QA.

Fallback when RAG is unavailable: in offline or low-latency settings (on-device, Speech-in Speech-out), run CONFQA without the dampener to preserve correctness while suppressing hallucinations; add the dampener to drive hallucinations under 5% when maximum reliability is required.

5.3 Deep dive on CONFQA

Since the supervised model CONFQA is critical to deciding when to trigger RAG, we now examine its design choices and generalizability.

Ablation study: We compare CONFQA with several alternatives in Table 4 (full results in Table 7 Appendix D.2). First, dropping the dampener from

Model	SimpleQA							CRAG						
	Upper	Corr	Miss	Incor	Fac	L-P50	L-P90	Upper	Corr	Miss	Incor	Fac	L-P50	L-P90
LLM-only	20.0	20.0	44.1	35.9	-15.8	480	896	58.7	58.7	15.6	25.7	33.0	480	896
RAG-everywhere	100.0	78.1	11.5	10.5	67.6	1,900	2,780	100.0	61.1	15.1	23.8	37.3	1,900	2,780
CONFRAG	95.1	77.2	11.4	11.5	65.7	1,802	2,650	95.6	62.3	14.2	23.5	38.8	1,278	1,955

Table 3: CONFRAG achieves similar quality to RAG-everywhere, while cutting latency. **Upper**: upper bound of percentage of correct. **L-P50**: P50 latency and **L-P90**: P90 latency.

Model	DBPedia (in-domain)				IMDb (out-of-domain)				SimpleQA (out-of-domain)				CRAG (out-of-domain)			
	Corr	Miss	Incor	Fac	Corr	Miss	Incor	Fac	Corr	Miss	Incor	Fac	Corr	Miss	Incor	Fac
Llama-3.1 Dampen	47.0	26.8	26.2	20.8	40.7	36.2	23.2	17.5	16.8	48.0	35.2	-18.4	57.5	22.3	20.2	37.2
CONFQA	31.5	63.3	5.2	26.3	32.5	63.3	4.2	28.3	4.9	93.1	2.1	2.8	39.4	56.2	4.4	35.0
* No-dampener in Train	36.0	50.2	13.8	22.2	34.2	50.7	15.2	19.0	5.8	87.5	6.7	-0.9	46.0	44.2	9.8	36.2
* MMLU-as-source	8.2	89.8	2.0	6.2	16.7	82.0	1.3	15.4	0.6	98.8	0.5	0.1	7.0	92.7	0.3	6.7
* GT-as-label	48.0	2.8	49.2	-1.2	41.2	4.3	54.5	-13.3	17.8	13.7	68.4	-50.6	53.7	14.2	32.1	21.6

Table 4: Ablation study for CONFQA when applying dampener in inference for all models, showing effectiveness of our fine-tuned model. All numbers are in percentage (%) and full results in Table 7.

Model	Long Fact				Alpaca Fact				Biography			
	Prec	Rec	F1	Miss	Prec	Rec	F1	Miss	Prec	Rec	F1	Miss
Llama3.1	64.5	65.4	64.3	0	62.3	71.0	63.8	0	35.4	40.3	37.1	0
CONFQA	67.0	67.7	66.7	0.8	62.2	71.1	63.8	0.4	42.0	46.5	42.6	12.6

Table 5: CONFQA matches or exceeds baseline for long-form generation.

Model	MMLU (5-shot)	MMLU Pro
Llama3.1	82.7	66.3
CONFQA	82.8	65.4

Table 6: CONFQA does not regress on MMLU.

training raises correctness and lowers the missing rate, but increased hallucinations and thus lower factuality—confirming the dampener’s importance in training (Intuition #2). Second, generating training data from MMLU instead of DBPedia yields low hallucination (< 2%) but far lower correctness (Intuition #3). Third, using the ground truth directly as the label (*GT-as-label*) gives the highest correctness and lowest missing rate, but hallucinates much more (up to 68.4%), consistent with observations in previous work (Lin et al., 2024; Gekhman et al., 2024).

Generalizability to long-form benchmarks: We test *long-form factuality benchmarks*, where answers contain multiple factual statements (see Table 1 and Appendix A). We evaluate with *VeriScore* (Song et al., 2024), which computes precision, recall, and F1, using their fine-tuned models for claim extraction and verification; the minimum facts for perfect recall is set to the median extracted claims per dataset. We omit the dampener prompt to avoid regressing recall. Table 5 shows CONFQA matches or exceeds baseline precision and re-

call, except on 13% long-tail biography questions, which it declines to answer due to low confidence.

No regression on general benchmarks: Finally, we evaluated CONFQA on MMLU’s multi-choice questions on general knowledge and reasoning (Hendrycks et al., 2021; Wang et al., 2024), again without the dampener prompt. Table 6 shows scores mostly similar to baseline.

6 Conclusion

We presented CONFRAG, a confidence-guided RAG architecture that triggers RAG only when an LLM is unlikely to be correct. Built on CONFQA – a lightweight fine-tune that teaches a model to answer when confident and otherwise say “I am unsure” – it cuts hallucinations to 5% or below and matches RAG-everywhere accuracy (slight under-performance on SimpleQA and slight improvement on CRAG) while reducing retrievals by 5-39%.

CONFRAG is practical to deploy and maintain: CONFQA retrains in one epoch over 3K auto-generated knowledge-graph facts and transfers across model families (Appendix D.5), requiring no access to model internals or per-domain labeling. As base models and retrievers improve, integrating internal and external knowledge this way should further improve both factuality and latency.

7 Limitations

Our experiments focus on SFT, and we leave DPO-based fine-tuning for future work. We empirically compared DBPedia and MMLU, where the former contains only simple factual questions, and the latter contains questions ranging from factual to reasoning. A comprehensive study regarding the effect of sources with different coverage in this spectrum can further improve understanding. SFT requires access to LLM parameters, thus restricting the application of the proposed framework to open-sourced models. Lastly, our cost/latency results come from a replay of the benchmark on production pipeline rather than a live A/B test, and savings depend on traffic mix and retrieval pricing.

