RETHINKING FAITHFULNESS IN LLMS: A PAIRWISE CONTEXT-SENSITIVE PERSPECTIVE

Anonymous authors
Paper under double-blind review

ABSTRACT

Large language models (LLMs) are expected to answer questions faithfully based on the provided context, abstaining when the context information is insufficient to answer the questions. Existing faithfulness evaluations typically assess each question-context instance in isolation; however, such instance-level evaluation fails to capture a fundamental requirement of faithful behavior: the ability to adapt model responses to changes in available contexts. In particular, a model should provide correct answers when sufficient evidence is present and abstain when it is not. In this work, we propose a Pairwise Faithfulness Benchmark (PFaithBench) that evaluates whether a model can switch between answering and abstaining for the same question under supporting versus non-supporting contexts. Our evaluations across thirty-nine models with seven model families demonstrate that faithfulness fundamentally involves a trade-off between answering and abstaining, and that most current models exhibit a strong bias toward answering, with most faithfulness errors arising from over-answering, i.e., models tend to fabricate a response even when the provided context is insufficient. We further conduct a series of studies on faithfulness training under different data constructions. Our results show that training outcomes are highly sensitive to the specific composition of answering and abstaining data. Constructing answering and abstaining data from mismatched sources can cause models to rely on dataset-specific shortcuts rather than actual context sufficiency. Moreover, increasing answer-supervised data improves answering performance but exacerbates over-answering, while increasing abstaining data reduces hallucination but leads to over-abstention.

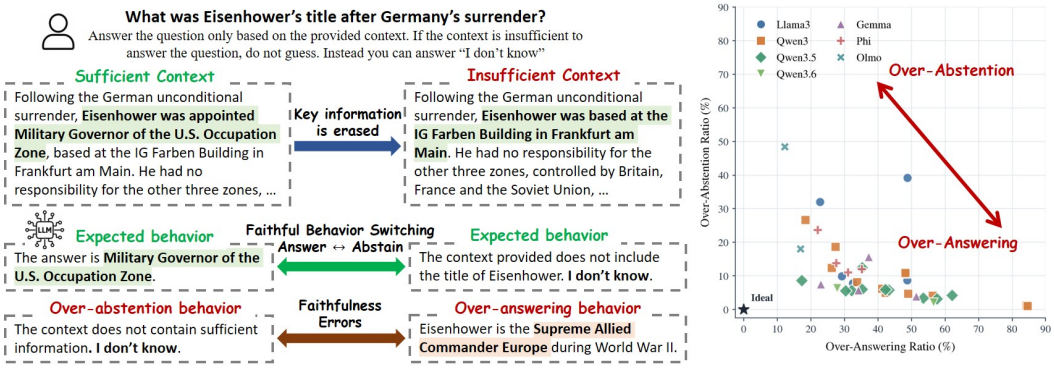

Figure 1: **Overview of pairwise faithfulness behavior**. *Left:* Given the same question with sufficient and insufficient contexts, a faithful model should exhibit context-sensitive behavior: answering when sufficient evidence is available and abstaining when key information is missing. Over-abstention and over-answering are two sources of faithfulness errors. *Right:* Evaluations across seven families with thirty-nine models: most models exhibit bias toward over-answering, where one point is one model.

1 INTRODUCTION

Large language models (LLMs) are increasingly deployed in knowledge-intensive applications (Bhathena et al., 2026; Kalra et al., 2024), often with external contexts provided through retrieval-

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

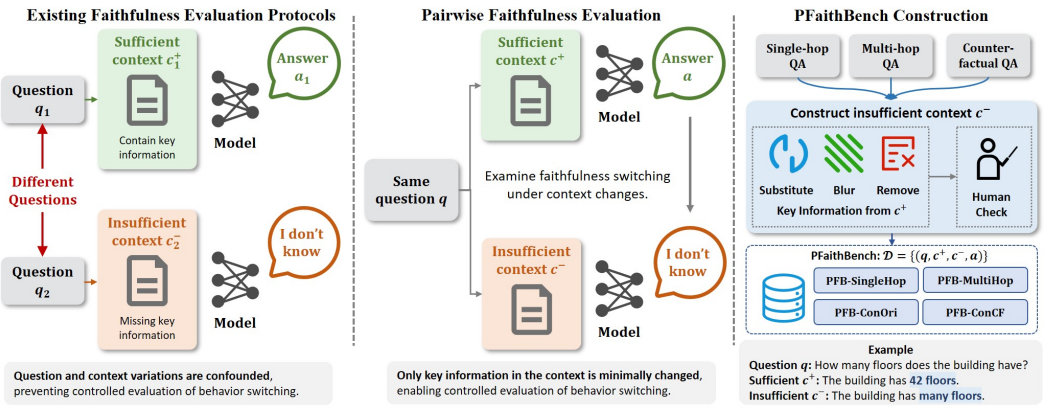

Figure 2: **Pairwise faithfulness evaluation and PFaithBench construction.** *Left:* Existing evaluation protocols assess individual question-context instances, confounding question and context variations. *Middle:* Pairwise evaluation keeps the question fixed while varying contextual sufficiency, enabling controlled assessment of faithfulness switching. *Right:* PFaithBench is constructed from multi-hop, single-hop, and counterfactual QA by modifying key information in sufficient contexts.

augmented generation (RAG) (Gao et al., 2023) or domain-specific knowledge bases (Edge et al., 2024). In such setting, a fundamental requirement for reliable deployment of LLMs is *faithfulness*: models should generate responses that are strictly grounded in the provided context, answering when sufficient evidence context is available and abstaining when it is not (Huang et al., 2025; Ming et al., 2025; Bi et al., 2025). However, LLM may fabricate fluent but unsupported responses when context is insufficient, leading to over-answering and hallucinations, or abstain despite sufficient context, resulting in over-abstention (Joren et al., 2025). This raises a fundamental question: *how should LLM behavior be characterized when the provided context is sufficient versus insufficient?*

Existing evaluation protocols typically assess faithfulness at the level of individual question-context instances. Prior work considers a range of contextual conditions, from sufficient context for evaluating answer correctness (Adlakha et al., 2024; Ho et al., 2020) to insufficient context for assessing abstaining behavior (Ming et al., 2025), as well as settings involving noisy (Chen et al., 2024), counterfactual (Bi et al., 2025), or time-related (Chen et al., 2021) context. However, as illustrated in Fig. 2, these conditions are usually evaluated independently and on different sets of questions. This entangles question variation with context variation, preventing a controlled assessment of whether models appropriately switch between answering and abstaining as contextual sufficiency changes.

In this work, we propose a *Pairwise Faithfulness Benchmark* (PFaithBench), which characterizes model behavior under sufficient and insufficient contextual conditions for the same question. As illustrated in Fig. 2, by keeping the question fixed while varying contextual sufficiency, PFaithBench enables a controlled assessment of whether models appropriately switch between answering and abstaining, without confounding from the variation of questions. We characterize two primary types of faithfulness errors: (i) *Over-answering* (OverAns), where models answer despite insufficient context, and (ii) *Over-abstention* (OverAbs), where models abstain despite sufficient context. As shown in Fig. 1, our evaluations across 39 models from seven families find that most existing models exhibit a bias toward over-answering, and a portion of over-answering stems from internal knowledge leakage, where models rely on internal parametric knowledge beyond the provided context.

Additionally, we conduct a series of data-centric analyses to investigate how training data composition affects faithfulness behaviors of models. We find that constructing answering and abstaining data from mismatched sources can induce dataset-specific shortcuts, causing models to rely on source-specific signals rather than contextual sufficiency. We further observe a clear trade-off between answering and abstaining supervision: increasing answering data improves answering performance but exacerbates over-answering, whereas increasing abstention data reduces unsupported answering but leads to over-abstention. These findings highlight the importance of proper context coupling and a balanced composition of answering and abstaining data for effective faithfulness training.

Our main contributions of this work are summarized as follows:

Table 1: **Evaluation of Over-answering and Over-abstention on PFaithBench across diverse LLMs.** Lightred indicates OverAns exceeds OverAbs by more than 5%, lightblue indicates OverAbs exceeds OverAns by more than 5%, and lightorange indicates an absolute difference within 5%.