References

- 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.
- Akari Asai, Zexuan Zhong, Danqi Chen, Pang Wei Koh, Luke Zettlemoyer, Hannaneh Hajishirzi, and Wen tau Yih. 2024. [Reliable, adaptable, and attributable language models with retrieval](#). *Preprint*, arXiv:2403.03187.
- Lida Chen, Zujie Liang, Xintao Wang, Jiaqing Liang, Yanghua Xiao, Feng Wei, Jinglei Chen, Zhenghong Hao, Bing Han, and Wei Wang. 2025a. [Teaching large language models to express knowledge boundary from their own signals](#). In *Proceedings of the 3rd Workshop on Towards Knowledgeable Foundation Models (KnowFM)*, pages 26–39, Vienna, Austria. Association for Computational Linguistics.
- Mingyang Chen, Linzhuang Sun, Tianpeng Li, Haoze Sun, Yijie Zhou, Chenzheng Zhu, Haofen Wang, Jeff Z. Pan, Wen Zhang, Huajun Chen, Fan Yang, Zenan Zhou, and Weipeng Chen. 2025b. [Research: Learning to reason with search for llms via reinforcement learning](#). *Preprint*, arXiv:2503.19470.
- Qinyuan Cheng, Tianxiang Sun, Xiangyang Liu, Wenwei Zhang, Zhangyue Yin, Shimin Li, Linyang Li, Zhengfu He, Kai Chen, and Xipeng Qiu. 2024. [Can ai assistants know what they don’t know?](#) *Preprint*, arXiv:2401.13275.
- Yung-Sung Chuang, Yujia Xie, Hongyin Luo, Yoon Kim, James Glass, and Pengcheng He. 2024. [Dola: Decoding by contrasting layers improves factuality in large language models](#). *Preprint*, arXiv:2309.03883.
- Roi Cohen, Konstantin Dobler, Eden Biran, and Gerard de Melo. 2024. [I don’t know: Explicit modeling of uncertainty with an \[idk\] token](#). *Preprint*, arXiv:2412.06676.
- 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.
- Jinhao Duan, Hao Cheng, Shiqi Wang, Alex Zavalny, Chenan Wang, Renjing Xu, Bhavya Kailkhura, and Kaidi Xu. 2024. [Shifting attention to relevance: Towards the predictive uncertainty quantification of free-form large language models](#). *Preprint*, arXiv:2307.01379.
- Yann Dubois, Chen Xuechen Li, Rohan Taori, Tianyi Zhang, Ishaan Gulrajani, Jimmy Ba, Carlos Guestrin, Percy S Liang, and Tatsunori B Hashimoto. 2023. [Alpacafarm: A simulation framework for methods that learn from human feedback](#). In *Advances in Neural Information Processing Systems*, volume 36, pages 30039–30069. Curran Associates, Inc.
- Nouha Dziri, Sivan Milton, Mo Yu, Osmar Zaiane, and Siva Reddy. 2022. [On the origin of hallucinations in conversational models: Is it the datasets or the models?](#) In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5271–5285, Seattle, United States. Association for Computational Linguistics.
- Wenqi Fan, Yujian Ding, Liangbo Ning, Shijie Wang, Hengyun Li, Dawei Yin, Tat-Seng Chua, and Qing Li. 2024. [A survey on rag meeting llms: Towards retrieval-augmented large language models](#). *Preprint*, arXiv:2405.06211.
- Shangbin Feng, Weijia Shi, Yike Wang, Wenxuan Ding, Vidhisha Balachandran, and Yulia Tsvetkov. 2024. [Don’t hallucinate, abstain: Identifying llm knowledge gaps via multi-llm collaboration](#). *Preprint*, arXiv:2402.00367.
- Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, Meng Wang, and Haofen Wang. 2024. [Retrieval-augmented generation for large language models: A survey](#). *Preprint*, arXiv:2312.10997.
- 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.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, and 542 others. 2024. [The llama 3 herd of models](#). *Preprint*, arXiv:2407.21783.
- Sheng Guo, Weilin Huang, Haozhi Zhang, Chenfan Zhuang, Dengke Dong, Matthew R. Scott, and Dinglong Huang. 2018. Curriculumnet: Weakly supervised learning from large-scale web images. In *ECCV*.

- Shailja Gupta, Rajesh Ranjan, and Surya Narayan Singh. 2024. [A comprehensive survey of retrieval-augmented generation \(rag\): Evolution, current landscape, and future directions](#). *Preprint*, arXiv:2410.12837.
- Junxian He, Graham Neubig, and Taylor Berg-Kirkpatrick. 2021. Efficient nearest neighbor language models. In *EMNLP*.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. [Measuring massive multitask language understanding](#). *Preprint*, arXiv:2009.03300.
- Yucheng Hu and Yuxing Lu. 2024. [Rag and rau: A survey on retrieval-augmented language model in natural language processing](#). *Preprint*, arXiv:2404.19543.
- Yizheng Huang and Jimmy Huang. 2024. [A survey on retrieval-augmented text generation for large language models](#). *Preprint*, arXiv:2404.10981.
- Soyeong Jeong, Jinheon Baek, Sukmin Cho, Sung Ju Hwang, and Jong C. Park. 2024. [Adaptive-rag: Learning to adapt retrieval-augmented large language models through question complexity](#). *Preprint*, arXiv:2403.14403.
- Pengcheng Jiang, Jiacheng Lin, Lang Cao, Runchu Tian, SeongKu Kang, Zifeng Wang, Jimeng Sun, and Jiawei Han. 2025. Deepretrieval: Hacking real search engines and retrievers with llms via reinforcement learning. In *COLM*.
- Zhengbao Jiang, Frank F. Xu, Luyu Gao, Zhiqing Sun, Qian Liu, Jane Dwivedi-Yu, Yiming Yang, Jamie Callan, and Graham Neubig. 2023. [Active retrieval augmented generation](#). *Preprint*, arXiv:2305.06983.
- Bowen Jin, hansi Zeng, Zhenrui Yue, Jinsung Yoon, Serkan O. Arik, Dong Wang, Hamed Azmani, and Jiawei Han. 2025. Search-r1: Training llms to reason and leverage search engines with reinforcement learning. In *COLM*.
- Saurav Kadavath, Tom Conerly, Amanda Askell, Tom Henighan, Dawn Drain, Ethan Perez, Nicholas Schiefer, Zac Hatfield-Dodds, Nova DasSarma, Eli Tran-Johnson, Scott Johnston, Sheer El-Showk, Andy Jones, Nelson Elhage, Tristan Hume, Anna Chen, Yuntao Bai, Sam Bowman, Stanislav Fort, and 17 others. 2022. [Language models \(mostly\) know what they know](#).
- Adam Tauman Kalai, Ofir Nachum, Santosh S. Vempala, and Edwin Zhang. 2025. [Why lms hallucinate?](#) *Preprint*, arXiv:2509.04664.
- Sanyam Kapoor, Nate Gruver, Manley Roberts, Katherine Collins, Arka Pal, Umang Bhatt, Adrian Weller, Samuel Dooley, Micah Goldblum, and Andrew Gordon Wilson. 2025. [Large language models must be taught to know what they don’t know](#). *Preprint*, arXiv:2406.08391.
- Jiyeon Kim, Hyunji Lee, Hyowon Cho, Joel Jang, Hyeonbin Hwang, Seungpil Won, Youbin Ahn, Do-haeng Lee, and Minjoon Seo. 2025. Knowledge entropy decay during lm pre-training hinders new knowledge acquisition. In *ICLR*.
- Lorenz Kuhn, Yarin Gal, and Sebastian Farquhar. 2023. [Semantic uncertainty: Linguistic invariances for uncertainty estimation in natural language generation](#). *Preprint*, arXiv:2302.09664.
- Nathan Lambert, Jacob Morrison, Valentina Pyatkin, Shengyi Huang, Hamish Ivison, Faeze Brahman, Lester James V. Miranda, Alisa Liu, Nouha Dziri, Shane Lyu, Yuling Gu, Saumya Malik, Victoria Graf, Jena D. Hwang, Jiangjiang Yang, Ronan Le Bras, Oyvind Tafjord, Chris Wilhelm, Luca Soldaini, and 4 others. 2025. [Tulu 3: Pushing frontiers in open language model post-training](#). *Preprint*, arXiv:2411.15124.
- Deren Lei, Yaxi Li, Mengya Hu, Mingyu Wang, Vincent Yun, Emily Ching, and Eslam Kamal. 2023. [Chain of natural language inference for reducing large language model ungrounded hallucinations](#). *Preprint*, arXiv:2310.03951.
- Bo Li, Tian Tian, Zhenghua Xu, Hao Cheng, Shikun Zhang, and Wei Ye. 2025a. Investigating the factual knowledge boundary of llms with retrieval augmentation. In *COLING*.
- Bo Li, Tian Tian, Zhenghua Xu, Hao Cheng, Shikun Zhang, and Wei Ye. 2025b. [Modeling uncertainty trends for timely retrieval in dynamic rag](#). *Preprint*, arXiv:2511.09980.
- Jiaqi Li, Yixuan Tang, and Yi Yang. 2025c. Know the unknown: An uncertainty-sensitive method for llm instruction tuning. In *ACL*.
- Kenneth Li, Oam Patel, Fernanda Viegas, Hanspeter Pfister, and Martin Wattenberg. 2023. Inference-time intervention: Eliciting truthful answers from a language model. In *NeurIPS*.
- Xiaoxi Li, Guanting Dong, Jiajie Jin, Yuyao Zhang, Yujia Zhou, Yutao Zhu, Peitian Zhang, and Zhicheng Dou. 2025d. [Search-o1: Agentic search-enhanced large reasoning models](#). *Preprint*, arXiv:2501.05366.
- Sheng-Chieh Lin, Luyu Gao, Barlas Oguz, Wenhan Xiong, Jimmy Lin, Wen-tau Yih, and Xilun Chen. 2024. [Flame : Factuality-aware alignment for large language models](#). *Preprint*, arXiv:2405.01525.
- Stephanie Lin, Jacob Hilton, and Owain Evans. 2022. [Teaching models to express their uncertainty in words](#). *Preprint*, arXiv:2205.14334.
- Sabrina J. Mielke, Arthur Szlam, Emily Dinan, and Y-Lan Boureau. 2022. [Reducing conversational agents’ overconfidence through linguistic calibration](#). *Transactions of the Association for Computational Linguistics*, 10:857–872.

- Sewon Min, Kalpesh Krishna, Xinxu Lyu, Mike Lewis, Wen tau Yih, Pang Wei Koh, Mohit Iyyer, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2023. [Factscore: Fine-grained atomic evaluation of factual precision in long form text generation](#). *Preprint*, arXiv:2305.14251.
- Abhika Mishra, Akari Asai, Vidhisha Balachandran, Yizhong Wang, Graham Neubig, Yulia Tsvetkov, and Hannaneh Hajishirzi. 2024. Fine-grained hallucination detection and editing for language models. In *COLM*.
- Tu Anh Nguyen, Eugene Kharitonov, Jade Copet, Yossi Adi, Wei-Ning Hsu, Ali Elkahky, Paden Tomasello, Robin Algayres, Benoit Sagot, Abdelrahman Mohamed, and Emmanuel Dupoux. 2022. [Generative spoken dialogue language modeling](#). *Preprint*, arXiv:2203.16502.
- Bo Ni, Zheyuan Liu, Leyao Wang, Yongjia Lei, Yuying Zhao, Xueqi Cheng, Qingkai Zeng, Xin Luna Dong, Yinglong Xia, Krishnamurthy Kenthapadi, Ryan Rossi, Franck Dernoncourt, Md Mehrab Tanjim, Neseen Ahmed, Xiaorui Liu, Wenqi Fan, Erik Blasch, Yu Wang, Meng Jiang, and Tyler Derr. 2025. [Towards trustworthy retrieval augmented generation for large language models: A survey](#). *Preprint*, arXiv:2502.06872.
- Shiyu Ni, Keping Bi, Jiafeng Guo, and Xueqi Cheng. 2024. [When do LLMs need retrieval augmentation? mitigating LLMs’ overconfidence helps retrieval augmentation](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 11375–11388, Bangkok, Thailand. Association for Computational Linguistics.
- Baolin Peng, Michel Galley, Pengcheng He, Hao Cheng, Yujia Xie, Yu Hu, Qiuyuan Huang, Lars Liden, Zhou Yu, Weizhu Chen, and Jianfeng Gao. 2023. [Check your facts and try again: Improving large language models with external knowledge and automated feedback](#). *Preprint*, arXiv:2302.12813.
- Boci Peng, Yun Zhu, Yongchao Liu, Xiaohe Bo, Haizhou Shi, Chuntao Hong, Yan Zhang, and Siliang Tang. 2024. [Graph retrieval-augmented generation: A survey](#). *Preprint*, arXiv:2408.08921.
- Qwen, :, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jingren Zhou, and 25 others. 2025. [Qwen2.5 technical report](#). *Preprint*, arXiv:2412.15115.
- Vipula Rawte, Swagata Chakraborty, Agnibh Pathak, Anubhav Sarkar, S.M Towhidul Islam Tonmoy, Aman Chadha, Amit Sheth, and Amitava Das. 2023. The troubling emergence of hallucination in large language models - an extensive definition, quantification, and prescriptive remediations. In *EMNLP*.
- Yixiao Song, Yekyung Kim, and Mohit Iyyer. 2024. [VeriScore: Evaluating the factuality of verifiable claims in long-form text generation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 9447–9474, Miami, Florida, USA. Association for Computational Linguistics.
- Weihang Su, Yichen Tang, Qingyao Ai, Zhijing Wu, and Yiqun Liu. 2024. [Dragin: Dynamic retrieval augmented generation based on the information needs of large language models](#). *Preprint*, arXiv:2403.10081.
- Kai Sun, Yifan Ethan Xu, Hanwen Zha, Yue Liu, and Xin Luna Dong. 2024. Head-to-tail: How knowledgeable are large language models (llms)? aka will llms replace knowledge graphs? In *NAACL-HLT*.
- Zhiqing Sun, Xuezhi Wang, Yi Tay, Yiming Yang, and Denny Zhou. 2023. [Recitation-augmented language models](#). *Preprint*, arXiv:2210.01296.
- Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, Louis Rouillard, Thomas Mesnard, Geoffrey Cideron, Jean bastien Grill, Sabela Ramos, Edouard Yvinec, Michelle Casbon, Etienne Pot, Ivo Penchev, and 197 others. 2025. [Gemma 3 technical report](#). *Preprint*, arXiv:2503.19786.
- Katherine Tian, Eric Mitchell, Allan Zhou, Archit Sharma, Rafael Rafailov, Huaxiu Yao, Chelsea Finn, and Christopher Manning. 2023. [Just ask for calibration: Strategies for eliciting calibrated confidence scores from language models fine-tuned with human feedback](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 5433–5442, Singapore. Association for Computational Linguistics.
- S. M Towhidul Islam Tonmoy, S M Mehedi Zaman, Vinija Jain, Anku Rani, Vipula Rawte, Aman Chadha, and Amitava Das. 2024. [A comprehensive survey of hallucination mitigation techniques in large language models](#). *Preprint*, arXiv:2401.01313.
- Harsh Trivedi, Niranjan Balasubramanian, Tushar Khot, and Ashish Sabharwal. 2023. [Interleaving retrieval with chain-of-thought reasoning for knowledge-intensive multi-step questions](#). *Preprint*, arXiv:2212.10509.
- Neeraj Varshney, Wenlin Yao, Hongming Zhang, Jian-shu Chen, and Dong Yu. 2023. [A stitch in time saves nine: Detecting and mitigating hallucinations of llms by validating low-confidence generation](#). *Preprint*, arXiv:2307.03987.
- Yubo Wang, Xueguang Ma, Ge Zhang, Yuansheng Ni, Abhranil Chandra, Shiguang Guo, Weiming Ren, Aaran Arulraj, Xuan He, Ziyang Jiang, Tianle Li, Max Ku, Kai Wang, Alex Zhuang, Rongqi Fan, Xiang Yue, and Wenhui Chen. 2024. [Mmlu-pro: A more robust and challenging multi-task language understanding benchmark](#). *Preprint*, arXiv:2406.01574.