Family	LLM	PFB-MultiHop		PFB-SingleHop		PFB-ConOri		PFB-ConCF		Avg.	
		OverAns ↓	OverAbs ↓	OverAns ↓	OverAbs ↓	OverAns ↓	OverAbs ↓	OverAns ↓	OverAbs ↓	OverAns ↓	OverAbs ↓
Llama3	Llama-3.2-1B-Instruct	53.90	35.50	52.76	40.81	44.25	37.04	44.37	43.11	48.82	39.12
	Llama-3.2-3B-Instruct	37.40	23.90	12.99	43.18	19.47	23.51	21.11	37.42	22.74	32.00
	Llama-3.1-8B-Instruct	56.60	6.50	45.41	6.04	45.89	3.79	47.03	18.46	48.73	8.70
	Llama-3.1-70B-Instruct	37.20	7.40	21.00	5.64	28.19	7.71	30.47	18.58	29.21	9.83
	Llama-3.3-70B-Instruct	37.20	7.20	27.69	3.54	32.24	6.45	33.12	13.78	32.56	7.74
Qwen3	Qwen3-0.6B (NoThink)	57.00	17.80	49.87	11.29	44.50	5.44	41.85	7.96	48.30	10.62
	Qwen3-1.7B (NoThink)	57.00	17.90	49.48	10.76	43.11	5.69	43.36	8.98	48.24	10.83
	Qwen3-4B (NoThink)	32.90	25.80	20.34	18.24	28.19	8.22	27.94	22.12	27.34	18.60
	Qwen3-8B (NoThink)	23.20	31.70	20.87	15.09	14.79	13.91	14.92	45.76	18.44	26.62
	Qwen3-14B (NoThink)	36.70	13.50	22.31	6.69	21.11	6.32	24.78	22.88	26.23	12.35
	Qwen3-32B (NoThink)	41.20	7.10	35.04	2.76	27.69	5.31	31.35	17.19	33.82	8.09
	Qwen3-0.6B (Think)	84.90	2.10	88.06	1.05	80.78	1.14	83.44	0.51	84.30	1.20
	Qwen3-1.7B (Think)	85.60	1.90	89.11	0.52	81.80	0.63	82.17	1.01	84.67	1.02
	Qwen3-4B (Think)	46.30	3.70	41.34	2.76	40.46	4.55	41.21	8.85	42.33	4.96
	Qwen3-8B (Think)	46.70	5.20	35.30	3.15	40.71	4.93	42.23	11.38	41.23	6.16
	Qwen3-14B (Think)	45.80	3.70	44.49	1.71	49.05	3.41	56.89	9.73	49.06	4.64
	Qwen3-32B (Think)	52.00	3.30	47.90	1.31	60.30	3.79	65.36	7.71	56.39	4.03
Qwen3.5	Qwen3.5-0.8B (NoThink)	53.80	12.00	36.22	15.22	26.17	9.61	24.91	12.77	35.27	12.40
	Qwen3.5-2B (NoThink)	60.50	5.30	46.46	4.33	31.23	5.18	35.02	8.22	43.30	5.76
	Qwen3.5-4B (NoThink)	52.00	4.60	31.36	2.62	29.46	5.06	28.45	11.76	35.32	6.01
	Qwen3.5-9B (NoThink)	47.00	3.80	30.84	2.23	26.04	5.94	24.65	10.37	32.13	5.58
	Qwen3.5-27B (NoThink)	48.80	3.20	27.82	1.57	22.00	5.69	22.88	11.76	30.38	5.56
	Qwen3.5-0.8B (Think)	76.80	2.60	62.47	3.02	55.37	4.42	53.86	6.83	62.12	4.22
	Qwen3.5-2B (Think)	63.80	3.00	37.80	2.89	33.50	6.57	34.01	11.25	42.28	5.93
	Qwen3.5-4B (Think)	74.50	1.80	55.25	0.79	49.94	3.92	50.82	6.19	57.63	3.18
	Qwen3.5-9B (Think)	32.20	5.50	18.11	5.12	8.47	8.09	10.49	15.80	17.32	8.63
	Qwen3.5-27B (Think)	80.10	1.10	50.52	0.92	40.08	4.42	43.36	6.95	53.52	3.35
Qwen3.6	Qwen3.6-27B (NoThink)	41.30	4.40	26.25	1.71	22.38	6.57	21.49	12.90	27.85	6.39
	Qwen3.6-27B (Think)	79.10	0.90	54.46	0.26	45.26	3.29	47.41	4.05	56.56	2.12
Gemma	Gemma-2-2B-It	44.00	30.90	29.53	13.12	38.56	6.70	36.92	11.76	37.25	15.62
	Gemma-2-9B-It	70.90	5.00	42.52	2.76	46.02	2.91	46.27	4.80	51.43	3.87
	Gemma-2-27B-It	53.40	6.40	30.31	2.49	23.77	4.93	29.08	8.85	34.14	5.67
	Gemma-4-31B-it	38.80	5.60	19.95	2.23	15.68	8.85	17.57	13.27	23.00	7.49
Phi	Phi-3-mini-128k-instruct	34.80	26.40	17.85	29.13	17.57	13.02	17.95	25.92	22.04	23.62
	Phi-3.5-mini-instruct	55.30	14.90	25.85	16.14	30.09	4.93	29.58	12.01	35.21	12.00
	Phi-4-mini-instruct	38.70	16.10	22.57	16.40	23.26	7.33	25.92	15.17	27.61	13.75
	Phi-4	42.30	6.30	30.18	4.99	24.78	8.47	27.18	24.27	31.11	11.01
Olmo	Olmo-3-7B-Instruct	19.60	52.60	6.69	54.07	10.37	30.85	12.26	56.01	12.23	48.38
	Olmo-3.1-32B-Instruct	21.90	20.20	13.25	14.44	14.03	10.49	18.46	27.05	16.91	18.05

- We propose PFaithBench, which assesses whether models can appropriately switch between answering and abstaining for the same question under sufficient and insufficient contexts.
- We identify two primary sources of faithfulness errors: over-answering and over-abstention, demonstrating a prevalent tendency toward over-answering among existing models.
- We conduct data-centric analyses, highlighting the importance of proper context coupling and balanced answering-abstention data composition for faithfulness training.

2 Related Work

Faithfulness evaluation. Faithfulness has emerged as a key requirement for LLMs deployed in information-sensitive RAG systems (Lewis et al., 2020; Gao et al., 2023), where models are expected to strictly ground responses in the provided context rather than parametric knowledge (Ji et al., 2023). Representative benchmarks include FaithEval (Ming et al., 2025), which examines unanswerable, inconsistent, and counterfactual contexts; RGB (Chen et al., 2024), which investigates robustness to noisy retrieval; ClashEval (Wu et al., 2024) and ContextDPO (Bi et al., 2025), which study conflicts between internal knowledge and external evidence; TimeQA (Chen et al., 2021), which constructs time-sensitive contexts; and 2WikiMultiHopQA (Ho et al., 2020) and HotpotQA (Yang et al., 2018), which focus on answer correctness in multi-hop QA. However, existing evaluations typically assess different questions under different contextual conditions, confounding question and context variations and obscuring models' ability to switch between answering and abstaining. This motivates our pairwise evaluation, which constructs paired sufficient and insufficient contexts for the same question.

Faithfulness training. Existing studies have explored various approaches to improve faithfulness, including context-grounded fine-tuning, such as RA-DIT (Lin et al., 2023) and RAFT (Zhang et al., 2024b); abstention-aware supervision, such as R-Tuning (Zhang et al., 2024a), Alignment for Honesty (Yang et al., 2024b), and CanAISayNo (Cheng et al., 2024); preference optimization, such as ContextDPO (Bi et al., 2025); and reinforcement learning, such as FaithRL (Gui et al., 2026), KnowRL (Ren et al., 2025), and Faithfulness-Aware Step-Level RL (Nie et al., 2026). Despite these advances, they primarily focus on teaching models to exhibit appropriate behaviors under specific contextual conditions. In this work, we adopt the widely used GRPO (Shao et al., 2024) as our

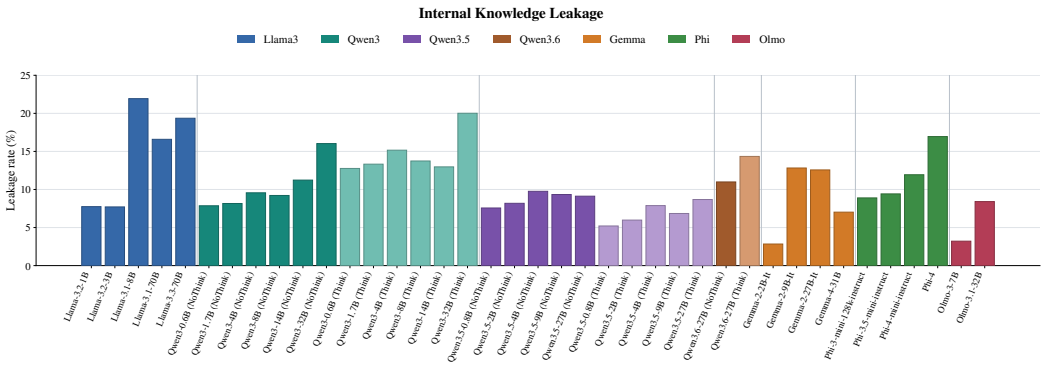

Figure 3: **Internal knowledge leakage across models.** Leakage measures the extent to which models produce answers by relying on internal parametric knowledge rather than the provided context. Colors indicate different model families, and model sizes increase from left to right within each family.

training algorithm and conduct a series of data-centric analyses to investigate how training data construction and composition affect context-sensitive answering and abstaining behaviors of models.

3 Pairwise Faithfulness Evaluation

To begin, we formalize the pairwise faithfulness evaluation (Sec. 3.1), introduce the corresponding evaluation metrics (Sec. 3.2), and describe the construction of PFaithBench (Sec. 3.3).

3.1 Problem Definition

Pairwise faithfulness evaluation considers question answering (QA) under varying contextual conditions. Given a question q and its ground-truth answer a, let c^+ and c^- denote two types of contexts associated with q: a sufficient context c^+ that contains supported key information to answer the question, and an insufficient context c^- lacks such information. Given a model π_θ, which takes a question-context pair (q, c) as input and samples a response $y \sim \pi_\theta(\cdot|q, c)$, pairwise faithfulness requires the model to exhibit context-sensitive behaviors as follows:

$$y^+ \sim \pi_\theta(\cdot|q, c^+) \rightarrow \text{answer correctly}, \qquad y^- \sim \pi_\theta(\cdot|q, c^-) \rightarrow \text{abstain appropriately}. \tag{1}$$

That is, the model should generate an answer when sufficient evidence is available, and abstain (e.g., by responding "I don't know") when the context is insufficient. Ideally, a perfectly faithful model should appropriately switch between these two behaviors when c^+ and c^- differ only in the key information required to answer the question. We formalize this setting as *pairwise faithfulness evaluation*, where each question q is associated with a pair of contexts (c^+, c^-), with c^- constructed by minimally modifying the key information in c^+. Accordingly, each data is represented as a tuple (q, c^+, c^-, a), where a denotes the ground-truth answer to q. The ground-truth answer a enables us to assess response correctness under sufficient context and identify correct answers generated despite insufficient context, potentially indicating reliance on internal parametric knowledge.