- Jason Wei, Nguyen Karina, Hyung Won Chung, Yunxin Joy Jiao, Spencer Papay, Amelia Glaese, John Schulman, and William Fedus. 2024a. [Measuring short-form factuality in large language models](#). *Preprint*, arXiv:2411.04368.
- Jerry Wei, Chengrun Yang, Xinying Song, Yifeng Lu, Nathan Hu, Jie Huang, Dustin Tran, Daiyi Peng, Ruibo Liu, Da Huang, Cosmo Du, and Quoc V. Le. 2024b. [Long-form factuality in large language models](#). In *Advances in Neural Information Processing Systems*, volume 37, pages 80756–80827. Curran Associates, Inc.
- Zhepei Wei, Wei lin Chen, and Yu Meng. 2025. [Instruc-trag: Instructing retrieval augmented generation via self-synthesized rationales](#). In *ICLR*.
- Yiqing Xie, Wenxuan Zhou, Pradyot Prakash, Di Jin, Yuning Mao, Quintin Fettes, Arya Talebzadeh, Sinong Wang, Han Fang, Carolyn Rose, Daniel Fried, and Hejia Zhang. 2025. [Improving model factuality with fine-grained critique-based evaluator](#). *Preprint*, arXiv:2410.18359.
- Zhifei Xie and Changqiao Wu. 2024. [Mini-omni: Language models can hear, talk while thinking in streaming](#). *Preprint*, arXiv:2408.16725.
- Miao Xiong, Zhiyuan Hu, Xinyang Lu, YIFEI LI, Jie Fu, Junxian He, and Bryan Hooi. 2024. [Can LLMs express their uncertainty? an empirical evaluation of confidence elicitation in LLMs](#). In *The Twelfth International Conference on Learning Representations*.
- Yasin Abbasi Yadkori, Ilja Kuzborskij, David Stutz, András György, Adam Fisch, Arnaud Doucet, Iuliya Beloshapka, Wei-Hung Weng, Yao-Yuan Yang, Csaba Szepesvári, Ali Taylan Cemgil, and Nenad Tomasev. 2024. [Mitigating llm hallucinations via conformal abstention](#). *Preprint*, arXiv:2405.01563.
- Xiao Yang, Kai Sun, Hao Xin, Yushi Sun, Nikita Bhalla, Xiangsen Chen, Sajal Choudhary, Rongze Daniel Gui, Ziran Will Jiang, Ziyu Jiang, Lingkun Kong, Brian Moran, Jiaqi Wang, Yifan Ethan Xu, An Yan, Chenyu Yang, Eting Yuan, Hanwen Zha, Nan Tang, and 8 others. 2024a. [Crag-comprehensive rag benchmark](#). *Preprint*, arXiv:2406.04744.
- Yuqing Yang, Ethan Chern, Xipeng Qiu, Graham Neubig, and Pengfei Liu. 2024b. [Alignment for honesty](#). *Preprint*, arXiv:2312.07000.
- Fan Yin, Jayanth Srinivasa, and Kai-Wei Chang. 2024. [Characterizing truthfulness in large language model generations with local intrinsic dimension](#). *Preprint*, arXiv:2402.18048.
- Wenhao Yu, Chenguang Zhu, Zaitang Li, Zhiting Hu, Qingyun Wang, Heng Ji, and Meng Jiang. 2022. [A survey of knowledge-enhanced text generation](#). *ACM Computing Surveys*, 54(11s):1–38.
- Hanning Zhang, Shizhe Diao, Yong Lin, Yi R. Fung, Qing Lian, Xingyao Wang, Yangyi Chen, Heng Ji, and Tong Zhang. 2024. [R-tuning: Instructing large language models to say ‘i don’t know’](#). *Preprint*, arXiv:2311.09677.
- Jiaxin Zhang, Zhuohang Li, Kamalika Das, Bradley Malin, and Sricharan Kumar. 2023. [SAC³: Reliable hallucination detection in black-box language models via semantic-aware cross-check consistency](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 15445–15458, Singapore. Association for Computational Linguistics.
- Penghao Zhao, Hailin Zhang, Qinhan Yu, Zhengren Wang, Yunteng Geng, Fangcheng Fu, Ling Yang, Wentao Zhang, Jie Jiang, and Bin Cui. 2024. [Retrieval-augmented generation for ai-generated content: A survey](#). *Preprint*, arXiv:2402.19473.
- Chunting Zhou, Graham Neubig, Jiatao Gu, Mona Diab, Paco Guzman, Luke Zettlemoyer, and Marjan Ghazvininejad. 2021. [Detecting hallucinated content in conditional neural sequence generation](#). *Preprint*, arXiv:2011.02593.
- Kaitlyn Zhou, Dan Jurafsky, and Tatsunori Hashimoto. 2023. Navigating the grey area: How expressions of uncertainty and overconfidence affect language models. In *EMNLP*.
- Yujia Zhou, Yan Liu, Xiaoxi Li, Jiajie Jin, Hongjin Qian, Zheng Liu, Chaozhuo Li, Zhicheng Dou, Tsung-Yi Ho, and Philip S. Yu. 2024. [Trustworthiness in retrieval-augmented generation systems: A survey](#). *Preprint*, arXiv:2409.10102.
- Yijie Zhu, Haojie Zhou, Wanting Hong, Tailin Liu, and Ning Wang. 2025. [Reap: Enhancing rag with recursive evaluation and adaptive planning for multi-hop question answering](#). *Preprint*, arXiv:2511.09966.

A Benchmark Setup Details

To evaluate the performance of the fine-tuning, We consider 8 benchmarks in the main content as described in Section 4.1 Table 1, with details described here.

- **Head-to-Tail (*short-form with simple questions*)** (Sun et al., 2024): We leveraged the data scraping scripts from the **Head-to-Tail** benchmark, and generated simple questions and their answers according to content from DBPedia ¹ (general knowledge graph) and IMDb ² (data in the *Movie* domain). From each dataset we randomly sampled 200 entities for *head* entities, 200 for *torso* entities, and 200 for *tail* entities. Here we follow the definition in (Sun et al., 2024) for head, torso and tail: we rank all entities by their traffic; head entities are top-popular entities that together account for 1/3 of traffic, tail entities are unpopular entities that together account for 1/3 of traffic, and torso entities are the remaining medium-popular entities. Together, we have 1200 question-answer pairs, 600 from each source.
- **SimpleQA (*short-form with simple questions*)** (Wei et al., 2024a): **SimpleQA** is a benchmark released by *OpenAI* to measure LLM factuality. It contains 4326 manually crafted short, fact-seeking questions, covering diverse topics such as science, technology, history, and entertainment.
- **CRAG (*short-form with simple and complex questions*)** (Yang et al., 2024a): CRAG is a benchmark to test RAG capabilities. It contains 4,409 training and 1,335 evaluation questions covering five domains (general, finance, sports, music, movie), entities of different popularities (head, torso, tail), facts of different dynamisms (static, slow-changing, fast-changing, real-time), and eight question types (simple, condition, set, comparison, aggregation, multi-hop, post-processing, false premise). We selected the 642 *static* questions from the evaluation data set, with 97 questions for head entities, 99 for torso, 90 for tail entities and 356 for facts from the web (mostly popular); we excluded false-premise

and dynamic questions from the sampling as it presents different challenges.

- **LongFact (*long-form*)** (Wei et al., 2024b): Aiming to measure the factuality of long-form responses consisting of at least several paragraphs, LongFact has 2,280 factual questions covering 38 topics, generated by prompting GPT-4. Following Wei et al. (2024b), we use the 250 prompts from the LongFact-Objects dataset in our experiments.
- **AlpacaFact (*long-form*)** (Lin et al., 2024): Initially sourced from diverse interactions with real-world users, the 805 instructions in AlpacaFarm (Dubois et al., 2023) served as a benchmark for evaluating the ability of different LLMs to follow instructions. Following Lin et al. (2024), we used a subset of 241 fact-seeking instructions in this work.
- **Biography (*long-form*)** (Min et al., 2023): To validate the effectiveness of FActScore, Min et al. (2023) created a collection of prompts named Biography by applying the template “Tell me a bio of [Person Name]” to 183 notable individuals listed on Wikipedia. Given its extensive use in recent literature, we have included this prompt set for our experiments as well.
- **MMLU (*General knowledge*)**: The MMLU (Hendrycks et al., 2021) dataset covers 57 subjects, including areas such as mathematics, history, law, and medicine. It contains two subsets: the MMLU 5-shot dataset contains 14,042 multi-choice questions to evaluate general knowledge and problem-solving tasks; the MMLU Pro (Wang et al., 2024) dataset contains 12,032 multi-choice questions to stress-test reasoning, disambiguation, and factual accuracy.

For short-form questions we consider factuality score as defined in Section 4.1 rather than F1 score, where F1-score is more lenient for incorrect answers (hallucinations), but factuality strongly prefers missing answers to hallucinations. For example, consider a model that answers 10% questions correctly (correct% = 10%) and the rest of the questions incorrectly (incorrect% = 90%); the F1-score is 10% (not punishing hallucinations much) while the factuality is -80%. Now consider another model that answers 10% questions correctly

¹dbpedia.org

²imdb.com

and admits "I am unsure about the answer" for the rest of the questions; the F1-score is 18.2%, only slightly higher than 10%, but the factuality is 10%, significantly higher than -80%. Finally, consider a set of binary questions. A model that does not know any answer will get factuality of 0% no matter whether it admits so or makes random guesses ($0 = 50\% - 50\%$); however, if we use F-measure, the former gets a F-measure of 0 and the latter gets a F-measure of 50%, encouraging random guesses and thus hallucinations.

For long-form responses we use the automatic evaluation metric, *VeriScore* (Song et al., 2024), for measuring the factuality (results in Section 5.3 Table 5). Following FActScore (Min et al., 2023) and SAFE (Wei et al., 2024b), VeriScore extracts more sensible and verifiable claims from each sentence and uses Google search snippet instead of Wikipedia as the source of knowledge. This approach allows VeriScore to be applied to more diverse topics and requires fewer but more meaningful claims to be checked. We report the F1 score from VeriScore, which represents the harmonic mean of precision and recall. In line with (Song et al., 2024), we set the minimum number of facts required for perfect recall based on the median number of extracted claims per dataset, using their fine-tuned models for claim extraction and verification.

B Analysis of LLM’s Self-Reported Confidence

Using three benchmarks, we reproduce the observations in (Wei et al., 2024a; Xiong et al., 2024), that LLMs’ self-reported confidence is systematically overestimated, making it a poor signal for deciding when to trigger RAG.

We prompted the LLM to output a confidence score between 0 and 1 along with its answer (Prompt 3 in Appendix C). After removing missing answers, we binned the reported confidences into equal-sized quantiles, and plotted average accuracy per bin to assess *calibration*: whether a confidence of 0.8 corresponds to 80% accuracy? Figure 4 shows the calibration, leading to three observations.

1. The self-reported confidence is mostly positively correlated with QA accuracy, but LLMs tend to be over-confident (the correlation curves are below the ideal calibration dashed line). For example, when Llama-3.1-70B pre-

dicts a confidence of 80% on CRAG, the real accuracy is only 33%.

2. Notably, for the same model series, the smaller model is often more confident than the larger model (with an exception of Claude3.5 on CRAG), demonstrating the interesting correlation between ignorance and self-assurance.
3. Finally, the overconfidence is more pronounced when answering SimpleQA questions than on other benchmarks. A sample of 50 questions from SimpleQA shows that the questions are often nuanced for fairly popular entities (e.g. “*What was the first line after the salutation in the letter sent to Sardar Patel by Abhay Charan De?*”, “*In which month and year was Service Pack 3 for Windows Fundamentals for Legacy PCs released?*”), possibly causing LLMs to be over-confident.

Influence of entity popularity on confidence:

Figure 5 shows the calibration on the Head-to-Tail and CRAG benchmarks, where questions are categorized by entity popularity into *Head*, *Torso*, *Tail* (plus *Web* for CRAG).

Interestingly, we found for simple questions on Head-to-Tail, models are better calibrated for head entities than torso or tail entities (Figure 5 bottom panels). However, on more complex questions on CRAG, models are better calibrated for tail entities than torso or head entities (Figure 5 top panels). This shows two different dimensions that can affect the model confidence: entity popularity and question nuances.