Under this formulation, model responses can be categorized into three types: (i) abstention (IDK), where the model chooses not to answer the question (e.g., by responding "I don't know"); (ii) correct (C), where the response matches the ground-truth a; and (iii) wrong (W), where the response does not match the ground-truth a and lacks sufficient supporting evidence from the provided context.

3.2 Metrics of Pairwise Faithfulness

Based on the response categories $y \in \{\text{IDK}, \text{C}, \text{W}\}$ of the model, beyond simple accuracy, a set of metrics can be defined to quantify pairwise faithfulness behavior as follows:

Over-answering vs. over-abstention. *Over-answering* (OverAns) measures the tendency of the model to produce a non-abstaining response under insufficient context, while *Over-abstention* (OverAbs) measures the tendency of the model to abstain despite sufficient context:

$$\text{OverAns} = \mathbb{P}(y^- \in \{\text{C}, \text{W}\}), \qquad \text{OverAbs} = \mathbb{P}(y^+ = \text{IDK}). \tag{2}$$

Table 2: **Evaluation of pairwise accuracy DFAcc and SFAcc on PFaithBench across diverse LLMs.** The Gap $=$ DFAcc $-$ SFAcc indicates the discrepancy between appropriate answer-or-abstain decisions and correct answering. Lightred indicates large gaps (Gap ≥ 10), while lightblue indicates small gaps (Gap < 10). The best results in each family are highlighted in **bold**.

Family	LLM	PFB-MultiHop		PFB-SingleHop		PFB-ConOri		PFB-ConCF		Avg.		
		DFAcc ↑	SFAcc ↑	DFAcc ↑	SFAcc ↑	DFAcc ↑	SFAcc ↑	DFAcc ↑	SFAcc ↑	DFAcc ↑	SFAcc ↑	Gap ↓
Llama3	Llama-3.2-1B-Instruct	21.50	5.60	17.06	6.69	27.81	7.21	26.68	3.79	23.26	5.82	17.44
	Llama-3.2-3B-Instruct	42.20	26.00	46.19	25.85	59.29	43.87	47.91	15.93	48.90	27.91	20.99
	Llama-3.1-8B-Instruct	37.70	30.00	49.87	39.50	50.95	45.13	42.10	19.09	45.15	33.43	11.72
	Llama-3.1-70B-Instruct	**56.10**	**50.00**	**73.62**	**68.90**	**64.85**	**59.04**	54.61	22.50	**62.30**	**50.11**	12.19
	Llama-3.3-70B-Instruct	56.00	49.30	69.16	64.70	62.20	57.77	**55.25**	**26.30**	60.65	49.52	11.13
Qwen3	Qwen3-0.6B (NoThink)	29.80	12.30	40.55	28.35	51.20	42.86	51.45	25.54	43.25	27.26	15.99
	Qwen3-1.7B (NoThink)	28.40	11.40	41.47	26.90	52.47	41.97	49.30	24.78	42.91	26.26	16.65
	Qwen3-4B (NoThink)	44.10	32.90	63.12	56.56	63.72	59.29	55.12	29.58	56.52	44.58	11.93
	Qwen3-8B (NoThink)	47.50	37.90	64.83	60.63	71.81	69.15	46.14	34.64	57.57	50.58	6.99
	Qwen3-14B (NoThink)	50.30	41.90	**71.13**	**69.03**	**72.69**	**69.79**	**57.14**	**41.21**	**62.82**	**55.48**	7.33
	Qwen3-32B (NoThink)	**52.80**	**46.80**	62.60	60.50	67.26	63.84	55.50	35.65	59.54	51.70	7.84
	Qwen3-0.6B (Think)	14.50	8.40	11.42	9.58	18.58	16.43	16.43	9.86	15.23	11.07	4.17
	Qwen3-1.7B (Think)	13.30	7.40	10.89	9.19	17.95	15.17	17.70	10.87	14.96	10.66	4.30
	Qwen3-4B (Think)	50.90	42.00	56.43	52.36	**55.50**	50.32	**51.71**	**32.49**	53.63	44.29	9.34
	Qwen3-8B (Think)	49.20	41.80	**62.07**	**57.87**	55.12	**50.95**	50.70	31.23	**54.27**	**45.46**	8.81
	Qwen3-14B (Think)	**51.30**	**42.80**	53.94	50.39	48.04	42.86	38.56	24.02	47.96	40.02	7.94
	Qwen3-32B (Think)	45.60	39.50	51.05	48.82	36.54	33.88	30.34	18.08	40.88	35.07	5.81
Qwen3.5	Qwen3.5-0.8B (NoThink)	35.80	22.40	50.39	30.45	65.49	55.63	64.10	30.21	53.94	34.67	19.27
	Qwen3.5-2B (NoThink)	35.40	26.10	50.00	39.63	64.48	56.01	58.91	31.35	52.20	38.27	13.92
	Qwen3.5-4B (NoThink)	43.90	36.30	66.27	60.37	65.99	62.45	62.07	33.38	59.56	48.12	11.44
	Qwen3.5-9B (NoThink)	**49.80**	**42.60**	67.19	64.57	68.27	63.34	66.37	**39.57**	62.91	52.52	10.39
	Qwen3.5-27B (NoThink)	48.30	40.80	**70.60**	**68.64**	**72.44**	**68.65**	**66.75**	37.04	**64.52**	**53.78**	10.74
	Qwen3.5-0.8B (Think)	21.50	13.00	35.17	27.82	41.21	34.13	42.10	21.74	35.00	24.18	10.82
	Qwen3.5-2B (Think)	33.90	27.80	59.58	54.46	60.94	56.26	56.51	31.86	52.73	42.59	10.14
	Qwen3.5-4B (Think)	24.10	18.30	44.23	40.81	46.65	42.98	44.88	26.30	39.96	32.10	7.87
	Qwen3.5-9B (Think)	**63.10**	**55.30**	**76.77**	**74.15**	**83.57**	**79.90**	**75.47**	**48.17**	**74.73**	**64.38**	10.35
	Qwen3.5-27B (Think)	19.10	13.00	48.69	45.93	56.13	51.20	51.45	28.45	43.84	34.64	9.20
Qwen3.6	Qwen3.6-27B (NoThink)	**54.40**	**47.00**	**72.05**	**69.42**	**71.43**	**67.38**	**67.26**	**36.92**	**66.28**	**55.18**	11.10
	Qwen3.6-27B (Think)	20.60	16.60	45.41	44.09	51.83	48.42	49.05	25.92	41.72	33.76	7.97
Gemma	Gemma-2-2B-It	31.70	7.60	59.58	40.42	56.13	42.48	54.11	23.39	50.38	28.47	21.91
	Gemma-2-9B-It	25.30	19.00	54.86	50.92	51.45	46.90	50.06	28.82	45.42	36.41	9.01
	Gemma-2-27B-It	40.80	31.90	67.45	62.86	71.55	66.88	63.84	**35.02**	60.91	49.16	11.75
	Gemma-4-31B-it	**56.20**	**49.60**	**77.82**	**75.33**	**75.73**	**70.67**	**69.66**	34.77	**69.85**	**57.59**	12.26
Phi	Phi-3-mini-128k-instruct	41.60	26.00	54.46	41.08	69.53	61.82	59.17	25.54	56.19	38.61	17.58
	Phi-3.5-mini-instruct	32.70	19.30	59.58	43.04	65.49	55.88	60.68	26.93	54.61	36.29	18.32
	Phi-4-mini-instruct	47.40	36.90	62.34	51.57	**70.04**	**64.48**	**62.07**	**27.43**	**60.46**	45.10	15.37
	Phi-4	**52.00**	**43.50**	**65.35**	**60.10**	67.26	60.56	53.48	23.89	59.52	**47.01**	12.51
Olmo	Olmo-3-7B-Instruct	32.40	22.40	40.55	33.73	59.67	51.83	34.26	21.62	41.72	32.39	9.33
	Olmo-3.1-32B-Instruct	**59.00**	**49.20**	**73.23**	**69.16**	**75.73**	**69.79**	**59.17**	**34.77**	**66.78**	**55.73**	11.05

These two metrics capture the primary sources of faithfulness errors, reflecting the opposing tendencies of the model to answer too aggressively or abstain too conservatively.

Internal knowledge leakage. We quantify the extent to which models rely on internal parametric knowledge rather than the provided context by measuring the probability that a model produces the ground-truth answer despite insufficient context:

$$\text{Leakage} = \mathbb{P}(y^- = \text{C}). \tag{3}$$

Ideally, a faithful model should follow the instruction of answering the question strictly based on the provided context rather than its internal parametric knowledge. However, models may still rely on memorized knowledge to answer questions when contextual evidence is insufficient (Ming et al., 2025; Gui et al., 2026), a phenomenon we refer to as *internal knowledge leakage*.