Consistency vs. Accuracy: To measure consistency, we ask LLM the same question 20 times with the temperature set to 1.0, select the most frequent response as the final answer, and calculate its frequency among the 20 times as the consistency score. To be robust against minor differences, we determine the “most frequent” answer based on semantic similarity rather than exact string match.

Figure 6 shows that consistency is mostly better calibrated than self-reported confidence, and largely overlays with the ideal calibration on CRAG. In addition, the calibration curve is more linear compared to self-reported confidence in Figure 4.




Figure 4: LLMs’ self-reported confidences are correlated with QA accuracy, but often over-confident.




Figure 5: Correlation between LLM’s self-reported confidences and average accuracies on the CRAG dataset and the Head-to-Tail dataset, categorized by question types.

C Prompts

We include a list of prompts used in experiments for reference. Each prompt is cited in the main paper where it is applied. For clarity, we also provide brief descriptions of each prompt here.

Prompt 1: Simple question answer generation prompt

As Assistant AI, you help answer factual questions. Please keep your responses short and concise and directly provide the answer to the user question without reasoning. Answer only if you are confident; otherwise, respond with ‘I am unsure about the answer’.

Prompt 1 is used in several places:

- *Training-data generation*: to prompt Llama-3.1-70B to answer the 3K DBPedia-based questions.
- *SFT*: as the system prompt, unless we state no prompt was used in training.

- *Benchmarks*: as the system prompt at inference, unless we state no prompt was used.

Prompt 2 is used in two places:

- Training-data generation: to grade whether the model’s answer matches the ground truth.
- Benchmarks: to grade model responses when evaluating different models.

Prompt 3 is used in Appendix B to make the model output a confidence score alongside its answer, studying how self-reported confidence and answer consistency each relate to answer accuracy in Appendix B.

Prompt 4 and 5 are from Paper (Ni et al., 2024), and are used to generate results for row Prompt and Prompt+ in Table 2.




Figure 6: LLM’s answer consistency is often better calibrated with QA accuracy than self-reported confidences.

Prompt 2: Answer-grading prompt. Used in Section 3.1 to generate training labels: it grades whether the model’s answer matches the ground truth. If so, the ground truth becomes the label; otherwise, “*I am unsure about the answer*” is used.

You need to check whether the prediction of a question-answering system to a question is correct. You should make the judgment based on a list of ground truth answers provided to you. Your response should be “correct” if the prediction is correct or “incorrect” if the prediction is wrong.

Example 1:

Question: Who authored The Taming of the Shrew (published in 2002)? Ground truth: ["William Shakespeare", "Roma Gill"]
Prediction: W Shakespeare
Correctness: correct

Example 2:

Question: Who authored The Taming of the Shrew (published in 2002)? Ground truth: ["William Shakespeare", "Roma Gill"]
Prediction: Roma Gill and W Shakespeare
Correctness: correct

Example 3:

Question: Who authored The Taming of the Shrew (published in 2002)? Ground truth: ["William Shakespeare", "Roma Gill"]
Prediction: Roma Shakespeare
Correctness: incorrect

Example 4:

Question: What country is Maharashtra Metro Rail Corporation Limited located in?
Ground truth: ["India"]
Prediction: Maharashtra
Correctness: incorrect

Example 5:

Question: What’s the job of Song Kang-ho in Parasite (2019)? Ground truth: ["actor"]
Prediction: He plays the role of Kim Ki-taek, the patriarch of the Kim family.
Correctness: correct

Prompt 3: Self-reported confidence prompt.

Please answer the following question truthfully to your best knowledge. Provide a confidence score between 0 to 1 representing your confidence in the answer.
Question: {{{question}}}
Your response must be in the following JSON format.
{"answer": "Your answer here", "confidence_score": float number}
Remember, your answer MUST be a valid json string with two keys: "answer", "confidence_score".

Prompt 4: Vanilla Prompt from Paper (Ni et al., 2024).

Answer the following question based on your internal knowledge with one or few words. If you are sure the answer is accurate and correct, please say "certain" after the answer. If you are not confident with the answer, please say "uncertain".

Prompt 5: Punish+Explain Prompt from Paper (Ni et al., 2024).

Answer the following question based on your internal knowledge with one or few words and explain why you give this answer. If you are sure the answer is accurate and correct, please say "certain" after the answer. If you are not confident with the answer, please say "uncertain". You will be punished if the answer is not right but you say "certain".

D CONFQA experiment and full ablation study

We discuss additional details for CONFQA experiment and include the full ablation study as a complementary document to the main paper.

D.1 Fine-tuning implementation

To determine how many data samples and fine-tuning steps are needed for optimal performance, we run a simple scaling-law study.

We prepared 27K DBpedia question-answer pairs (9k each for head, torso and tail entities) and ran 10K steps on 4 hosts, 8 processes per host, with batch size 1. Around 100 steps gives the best performance, with more causing over-fitting. At this setting, 100 steps is one epoch for 3200 samples. We thus selected 3K high quality instances (1K each for head/torso/tail) and fine-tuned for one epoch. The 1K/1K/1K split follows the real-world distribution of entity popularity by definition (see the Head-to-Tail paper (Sun et al., 2024) section 2.1). Head/torso/tail entities are also discussed with the Head-to-Tail benchmark in Appendix A.

D.2 Full ablation study

We compare CONFQA with more alternatives than in the main content, as shown in Table 7.

- CONFQA No-dampener: same training as CONFQA. Does not apply the dampener in inference.

- Gen-as-label: the same strategy to choose questions the model can answer correctly as CONFQA, but use model generation as the true label, rather than the ground truth.
- IDK: the same as the IDK (DBPedia) in the main paper in Table 2.
- No-dampener: the same data labeling as CONFQA, but only use the question answer pairs, excluding the dampener prompt in the training input data. In other words, does not pass dampener as the system prompt.
- GT-as-label: feed in the original 3k rows of raw DBPedia data into the fine-tuning, i.e. all labels are ground truth.
- Fact-feeding: rather than using only the 3k sample of DBPedia data, also mixing 10k samples from Tulu3 data (Lambert et al., 2025).
- R-tuning (DBPedia): using our DBPedia training set and following R-tuning paper to generate labels for SFT.
- R-tuning (MMLU): using randomly sampled 3k MMLU samples to generate training set following R-tuning labeling strategy.
- MMLU-as-source: the same strategy as CONFQA, but use MMLU as data source. We use the same 3k samples from R-tuning (MMLU).

Table 7 reports results in two rows. Results on the top are evaluated using no system prompt, i.e. No-dampener applied during inference, only pass in questions to the models. The bottom rows are results with passing the dampener as the system prompt when doing inference.

CONFQA is overall more balanced that improves Missing rate to certain level, without impacting Correctness too much, while reducing Incorrectness to less than 5%. As shown in main paper Table 2, this leads to the optimal triggering F_{msr} .

We are not going to discuss each model one by one, but highlight two observations: (1) as we discussed in Section 5.3, generating training data from MMLU instead of DBPedia yields low hallucination ($< 2\%$) but far lower correctness. This is likely because MMLU’s diverse tasks reduce the model’s overall confidence (Intuition #3). (2) *Fact-feeding* combines our CONFQA fine-tuning data with Tulu facts, drops hallucinations but also correctness, similar to R-tuning. We suspect this is because what

our training data teach the LLM (saying unsure) is of different purpose from what the extra Tulu facts teach the LLM (feeding knowledge), when mixed together can offset each other and cause confusion.

D.3 Answer distributions for entity popularity

Figure 7 shows the distribution of correct, missing, and incorrect answers for entities of different popularity (Head, Torso, Tail), before and after fine-tuning, with and without dampening. It confirms that fine-tuning suppresses hallucinations (Intuition #1), and the dampener prompt further reduces hallucinations (Intuition #2). Additionally, it shows CONFQA suppresses more on long-tail facts, where it lacks confidence.

D.4 p-Value of CONFQA models

We compute p-Values for CONFQA model on the hallucination reduction metrics comparing with baseline Llama-3.1 Oracle and Llama3.1 Dampen, and report in Table 8. We compare CONFQA with Llama3.1 Dampen and CONFQA No-dampener with Llama-3.1 Oracle. We use McNemar p-Value since this is a paired test. For simplicity, we assume at least 70% of the questions are the same in baseline vs CONFQA. For factuality, we compute independent z-test p-Value. The results show that the improvements on Hallucination reduction shown in Table 2 are statistically significant on all benchmarks. Factuality improvements are also statistically significant for DBPedia, IMDb and SimpleQA benchmarks.

D.5 Gemma and QWen SFT

We conducted similar experiments on the QWen2.5-7B-Instruct model (Qwen et al., 2025) and the Gemma-3-4B-IT model (Team et al., 2025) through the Huggingface framework. Training data is generated using the logic described in Section 3.1:

- The same 3k of DBPedia simple question answer data, and labeled using the logic in 3.1.
- Evaluations are done using the same 4 short-form generation data sets.

Train using NVIDIA A100 80G 8 GPU. For both fine-tuning, we run one epoch, with learning rate $2e-4$ and gradient accumulation steps 2.

Table 9 shows experiments results using the QWen2.5-7B-Instruct model as baseline model. Table 10 shows experiment results from fine-tuning Gemma-3-4B-IT. A few observations based on these results:

Model	DBpedia (in-domain)				IMDb (out-of-domain)				SimpleQA (out-of-domain)				CRAG (out-of-domain)			
	Corr	Miss	Incor	Fac	Corr	Miss	Incor	Fac	Corr	Miss	Incor	Fac	Corr	Miss	Incor	Fac
Llama-3.1 Oracle	47.2	21.5	31.3	15.9	41.0	33.3	25.7	15.3	16.0	42.7	41.3	-25.3	60.3	11.7	28.0	32.3
CONFQA No-dampener	49.0	33.5	17.5	31.5	43.0	42.0	16.0	27.0	17.3	55.8	26.8	-9.5	57.0	19.6	23.4	33.6
* Gen-as-label	48.7	31.7	19.7	29.0	42.5	39.5	18.0	24.5	17.7	52.4	29.9	-12.3	57.6	18.4	24.0	33.6
* IDK (no-dampener)	44.5	40.3	15.2	29.3	40.7	45.8	13.5	27.2	14.4	65.0	20.6	-6.2	56.9	21.5	21.7	35.2
* No-dampener	42.0	34.7	23.3	18.7	40.2	38.0	21.8	18.4	12.0	66.4	21.6	-9.6	52.6	31.2	16.2	36.4
* GT-as-label	48.7	1.5	49.8	-1.1	42.0	0.2	57.8	-15.8	18.9	2.7	78.5	-59.6	58.1	5.3	36.6	21.5
* Fact-feeding	50.0	26.8	23.2	26.8	43.3	35.5	21.2	22.1	18.1	41.2	40.7	-22.6	56.9	16.5	26.6	30.3
* R-tuning (DBpedia)	53.7	6.7	39.7	14.0	44.5	<i>11.3</i>	44.2	0.3	22.5	<i>13.5</i>	64.0	-41.5	58.7	8.7	32.6	26.1
* R-tuning(MMLU)	50.2	19.2	30.7	19.5	45.5	28.2	26.3	19.2	20.3	38.0	41.7	-21.4	57.8	17.1	25.1	32.7
* MMLU-as-source	<i>50.5</i>	21.8	27.7	22.8	<i>44.2</i>	32.7	23.2	21.0	<i>20.4</i>	39.9	39.8	-19.4	56.1	18.8	25.1	31.0
Llama-3.1 Dampen	<i>47.0</i>	26.8	26.2	20.8	<i>40.7</i>	36.2	23.2	17.5	<i>16.8</i>	<i>48.0</i>	35.2	-18.4	57.5	22.3	20.2	37.2
CONFQA	31.5	63.3	5.2	26.3	32.5	63.3	4.2	28.3	4.9	93.1	2.1	2.8	39.4	56.2	4.4	35.0
* Gen-as-label (D)	28.5	65.7	5.8	22.7	27.7	69.8	2.5	25.2	3.1	96.0	1.9	<i>1.2</i>	32.7	64.0	3.3	29.4
* IDK	17.0	81.5	1.5	15.5	22.0	77.0	1.0	21.0	0.6	99.1	0.2	0.4	20.7	78.2	<i>1.1</i>	19.6
* No-dampener (D)	36.0	50.2	13.8	22.2	34.2	50.7	15.2	19.0	5.8	87.5	6.7	-0.9	46.0	44.2	9.8	36.2
* GT-as-label (D)	48.0	2.8	49.2	-1.2	41.2	4.3	54.5	-13.3	17.8	13.7	68.4	-50.6	53.7	14.2	32.1	21.6
* Fact-feeding (D)	20.7	76.7	2.7	18.0	25.5	70.7	3.8	21.7	2.5	94.7	2.8	-0.3	22.4	74.5	3.1	19.3
* R-tuning (DBpedia)	24.5	67.8	7.7	16.8	25.3	70.2	4.5	20.8	3.7	83.3	13.0	-9.3	31.6	55.0	13.4	18.2
* R-tuning(MMLU)	24.3	67.3	8.3	16.0	28.2	60.5	11.3	16.9	5.8	85.1	9.1	-3.3	31.3	56.5	12.1	19.2
* MMLU-as-source (D)	8.2	89.8	2.0	6.2	16.7	82.0	<i>1.3</i>	15.4	0.6	98.8	<i>0.5</i>	0.1	7.0	92.7	0.3	6.7

Table 7: Ablation study of CONFQA, showing effectiveness of our fine-tuned model to reduce hallucination and improve factuality (all numbers are in percentage (%); bold indicates the best performance and italic the second best.)




Figure 7: CONFQA suppresses more (i.e., gives more missing answers) on less popular entities; that is, missing rate for Tail is larger than Torso, which is in turn larger than Head.