Pairwise faithfulness accuracy. We define two accuracy metrics to capture pairwise faithfulness at different levels of strictness: *Decision Faithfulness Accuracy* (DFAcc) and *Strict Faithfulness Accuracy* (SFAcc). Specifically, DFAcc measures whether the model appropriately switches its behaviors between answering and abstaining under sufficient and insufficient contexts, while SFAcc additionally requires the answer to be correct under sufficient context:

$$\text{DFAcc} = \mathbb{P}(y^+ \in \{\text{C}, \text{W}\},\ y^- = \text{IDK}), \qquad \text{SFAcc} = \mathbb{P}(y^+ = \text{C},\ y^- = \text{IDK}). \tag{4}$$

Compared to DFAcc, SFAcc imposes a stricter requirement by considering answer correctness. The gap between DFAcc and SFAcc reflects the model's inability to produce correct answers under sufficient context despite making appropriate answering and abstaining decisions.

Margin faithfulness accuracy. We also evaluate model behavior under each contextual condition separately: *Answer Accuracy* (AnsAcc) under sufficient context and *Abstention Accuracy* (AbsAcc) under insufficient context:

$$\text{AnsAcc} = \mathbb{P}(y^+ = \text{C}), \qquad \text{AbsAcc} = \mathbb{P}(y^- = \text{IDK}). \tag{5}$$

Table 3: **Evaluation of AnsAcc and AbsAcc on PFaithBench across diverse LLMs.** Lightred indicates that the absolute difference $|\mathrm{AnsAcc} - \mathrm{AbsAcc}|$ is larger than 10%, while lightblue indicates that the absolute difference is at most 10%. The best results in each family are highlighted in **bold**.

Family	LLM	PFB-MultiHop		PFB-SingleHop		PFB-ConOri		PFB-ConCF		Avg.	
		AnsAcc ↑	AbsAcc ↑	AnsAcc ↑	AbsAcc ↑	AnsAcc ↑	AbsAcc ↑	AnsAcc ↑	AbsAcc ↑	AnsAcc ↑	AbsAcc ↑
Llama3	Llama-3.2-1B-Instruct	18.60	46.10	21.92	47.24	21.49	55.75	6.19	55.63	17.05	51.18
	Llama-3.2-3B-Instruct	45.90	62.60	32.28	**87.01**	57.65	**80.53**	18.20	**78.89**	38.51	**77.26**
	Llama-3.1-8B-Instruct	73.80	43.40	73.36	54.59	84.45	54.11	29.58	52.97	65.30	51.27
	Llama-3.1-70B-Instruct	**82.10**	**62.80**	86.88	79.00	85.46	71.81	30.34	69.53	71.19	70.79
	Llama-3.3-70B-Instruct	81.60	**62.80**	**88.71**	72.31	**86.35**	67.76	**34.51**	66.88	**72.79**	67.44
Qwen3	Qwen3-0.6B (NoThink)	36.40	43.00	59.84	50.13	80.03	55.50	43.62	58.15	54.97	51.70
	Qwen3-1.7B (NoThink)	38.20	43.00	58.79	50.52	76.61	56.89	40.20	56.64	53.45	51.76
	Qwen3-4B (NoThink)	56.80	67.10	73.10	**79.66**	85.59	71.81	36.92	72.06	63.10	72.66
	Qwen3-8B (NoThink)	55.50	**76.80**	79.27	79.13	82.68	**85.21**	39.19	**85.08**	64.16	**81.56**
	Qwen3-14B (NoThink)	70.30	63.30	88.98	77.69	89.38	78.89	**51.45**	75.22	75.03	73.77
	Qwen3-32B (NoThink)	**82.00**	58.80	**92.91**	64.96	**90.14**	72.31	47.79	68.65	**78.21**	66.18
	Qwen3-0.6B (Think)	61.90	15.10	78.87	11.94	88.24	19.22	48.42	16.56	69.36	15.70
	Qwen3-1.7B (Think)	63.00	14.40	78.35	10.89	88.12	18.20	48.67	17.83	69.53	15.33
	Qwen3-4B (Think)	81.20	53.70	89.11	58.66	87.61	**59.54**	**50.19**	**58.79**	77.03	57.67
	Qwen3-8B (Think)	81.00	53.30	88.71	**64.70**	87.48	59.29	48.04	57.77	76.31	**58.77**
	Qwen3-14B (Think)	82.90	**54.20**	91.34	55.51	86.85	50.95	49.18	43.11	77.57	50.94
	Qwen3-32B (Think)	**85.90**	48.00	**93.96**	52.10	**90.39**	39.70	48.17	34.64	**79.61**	43.61
Qwen3.5	Qwen3.5-0.8B (NoThink)	59.30	46.20	54.46	63.78	75.60	73.83	38.18	75.09	56.89	64.73
	Qwen3.5-2B (NoThink)	71.20	39.50	78.87	53.54	81.42	68.77	42.86	64.98	68.59	56.70
	Qwen3.5-4B (NoThink)	81.50	48.00	88.71	68.64	89.13	70.54	46.78	71.55	76.53	64.68
	Qwen3.5-9B (NoThink)	82.90	**53.00**	93.96	69.16	87.86	73.96	**53.35**	75.35	**79.52**	67.87
	Qwen3.5-27B (NoThink)	**84.10**	51.20	**94.23**	**72.18**	**89.38**	**78.00**	47.28	**77.12**	78.75	**69.62**
	Qwen3.5-0.8B (Think)	64.70	23.20	78.22	37.53	80.78	44.63	39.44	46.14	65.79	37.88
	Qwen3.5-2B (Think)	80.10	36.20	88.32	62.20	85.34	66.50	45.64	65.99	74.85	57.72
	Qwen3.5-4B (Think)	83.60	25.50	**92.91**	44.75	**87.99**	50.06	48.55	49.18	78.26	42.37
	Qwen3.5-9B (Think)	84.00	**67.80**	90.29	**81.89**	87.10	**91.53**	**52.34**	**89.51**	**78.43**	**82.68**
	Qwen3.5-27B (Think)	**84.20**	19.90	92.39	49.48	86.60	59.92	47.16	56.64	77.59	46.48
Qwen3.6	Qwen3.6-27B (NoThink)	83.90	**58.70**	94.36	**73.75**	88.37	**77.62**	47.41	**78.51**	78.51	**72.15**
	Qwen3.6-27B (Think)	**86.10**	20.90	**95.80**	45.54	**89.63**	54.74	**54.24**	52.59	**81.44**	43.44
Gemma	Gemma-2-2B-It	18.90	56.00	58.66	70.47	69.28	61.44	34.01	63.08	45.21	62.75
	Gemma-2-9B-It	75.00	29.10	88.58	57.48	**89.00**	53.98	**48.67**	53.73	75.31	48.57
	Gemma-2-27B-It	76.20	46.60	88.32	69.69	88.50	76.23	43.99	70.92	74.25	65.86
	Gemma-4-31B-it	**84.50**	**61.20**	**94.09**	**80.05**	85.46	**84.32**	42.10	**82.43**	**76.54**	**77.00**
Phi	Phi-3-mini-128k-instruct	49.00	**65.20**	54.33	**82.15**	77.37	**82.43**	28.19	**82.05**	52.22	**77.96**
	Phi-3.5-mini-instruct	56.10	44.70	60.37	74.15	81.04	69.91	32.62	70.42	57.53	64.79
	Phi-4-mini-instruct	65.50	61.30	69.55	77.43	**85.34**	76.74	**35.02**	74.08	63.85	72.39
	Phi-4	**79.70**	57.70	**85.83**	69.82	83.19	75.22	27.31	72.82	**69.00**	68.89
Olmo	Olmo-3-7B-Instruct	30.80	**80.40**	37.40	**93.31**	59.29	**89.63**	23.89	**87.74**	37.85	**87.77**
	Olmo-3.1-32B-Instruct	**66.90**	78.10	**79.66**	86.75	**82.68**	85.97	**40.58**	81.54	**67.46**	83.09

These two metrics are commonly adopted in existing instance-level evaluation protocols, which primarily focus on answer correctness and abstention accuracy under individual contextual conditions.

3.3 PFaithBench Construction

To systematically evaluate pairwise faithfulness, we construct PFaithBench $\mathcal{D} = \{(q, c^+, c^-, a)\}$, where each example (q, c^+, c^-, a) pairs a sufficient and an insufficient context for the same question. As illustrated in Fig. 2, we collect sufficient-context QA examples from three types of data sources: multi-hop, single-hop, and counterfactual QA. We then construct c^- by minimally substituting, removing, or blurring key information in c^+, followed by manual verification to ensure its quality.

Multi-hop factual QA. We collect (q, c^+, a) from HotpotQA (Yang et al., 2018) and 2WikiMultihopQA (Ho et al., 2020), where answering requires multiple supporting facts. Both datasets provide annotated supporting facts, allowing us to construct c^- by replacing one key supporting fact with an irrelevant fact, breaking the evidence dependency, resulting in the PFB-MultiHop subset.

Single-hop factual QA. We collect (q, c^+, a) from SQuAD (Rajpurkar et al., 2016). The insufficient context c^- is constructed using LLM rewriting to remove or blur key answer-supporting information while keeping the remaining context largely unchanged, resulting in the PFB-SingleHop subset.

Counterfactual QA. Counterfactual QA refers to questions that require answering based on contexts that contradict real-world knowledge (Ming et al., 2025). We collect $(q, c^+_{\mathrm{ori}}, a_{\mathrm{ori}}, c^+_{\mathrm{cf}}, a_{\mathrm{cf}})$ from ConFiQA (Bi et al., 2025), which provides paired original and counterfactual contexts for the same question q. We treat both c^+_{ori} and c^+_{cf} as sufficient contexts and construct their corresponding insufficient contexts by removing or blurring only the key answer-supporting information while preserving the remaining content, resulting in the PFB-ConOri and PFB-ConCF subsets.

For each QA type, we begin with 1,000 examples of the form (q, c^+, a). For multi-hop QA, annotated supporting facts allow us to construct insufficient contexts by replacing a key fact with an irrelevant one. For single-hop and counterfactual QA, two postgraduate students manually verify the LLM-

Table 4: **Impact of context coupling on faithfulness training.** The best value of each metric is **bold**. Among the three training settings, lightblue marks both AnsAcc and AbsAcc for the smallest $|\text{AnsAcc} - \text{AbsAcc}|$, and marks both OverAns and OverAbs for the smallest $|\text{OverAns} - \text{OverAbs}|$.

LLM	Setting	PFB-MultiHop						PFB-SingleHop					
		DFAcc ↑	SFAcc ↑	AnsAcc ↑	AbsAcc ↑	OverAns ↓	OverAbs ↓	DFAcc ↑	SFAcc ↑	AnsAcc ↑	AbsAcc ↑	OverAns ↓	OverAbs ↓
Llama-3.1-8B-Instruct	Base Model	37.70	30.00	73.80	43.40	56.60	6.50	49.87	39.50	73.36	54.59	45.41	6.04
	MHSuff-SHInsuff	34.20	27.40	**81.80**	36.60	63.40	**2.60**	37.14	31.50	32.28	**98.56**	**1.44**	61.68
	SHSuff-MHInsuff	2.70	0.00	0.00	**96.70**	**3.30**	95.60	16.01	14.70	**91.60**	16.80	83.20	**1.05**
	PairContext	**71.50**	**61.00**	76.40	81.90	18.10	10.60	**79.13**	**75.33**	86.22	86.35	13.65	7.74
Llama-3.2-3B-Instruct	Base Model	42.20	26.00	45.90	62.60	37.40	23.90	46.19	25.85	32.28	87.01	12.99	43.18
	MHSuff-SHInsuff	33.80	26.00	**73.00**	38.20	61.80	**6.30**	26.90	23.10	24.93	**97.51**	**2.49**	70.87
	SHSuff-MHInsuff	45.00	10.90	18.70	79.60	20.40	37.00	36.35	33.07	**82.55**	39.90	60.10	**4.20**
	PairContext	**68.20**	**53.40**	69.40	**80.00**	**20.00**	12.50	**67.45**	**62.86**	79.27	79.40	20.60	12.47
Qwen3-4B (NoThink)	Base Model	44.10	32.90	56.80	67.10	32.90	25.80	63.12	56.56	73.10	79.66	20.34	18.24
	MHSuff-SHInsuff	56.10	45.50	**74.20**	64.00	36.00	**8.70**	67.72	56.30	61.02	**93.44**	**6.56**	26.51
	SHSuff-MHInsuff	58.40	47.30	73.30	67.10	32.90	9.60	69.42	63.78	**86.48**	73.62	26.38	**5.25**
	PairContext	**68.80**	**60.50**	74.10	**84.20**	**15.80**	16.00	**77.95**	**74.15**	81.63	90.29	9.71	12.60
Qwen3-1.7B (NoThink)	Base Model	28.40	11.40	38.20	43.00	57.00	17.90	41.47	26.90	58.79	50.52	49.48	10.76
	MHSuff-SHInsuff	24.10	14.90	**61.10**	27.20	72.80	**4.20**	43.57	19.16	21.26	**93.96**	**6.04**	51.18
	SHSuff-MHInsuff	32.20	13.90	26.30	75.60	24.40	46.70	38.58	31.50	**74.54**	42.65	57.35	**5.12**
	PairContext	**61.50**	**43.70**	58.50	**77.20**	**22.80**	17.80	**58.01**	**50.52**	65.88	79.53	20.47	23.62

LLM	Setting	PFB-ConOri						PFB-ConCF					
		DFAcc ↑	SFAcc ↑	AnsAcc ↑	AbsAcc ↑	OverAns ↓	OverAbs ↓	DFAcc ↑	SFAcc ↑	AnsAcc ↑	AbsAcc ↑	OverAns ↓	OverAbs ↓
Llama-3.1-8B-Instruct	Base Model	50.95	45.13	84.45	54.11	45.89	3.79	42.10	19.09	29.58	52.97	47.03	18.46
	MHSuff-SHInsuff	66.75	61.44	66.62	**94.06**	**5.94**	27.69	42.98	24.78	25.92	**94.94**	**5.06**	53.86
	SHSuff-MHInsuff	15.80	14.29	**91.40**	16.81	83.19	**1.14**	12.39	7.59	**52.72**	14.79	85.21	**3.79**
	PairContext	**80.40**	**75.98**	85.84	89.00	11.00	9.10	**70.16**	**42.35**	47.03	87.74	12.26	19.22
Llama-3.2-3B-Instruct	Base Model	59.29	43.87	57.65	80.53	19.47	23.51	47.91	15.93	18.20	78.89	21.11	37.42
	MHSuff-SHInsuff	65.36	52.72	57.77	**93.30**	**6.70**	28.45	48.17	19.97	20.73	**95.07**	**4.93**	49.30
	SHSuff-MHInsuff	52.84	49.18	**88.24**	57.14	42.86	**4.80**	39.19	23.14	38.56	53.35	46.65	**21.74**
	PairContext	**78.00**	**73.32**	85.71	86.22	13.78	8.22	**63.46**	**40.83**	**46.40**	84.70	15.30	24.40
Qwen3-4B (NoThink)	Base Model	63.72	59.29	85.59	71.81	28.19	8.22	55.12	29.58	36.92	72.06	27.94	22.12
	MHSuff-SHInsuff	76.36	65.99	76.36	86.73	13.27	11.38	63.08	32.49	35.90	87.61	12.39	27.18
	SHSuff-MHInsuff	70.80	65.87	**86.09**	77.75	22.25	**6.95**	63.97	34.39	**41.09**	77.37	22.63	**15.68**
	PairContext	**80.66**	**75.60**	81.29	**94.18**	**5.82**	13.65	**68.65**	**39.70**	40.71	**94.56**	**5.44**	28.07
Qwen3-1.7B (NoThink)	Base Model	52.47	41.97	76.61	56.89	43.11	**5.69**	49.30	24.78	40.20	56.64	43.36	**8.98**
	MHSuff-SHInsuff	67.89	44.75	49.94	**91.40**	**8.60**	23.89	61.44	20.86	22.63	**91.15**	**8.85**	30.47
	SHSuff-MHInsuff	50.82	46.90	**84.20**	57.27	42.73	7.21	50.06	27.69	**47.16**	56.01	43.99	**8.98**
	PairContext	**68.39**	**62.83**	79.39	81.42	18.58	13.40	**61.69**	**34.26**	41.72	80.66	19.34	20.99

rewritten insufficient contexts, yielding 762 and 791 valid examples, respectively. The higher retention rates 76.2% and 79.1% demonstrate the effectiveness of our construction pipeline. PFaithBench thus comprises four evaluation subsets: PFB-MultiHop, PFB-SingleHop, PFB-ConOri, and PFB-ConCF, with 1,000, 762, 791 and 791 examples, respectively. Additionally, we use an LLM judge to classify model responses as $\{\text{IDK}, \text{C}, \text{W}\}$ for metric computation. More details refer to Appendix A.

3.4 Evaluation Results

A wide range of existing LLMs is taken into the evaluation, covering seven families with thirty-nine models across parameters from 0.6 billion to 70 billion. All models are explicitly prompted to answer solely based on the provided context. When the context is sufficient, the model is required to provide an answer along with its reasoning process; otherwise, the model should acknowledge the lack of information and respond "I don't know". The detailed evaluation setup is provided in Appendix A.5.

Most existing models exhibit a strong bias toward over-answering. As shown in Table 1, the majority of models $35/39$ demonstrate substantially higher OverAns than OverAbs, and only a few models, such as Olmo-3-7B, show the over-abstention-dominant behavior. indicating a tendency that models produce answers even when the context is insufficient. One possible reason is existing outcome-based post-training, such as reinforcement learning with verifiable rewards (RLVR) (Yang et al., 2026). Such methods primarily reward correct answers and penalize incorrect ones, while often treating appropriate abstention as an incorrect response. This may encourage models to attempt an answer even when the context does not include sufficient information (Yao et al., 2025).

Over-answering is closely associated with internal knowledge leakage. Figure 3 shows that all models exhibit some degree of internal knowledge leakage: when the provided context is insufficient, they draw on their internal knowledge to produce an answer. This may partly explain their tendency to over-answer. Within most model families, leakage also increases with model size, suggesting that larger models are more likely to rely on internal knowledge when context is insufficient. Although this behavior may make models appear more helpful, it also increases the risk of unsupported answers in information-sensitive settings, particularly when their internal knowledge is outdated or incorrect.