Model	Incor (p-Value)	Fac (p-Value)	Incor (p-Value)	Fac (p-Value)
	DBpedia (in-domain)		IMDb (out-of-domain)	
Llama-3.1 Oracle	31.3	15.9	25.7	15.3
Llama-3.1 Dampen	26.2	20.8	23.2	17.5
CONFQA No-dampener	17.5 (5.5E-12)	31.5 (2.1E-10)	16.0 (7.2E-08)	27.0 (7.0E-07)
CONFQA	5.2 (<1.0E-16)	26.3 (2.5E-02)	4.2 (<1.0E-16)	28.3 (8.5E-06)
	SimpleQA (out-of-domain)		CRAG (out-of-domain)	
Llama3.1 Oracle	41.3	-25.3	28.0	32.3
Llama3.1 Dampen	35.2	-18.4	20.2	37.2
CONFQA No-dampener	26.8 (<1.0E-16)	-9.5 (<1.0E-16)	23.4 (6.2E-03)	33.6 (6.2E-01)
CONFQA	2.1 (<1.0E-16)	2.8 (<1.0E-16)	4.4 (<1.0E-16)	35.0 (4.1E-01)

Table 8: Factuality improvement on short-form benchmarks with p-Value; CONFQA models can reduce hallucination (Incor) to less than 5% with the dampener prompt with significant difference. All numbers are in percentage (%).

- Our proposed method CONFQA can reduce hallucination (Incor) by 13-50% for Qwen 7B when applying the dampener prompt.
- The hallucination could be reduced to close or below 5% for Gemma 4B model in the same case.
- Qwen 7B has fairly low correctness, and CONFQA can further reduce it as it changes low-confidence answers into unsure answers. However, the factuality increases by 12-37%, showing that it reduces much more hallucinations than correct answers.
- Similar for Gemma model: fairly low correctness, as it is a even smaller model. Comparing to Qwen, the factuality increase is more effective for Gemma model: increases by 20-89%.
- Transferability: the fine-tuning on DBpedia atomic question answering pairs could extend to out-of-domain datasets.

Model	Corr	Miss	Incor	Fac.	Corr	Miss	Incor	Fact.
	DBPedia (in-domain)				IMDb (out-of-domain)			
QWen2.5	21.3	25.8	52.8	-31.5	19.3	15.7	65.0	-45.7
QWen2.5 Dampen	12.0	64.2	23.5	-11.5	14.7	56.7	28.3	-13.6
CONFQA No-dampener in Inf.	21.5	50.8	27.7	-6.2	19.5	31.5	49.0	-29.5
CONFQA	8.3	88.7	3.0	5.3	12.7	74.0	13.3	-0.6
	SimpleQA (out-of-domain)				CRAG (out-of-domain)			
QWen2.5	3.8	21.6	74.5	-70.7	22.7	30.1	47.2	-24.5
QWen2.5 Dampen	1.7	71.5	26.8	-25.1	15.7	57.8	26.5	-10.8
CONFQA No-dampener in Inf.	2.8	61.0	36.2	-33.4	24.3	41.1	34.6	-10.3
CONFQA	0.8	86.1	13.1	-12.3	8.3	85.2	6.5	1.8

Table 9: Overall factuality improvement on short-form benchmarks for QWen2.5-7B-Instruct; CONFQA can reduce hallucination to around 10% with the dampener prompt. All numbers are in percentage (%).

Model	Corr	Miss	Incor	Fac.	Corr	Miss	Incor	Fact.
	DBPedia (in-domain)				IMDb (out-of-domain)			
Gemma3	19.5	4.5	76.0	-56.5	17.2	2.0	80.8	-63.6
Gemma3 Dampen	20.5	4.3	75.2	-54.7	19.5	0.3	80.2	-60.7
CONFQA No-dampener in Inf.	22.3	38.8	38.8	-16.5	16.2	24.8	59.0	-42.8
CONFQA	6.0	92.2	1.8	4.2	9.0	85.2	5.8	3.2
	SimpleQA (out-of-domain)				CRAG (out-of-domain)			
Gemma3	4.0	1.8	94.2	-90.2	26.8	3.0	70.2	-43.4
Gemma3 Dampen	3.7	0.9	95.4	-91.7	29.1	2.6	68.2	-39.1
CONFQA No-dampener in Inf.	2.5	43.1	54.5	-52.0	19.3	39.3	41.4	-22.1
CONFQA	0.2	97.2	2.5	-2.3	3.3	95.0	1.7	1.6

Table 10: Overall factuality improvement on short-form benchmarks for Gemma-3-4B-IT; CONFQA can reduce hallucination to below 5% with the dampener prompt. All numbers are in percentage (%).

D.6 Dampener prompt wording

To understand how sensitive is the approach to the exact wording of the dampening prompt, we evaluated the fine-tuned CONFQA model with different wordings in the prompt. The CONFQA prompt is listed in Appendix Prompt 1. In this experiment, we only change the dampener part of the whole prompt, i.e. *Answer only if you are confident; otherwise, respond with 'I am unsure about the answer'*. In the experiment, Prompt 1: *Answer only if you are confident; otherwise, respond with 'I am unsure'*, Prompt 2: *Answer if you are confident; otherwise, respond with 'I am unsure about the answer'*.

Table 11 shows results from using different prompts in the evaluation. We see that Prompt 1 and 2 give very similar results compared to baseline CONFQA.

E Dynamic questions identification

In Section 3.2, we briefly mentioned that dynamic questions—those requiring up-to-date information—should be answered by the RAG pipeline. In real

application, we fine-tune a small-sized model (such as Llama-3.1 8B Instruct) to do this classification task: identify if a question is dynamic (whose answer may change over time) or static (whose answer is static).

Training data and prompt: Take CRAG training data as example, instead of using ground truth answer as label, we label the questions as either "dynamic" or "static". We use prompt to instruct the model to output only "dynamic" or "static". Respond with "dynamic" if the answer to the question could change over time, and respond with "static" if the answer to the question does not change over time. The same supervised fine-tuning is applied to the model. We normally observe 80-90% precision and recall of this method in identifying dynamic questions. We omit additional details as this is not relevant to the key content of this paper.

Prompt	Corr	Miss	Incor	Fac.	Corr	Miss	Incor	Fact.
	DBPedia (in-domain)				IMDb (out-of-domain)			
CONFQA	31.5	63.3	5.2	26.3	32.5	63.3	4.2	28.3
Prompt 1	30.7	64.0	5.3	25.4	31.2	65.2	3.7	27.5
Prompt 2	32.0	63.5	4.5	27.5	32.0	64.7	3.3	28.7
	SimpleQA (out-of-domain)				CRAG (out-of-domain)			
CONFQA	4.9	93.1	2.1	2.8	39.4	56.2	4.4	35.0
Prompt 1	5.0	93.3	1.7	3.3	39.9	55.9	4.2	35.7
Prompt 2	5.9	92.7	1.3	4.6	37.9	58.3	3.9	34.0

Table 11: Overall factuality difference with different wording in the dampener prompt. All numbers are in percentage (%).