Models exhibit imbalanced and asymmetric context-sensitive behavior. Tables 3 and 2 report the marginal AnsAcc and AbsAcc and joint DFAcc and SFAcc metrics, respectively. Many models show a substantial gap between AnsAcc and AbsAcc, indicating an imbalance between answering and

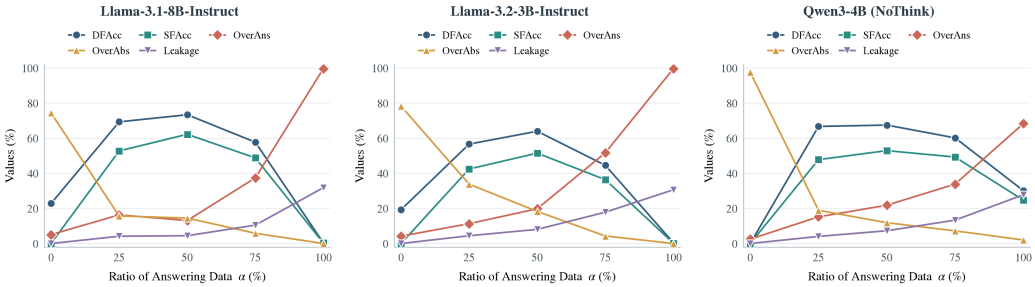

Figure 4: **Impact of answering–abstention data composition (asymmetric reward).** Faithfulness training under different answering data ratios $\alpha \in \{0, 25, 50, 75, 100\}$ across three LLMs. This is the average value across four subsets in PFaithBench. More results are provided in Appendix C.2.

abstention. Their joint performance is also weaker than their marginal performance, suggesting that strong accuracy under individual context conditions may not translate into pairwise faithful behavior.

4 Data-centric Analyses for Faithfulness Training

PFaithBench pairs sufficient and insufficient contexts (c^+, c^-) for each question, providing a natural setting to examine how data composition affects faithfulness. In this section, we investigate two factors: the coupling between sufficient and insufficient contexts, and their relative proportions.

4.1 Faithfulness Training Scheme

Existing faithfulness training has explored supervised fine-tuning (SFT) (Zhang et al., 2024a), direct preference optimization (DPO) (Bi et al., 2025), and reinforcement learning (RL) (Gui et al., 2026; Ren et al., 2025), primarily focusing on response correctness under given contexts. In our pairwise faithfulness setting, model responses are explicitly categorized into $\{\mathrm{C}, \mathrm{W}, \mathrm{IDK}\}$, enabling explicit rewards for answering and abstention. We therefore adopt GRPO (Shao et al., 2024) as the training algorithm, and define rewards that penalize both over-answering and over-abstention:

$$r(a, y^-) = \begin{cases} +1 & \text{if } y^- = \mathrm{IDK}, \\ -1 & \text{if } y^- \in \{\mathrm{C}, \mathrm{W}\}, \end{cases} \qquad r_{\mathrm{asym}}(a, y^+) = \begin{cases} +1 & \text{if } y^+ = \mathrm{C}, \\ 0 & \text{if } y^+ = \mathrm{W}, \\ -1 & \text{if } y^+ = \mathrm{IDK}, \end{cases} \tag{6}$$

Under sufficient context (c^+), this reward favors answering over abstaining, even when the answer is incorrect, while assigning the highest reward to a correct answer. Under insufficient context (c^-), it rewards abstention and penalizes answering. The subsequent analyses are conducted on this training scheme and training data is constructed using the same pipeline as in PFaithBench. We also examine a symmetric reward that penalizes incorrect answers and abstention equally under c^+ in Appendix C.2.

4.2 Confounding Effects from Mismatched Context Coupling

We investigate how the coupling between sufficient and insufficient contexts affects faithfulness training, particularly when these contexts are drawn from different data sources. The question is whether models learn to rely on contextual sufficiency or instead exploit dataset-specific shortcuts.

- **Mismatched context coupling:** Sufficient and insufficient examples come from different QA sources, $(q_A, c^+) \sim \mathcal{D}_{\mathrm{A}}$ and $(q_B, c^-) \sim \mathcal{D}_{\mathrm{B}}$. Thus, each question appears with only one context type, making sufficiency correlated with data source. We consider two settings: MHSuff-SHInsuff (multi-hop sufficient, single-hop insufficient) and SHSuff-MHInsuff (the reverse).
- Both contexts are observed for each question, $(q, c^+, c^-) \sim \mathcal{D}$, preserving their pairing. We sample these pairs from a balanced pool of multi-hop and single-hop QA data, denoted PairContext.

The experiments are conducted across four models of varying scales: Llama-3.1-8B-Instruct, Llama-3.2-3B-Instruct, and Qwen3-4B/1.7B. Table 4 reports the results across four subsets of PFaithBench.

Table 5: **Impact of faithfulness training on general model capabilities.**

Setting	AMC	AIME24	AIME25	MATH500	MMLU-Pro	GPQA-D	IFEval	Setting	AMC	AIME24	AIME25	MATH500	MMLU-Pro	GPQA-D	IFEval
Llama-3.1-8B-Instruct								$\alpha = 0$	16.42	3.02	0.42	31.80	34.21	26.26	78.82
Base Model	14.76	2.19	0.63	33.40	36.10	29.29	76.86	$\alpha = 25$	17.32	3.33	0.21	34.60	36.43	30.81	77.42
MHSuff-SHInsuff	19.28	3.96	0.52	38.60	35.61	25.76	80.34	$\alpha = 50$	18.37	3.33	0.52	34.60	36.18	25.25	78.87
SHSuff-MHInsuff	18.52	2.71	0.31	36.20	35.46	35.35	79.53	$\alpha = 75$	17.62	3.23	0.31	37.60	35.85	27.78	78.53
PairContext	19.43	4.17	0.31	38.00	36.58	30.30	77.81	$\alpha = 100$	18.52	4.17	0.31	36.80	36.58	27.78	78.39
Llama-3.2-3B-Instruct								$\alpha = 0$	15.06	5.42	0.10	34.60	28.99	26.77	74.14
Base Model	17.17	4.06	0.21	35.00	29.73	31.82	72.97	$\alpha = 25$	15.06	5.63	0.42	35.00	29.31	28.79	74.79
MHSuff-SHInsuff	16.87	5.94	0.31	36.40	29.93	26.77	73.61	$\alpha = 50$	18.07	4.90	0.42	36.20	29.34	24.24	72.53
SHSuff-MHInsuff	17.77	4.69	0.83	37.60	29.13	27.27	71.72	$\alpha = 75$	17.77	4.17	0.21	36.60	29.73	28.79	74.28
PairContext	17.77	5.42	70.42	35.00	29.08	25.25	74.83	$\alpha = 100$	18.98	5.21	0.42	34.60	30.27	26.77	73.49
Qwen3-4B (NoThink)								$\alpha = 0$	58.43	21.88	17.40	80.80	30.94	47.98	84.59
Base Model	59.04	21.98	18.23	79.80	34.88	44.44	82.79	$\alpha = 25$	57.83	22.29	18.02	80.00	32.75	46.46	84.16
MHSuff-SHInsuff	59.19	22.40	17.50	79.40	34.72	42.93	84.38	$\alpha = 50$	58.43	22.19	18.54	81.00	32.26	43.94	83.42
SHSuff-MHInsuff	57.98	23.44	18.54	81.00	32.77	44.44	84.06	$\alpha = 75$	57.23	22.08	18.23	81.40	35.41	42.42	84.42
PairContext	57.23	21.46	17.81	81.40	31.02	41.41	83.64	$\alpha = 100$	57.38	22.60	16.98	83.80	35.66	41.41	84.57

Mismatched context coupling may induce data-specific shortcuts. As shown in Table 4, mismatched training produces behavior aligned with data source. For example, Llama-3.1-8B-Instruct trained with MHSuff-SHInsuff shows high OverAns on PFB-MultiHop (63.40) and high OverAbs on PFB-SingleHop (61.68). Reversing the sources leads to high OverAbs on PFB-MultiHop (95.60) and high OverAns on PFB-SingleHop (83.20). This pattern suggests that models may use source-specific cues to decide whether to answer. PairContext achieves the best pairwise faithfulness in most cases, indicating the importance of preserving context pairs during faithfulness training.

Mismatched context coupling may exacerbate the imbalance between answering and abstention. From the Table 4, we observe that mismatched context coupling can produce extreme imbalances between AnsAcc and AbsAcc, with one becoming very high while the other remains very low. For example, under MHSuff-SHInsuff, Llama-3.1-8B-Instruct achieves high AnsAcc but low AbsAcc on PFB-MultiHop (81.80 vs. 36.60), while showing the reverse on PFB-SingleHop (32.28 vs. 98.56). This suggests that mismatched training may induce a source-conditioned answer-abstention bias.

4.3 Impact of Answering-Abstention Data Proportion

Beyond context coupling, we further investigate how the proportion of sufficient-context (answering) and insufficient-context (abstaining) data affects faithfulness training. Specifically, we vary the percentage of answering data as $\alpha \in \{0, 25, 50, 75, 100\}$ from abstention-only ($\alpha = 0$) to answering-only ($\alpha = 100$). We conduct training under these different data compositions across three models: Llama-3.1-8B-Instruct, Llama-3.2-3B-Instruct, and Qwen3-4B (NoThink).

Data composition shapes the balance between answering and abstention. Figure 4 shows that increasing α generally raises OverAns and Leakage while reducing OverAbs. At $\alpha = 0$, abstention dominates and SFAcc falls to zero; at $\alpha = 100$, over-answering becomes severe, reaching 100% for both Llama models. Across all three models, DFAcc and SFAcc peak near the balanced setting ($\alpha = 50$). These results suggest that a balanced mix of sufficient- and insufficient-context data helps faithfulness training improve models' ability to switch between answering and abstaining.

Faithfulness training does not affect general model capabilities. Table 5 further demonstrate whether faithfulness training affects models' general capabilities. Specifically, we evaluate the base and faithfulness-trained models across diverse benchmarks, covering mathematical reasoning AMC (Yang et al., 2024a), AIME24 (Zhang & Math-AI, 2024), AIME25 (Zhang & Math-AI, 2025), and MATH500 (Lightman et al., 2024); multi-domain knowledge MMLU-Pro (Wang et al., 2024) and GPQA-Diamond (Rein et al., 2023); and instruction-following IFEval (Zhou et al., 2023). We observe that faithfulness-trained models retain performance comparable to their untrained base model. These results suggest that the changes in pairwise faithfulness need not come at the expense of general capabilities, supporting the study of faithfulness as a distinct training capability.

5 Conclusion

In this work, we revisit faithfulness in LLMs from a pairwise context-sensitive perspective and introduce PFaithBench, a benchmark that evaluates whether models can adapt their behavior for the same question under sufficient and insufficient contexts. Unlike instance-level evaluation, our framework explicitly measures the ability of models to switch between answering and abstaining based on contextual sufficiency. Our evaluation identifies over-answering and over-abstention as the main faithfulness errors, with existing models biased toward over-answering, which is closely associated with internal knowledge leakage. Our data-centric analyses show that context coupling and the balance of answering and abstention data matter for faithfulness training.

AI USE STATEMENT

In this work, we used generative AI tools to aid or polish writing and to support research ideation and execution. All AI-assisted outputs were carefully reviewed and validated by the authors. In particular, writing suggestions were manually revised, and research-related ideas or execution suggestions were strictly assessed before being incorporated into the work. We take responsibility for the final content of this work, including text, claims, and artifacts produced with the aid of generative AI.

ETHICS STATEMENT

This work studies the faithfulness of large language models under varying contextual conditions and does not involve human-subject experiments or the collection of new personally identifiable or sensitive information. PFaithBench is constructed entirely from publicly available, open-source question-answering datasets, by modifying contextual information to create paired sufficient and insufficient contexts, with manual verification applied to ensure data quality. Our study is intended to improve the evaluation and understanding of context-grounded model behavior. We do not identify additional ethical concerns specific to the methodology beyond those generally associated with the development and evaluation of large language models. We have conducted this work in accordance with standard research integrity practices and the ICLR Code of Ethics.

REPRODUCIBILITY STATEMENT

We have made substantial efforts to ensure the reproducibility of our results. We provide necessary details of the PFaithBench construction and prompt design, as well as the experimental settings. We also release the source code and data at https://anonymous.4open.science/r/PFaithBench_Submission-ICLR. Together, the data construction, experimental settings, source code, and data provide the necessary resources for reproducing the results in this work.

REFERENCES

Marah Abdin, Jyoti Aneja, Hany Awadalla, Ahmed Awadallah, Ammar Ahmad Awan, Nguyen Bach, Amit Bahree, Arash Bakhtiari, Jianmin Bao, Harkirat Behl, Alon Benhaim, et al. Phi-3 technical report: A highly capable language model locally on your phone. *arXiv preprint arXiv:2404.14219*, 2024a.

Marah Abdin, Jyoti Aneja, Harkirat Behl, Sébastien Bubeck, Ronen Eldan, Suriya Gunasekar, Michael Harrison, Russell J Hewett, Mojan Javaheripi, Piero Kauffmann, et al. Phi-4 technical report. *arXiv preprint arXiv:2412.08905*, 2024b.

Abdelrahman Abouelenin, Atabak Ashfaq, Adam Atkinson, Hany Awadalla, Nguyen Bach, Jianmin Bao, Alon Benhaim, Martin Cai, Vishrav Chaudhary, Congcong Chen, et al. Phi-4-mini technical report: Compact yet powerful multimodal language models via mixture-of-loras. *arXiv preprint arXiv:2503.01743*, 2025.

Vaibhav Adlakha, Parishad BehnamGhader, Xing Han Lu, Nicholas Meade, and Siva Reddy. Evaluating correctness and faithfulness of instruction-following models for question answering. *Transactions of the Association for Computational Linguistics*, 12:681–699, 2024.

Hanoz Bhathena, Parin Rajesh Jhaveri, Rohan Mittal, Prateek Singh, Aymen Kallala, Rachneet Kaur, Yiqiao Jin, Zhen Zeng, Adwait Ratnaparkhi, and Denis Kochedykov. Mm-bizrag: Rethinking multimodal retrieval-augmented generation for general purpose enterprise q&a. In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (ACL 2026)*, pp. 1967–2000, 2026.

Baolong Bi, Shaohan Huang, Yiwei Wang, Tianchi Yang, Zihan Zhang, Haizhen Huang, Lingrui Mei, Junfeng Fang, Zehao Li, Furu Wei, et al. Context-dpo: Aligning language models for context-faithfulness. In *Findings of the Association for Computational Linguistics: ACL 2025*, pp. 10280–10300, 2025.

Jiawei Chen, Hongyu Lin, Xianpei Han, and Le Sun. Benchmarking large language models in retrieval-augmented generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 17754–17762, 2024.

Wenhu Chen, Xinyi Wang, and William Yang Wang. A dataset for answering time-sensitive questions. *arXiv preprint arXiv:2108.06314*, 2021.

Qinyuan Cheng, Tianxiang Sun, Xiangyang Liu, Wenwei Zhang, Zhangyue Yin, Shimin Li, Linyang Li, Zhengfu He, Kai Chen, and Xipeng Qiu. Can ai assistants know what they don't know? *arXiv preprint arXiv:2401.13275*, 2024.

Darren Edge, Ha Trinh, Newman Cheng, Joshua Bradley, Alex Chao, Apurva Mody, Steven Truitt, Dasha Metropolitansky, Robert Osazuwa Ness, and Jonathan Larson. From local to global: A graph rag approach to query-focused summarization. *arXiv preprint arXiv:2404.16130*, 2024.

Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yixin Dai, Jiawei Sun, Haofen Wang, Haofen Wang, et al. Retrieval-augmented generation for large language models: A survey. *arXiv preprint arXiv:2312.10997*, 2(1):32, 2023.

Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.

Runquan Gui, Yafu Li, Xiaoye Qu, Ziyan Liu, Yeqiu Cheng, and Yu Cheng. Learning to reason faithfully through step-level faithfulness maximization. *arXiv preprint arXiv:2602.03507*, 2026.

Xanh Ho, Anh-Khoa Duong Nguyen, Saku Sugawara, and Akiko Aizawa. Constructing a multi-hop qa dataset for comprehensive evaluation of reasoning steps. In *Proceedings of the 28th International Conference on Computational Linguistics*, pp. 6609–6625, 2020.

Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, et al. A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions. *ACM Transactions on Information Systems*, 43(2):1–55, 2025.

Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. Survey of hallucination in natural language generation. *ACM computing surveys*, 55(12):1–38, 2023.

Hailey Joren, Jianyi Zhang, Chun-Sung Ferng, Da-Cheng Juan, Ankur Taly, and Cyrus Rashtchian. Sufficient context: A new lens on retrieval augmented generation systems. In *International Conference on Learning Representations*, volume 2025, pp. 20310–20334, 2025.

Mandar Joshi, Eunsol Choi, Daniel S Weld, and Luke Zettlemoyer. Triviaqa: A large scale distantly supervised challenge dataset for reading comprehension. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 1601–1611, 2017.

Rishi Kalra, Zekun Wu, Ayesha Gulley, Airlie Hilliard, Xin Guan, Adriano Koshiyama, and Philip Colin Treleaven. Hypa-rag: A hybrid parameter adaptive retrieval-augmented generation system for ai legal and policy applications. In *Proceedings of the 1st workshop on customizable nlp: Progress and challenges in customizing nlp for a domain, application, group, or individual (customnlp4u)*, pp. 237–256, 2024.

Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, et al. Natural questions: a benchmark for question answering research. *Transactions of the Association for Computational Linguistics*, 2019.

Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in neural information processing systems*, 33: 9459–9474, 2020.

Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let's verify step by step. In *ICLR*, 2024.

Xi Victoria Lin, Xilun Chen, Mingda Chen, Weijia Shi, Maria Lomeli, Rich James, Pedro Rodriguez, Jacob Kahn, Gergely Szilvasy, Mike Lewis, et al. Ra-dit: Retrieval-augmented dual instruction tuning. *arXiv preprint arXiv:2310.01352*, 2023.

Alex Mallen, Akari Asai, Victor Zhong, Rajarshi Das, Daniel Khashabi, and Hannaneh Hajishirzi. When not to trust language models: Investigating effectiveness of parametric and non-parametric memories. In *Proceedings of the 61st annual meeting of the association for computational linguistics (volume 1: Long papers)*, pp. 9802–9822, 2023.

Yifei Ming, Senthil Purushwalkam, Shrey Pandit, Zixuan Ke, Xuan-Phi Nguyen, Caiming Xiong, and Shafiq Joty. Faitheval: Can your language model stay faithful to context, even if" the moon is made of marshmallows". In *The Thirteenth International Conference on Learning Representations*, 2025.

Shuo Nie, Hexuan Deng, Chao Wang, Ruiyu Fang, Xuebo Liu, Shuangyong Song, Yu Li, Min Zhang, and Xuelong Li. Stop rewarding hallucinated steps: Faithfulness-aware step-level reinforcement learning for small reasoning models. *arXiv preprint arXiv:2602.05897*, 2026.

Team Olmo, Allyson Ettinger, Amanda Bertsch, Bailey Kuehl, David Graham, David Heineman, Dirk Groeneveld, Faeze Brahman, Finbarr Timbers, Hamish Ivison, et al. Olmo 3. *arXiv preprint arXiv:2512.13961*, 2025.

Qwen Team. Qwen3.5: Towards native multimodal agents, February 2026a. URL https://qwen.ai/blog?id=qwen3.5.

Qwen Team. Qwen3.6-27B: Flagship-level coding in a 27B dense model, April 2026b. URL https://qwen.ai/blog?id=qwen3.6-27b.

Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. Squad: 100,000+ questions for machine comprehension of text. In *Proceedings of the 2016 conference on empirical methods in natural language processing*, pp. 2383–2392, 2016.

David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani, Julian Michael, and Samuel R Bowman. Gpqa: A graduate-level google-proof q&a benchmark. *arXiv preprint arXiv:2311.12022*, 2023.

Baochang Ren, Shuofei Qiao, Da Zheng, Huajun Chen, and Ningyu Zhang. Knowrl: Exploring knowledgeable reinforcement learning for factuality. *arXiv preprint arXiv:2506.19807*, 2025.

Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Yang Wu, et al. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024.

Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, et al. Gemma 2: Improving open language models at a practical size. *arXiv preprint arXiv:2408.00118*, 2024.

Harsh Trivedi, Niranjan Balasubramanian, Tushar Khot, and Ashish Sabharwal. Musique: Multihop questions via single-hop question composition. *Transactions of the Association for Computational Linguistics*, 10:539–554, 2022.

Yubo Wang, Xueguang Ma, Ge Zhang, Yuansheng Ni, Abhranil Chandra, Shiguang Guo, Weiming Ren, Aaran Arulraj, Xuan He, Ziyan Jiang, et al. Mmlu-pro: A more robust and challenging multi-task language understanding benchmark. *Advances in Neural Information Processing Systems*, 37: 95266–95290, 2024.

Kevin Wu, Eric Wu, and James Zou. Clasheval: Quantifying the tug-of-war between an llm's internal prior and external evidence. *Advances in neural information processing systems*, 37:33402–33422, 2024.

An Yang, Beichen Zhang, Binyuan Hui, Bofei Gao, Bowen Yu, Chengpeng Li, Dayiheng Liu, Jianhong Tu, Jingren Zhou, Junyang Lin, et al. Qwen2.5-math technical report: Toward mathematical expert model via self-improvement. *arXiv preprint arXiv:2409.12122*, 2024a.

An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*, 2025.

Yuqing Yang, Ethan Chern, Xipeng Qiu, Graham Neubig, and Pengfei Liu. Alignment for honesty. *Advances in Neural Information Processing Systems*, 37:63565–63598, 2024b.

Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D Manning. Hotpotqa: A dataset for diverse, explainable multi-hop question answering. In *Proceedings of the 2018 conference on empirical methods in natural language processing*, pp. 2369–2380, 2018.

Zhiqin Yang, Jingwen Fu, Yuhan Liu, Hengyu Liu, Yonggang Zhang, Kainan Cao, Zizhuo Zhang, Chenxin Li, Ruibin Yuan, Jiahao Pan, et al. Scaling large reasoning models beyond human supervision: A path toward superintelligence. *arXiv preprint arXiv:2608.31075*, 2026.

Zijun Yao, Yantao Liu, Yanxu Chen, Jianhui Chen, Junfeng Fang, Lei Hou, Juanzi Li, and Tat-Seng Chua. Are reasoning models more prone to hallucination? *arXiv preprint arXiv:2505.23646*, 2025.

Hanning Zhang, Shizhe Diao, Yong Lin, Yi Fung, Qing Lian, Xingyao Wang, Yangyi Chen, Heng Ji, and Tong Zhang. R-tuning: Instructing large language models to say 'i don't know'. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pp. 7113–7139, 2024a.

Tianjun Zhang, Shishir G Patil, Naman Jain, Sheng Shen, Matei Zaharia, Ion Stoica, and Joseph E Gonzalez. Raft: Adapting language model to domain specific rag. *arXiv preprint arXiv:2403.10131*, 2024b.

Yifan Zhang and Team Math-AI. American invitational mathematics examination (aime) 2024, 2024.

Yifan Zhang and Team Math-AI. American invitational mathematics examination (aime) 2025, 2025.

Jeffrey Zhou, Tianjian Lu, Swaroop Mishra, Siddhartha Brahma, Sujoy Basu, Yi Luan, Denny Zhou, and Le Hou. Instruction-following evaluation for large language models. *arXiv preprint arXiv:2311.07911*, 2023.

A More Details of PFaithBench

A.1 Details of Data Sources

The construction of our PFaithBench is based on four publicly available context-based QA datasets: HotpotQA, 2WikiMultiHopQA, SQuAD, and ConFiQA, covering diverse contextual conditions, including multi-hop reasoning, single-hop reading comprehension, and counterfactual contexts, enabling a comprehensive evaluation of pairwise faithfulness behaviors.

HotpotQA (Yang et al., 2018) is a large-scale multi-hop factual QA dataset that requires models to reason over multiple supporting passages to derive the correct answer. Each question is paired with a set of Wikipedia passages, among which two or more contain the necessary supporting facts. Although HotpotQA provides sentence-level supporting fact annotations, we use passage-level contexts to preserve the difficulty of context-based reasoning. For each question, we construct the insufficient context c^- by randomly replacing one of the supporting passages with an irrelevant passage. This modification breaks the multi-hop dependency while preserving the remaining contextual information.

2WikiMultihopQA (Ho et al., 2020) is another multi-hop QA dataset constructed from Wikipedia and Wikidata. Compared with HotpotQA, 2WikiMultihopQA emphasizes structured multi-hop reasoning, including comparison, inference, and compositional questions, with explicit reasoning paths derived from knowledge triples. Each question is paired with multiple supporting passages that contain the necessary evidence. We adopt the same strategy as in HotpotQA to construct insufficient contexts by randomly replacing one of the supporting passages with an irrelevant one, thereby breaking the multi-hop reasoning chain while preserving the most of contextual information.

SQuAD (Rajpurkar et al., 2016) (Stanford Question Answering Dataset) is a widely used single-hop reading comprehension benchmark, in which each question is grounded in a single Wikipedia passage and the answer is a contiguous span within the passage. SQuAD primarily evaluates a model's ability to understand and extract relevant information from a given context. In our work, we use SQuAD as the source of single-hop QA data and construct insufficient contexts by rewriting the original passage to remove or blur the answer-supporting span while preserving most of the surrounding content.

ConFiQA (Bi et al., 2025) is a counterfactual question answering benchmark designed to evaluate the ability of LLMs to faithfully follow the provided context, even when the context contradicts real-world knowledge. Each question in ConFiQA is associated with both an original context that aligns with real-world knowledge and a counterfactual context that introduces fictitious knowledge, along with their corresponding answers. This paired design makes ConFiQA particularly suitable for studying faithfulness. In our work, we treat both the original and counterfactual contexts as sufficient contexts, and construct corresponding insufficient contexts by removing or blurring the key information through LLM-based rewriting. This results in two evaluation subsets, PFB-ConOri and PFB-ConCF, where c^+ corresponds to original and counterfactual contexts respectively.

A.2 Statistics of PFaithBench

Table 7 summarizes statistics of PFaithBench. We additionally construct training sets for PFB-MultiHop and PFB-SingleHop for our data-centric analyses of faithfulness training in Sec. 4. All evaluation samples undergo human verification. Training samples follow the same construction pipeline but are not manually verified. PFB-ConOri and PFB-ConCF are used solely for evaluation.

A.3 Relationship between PFaithBench and Existing Benchmarks

Table 6 summarizes the relationship between PFaithBench and existing benchmarks from the perspectives of evaluation paradigm, context setting, and reasoning type. Existing benchmarks cover different aspects relevant to contextual faithfulness, including answerable, unanswerable, and counterfactual settings, as well as single-hop and multi-hop reasoning. However, they primarily evaluate each instance independently. In contrast, PFaithBench pairs sufficient and insufficient contexts for the same question, enabling pairwise evaluation of hether a model appropriately answers with sufficient context while abstaining with insufficient context. Meanwhile, PFaithBench covers answerable, unanswerable, and counterfactual settings across both single-hop and multi-hop reasoning, providing a unified benchmark for evaluating pairwise faithfulness under diverse contextual conditions.

Table 6: **Relationship between PFaithBench and existing benchmarks.**

Benchmark	Evaluation Paradigm	Answerable	Unanswerable	Counterfactual	Reasoning Type
FaithEval (Ming et al., 2025)	Instance-level	✓	✓	✓	Single-hop
ConFiQA (Bi et al., 2025)	Instance-level	✓		✓	Single-hop
TimeQA (Chen et al., 2021)	Instance-level	✓	✓		Multi-hop
RGB (Chen et al., 2024)	Instance-level	✓			Single-hop
ClashEval (Wu et al., 2024)	Instance-level	✓			Multi-hop
SQuAD (Rajpurkar et al., 2016)	Instance-level	✓			Single-hop
NQ (Kwiatkowski et al., 2019)	Instance-level	✓			Single-hop
TriviaQA (Joshi et al., 2017)	Instance-level	✓			Single-hop
PopQA (Mallen et al., 2023)	Instance-level	✓			Single-hop
MuSiQue (Trivedi et al., 2022)	Instance-level	✓			Multi-hop
2WikiMultihopQA (Ho et al., 2020)	Instance-level	✓			Multi-hop
HotpotQA (Yang et al., 2018)	Instance-level	✓			Multi-hop
PFaithBench	**Pairwise**	✓	✓	✓	**Both**

A.4 More Details of Constructing Insufficient Context c^-

For multi-hop QA, HotpotQA and 2WikiMultihopQA provide annotated supporting facts. We construct c^- by replacing one passage containing a key supporting fact with an irrelevant passage, thereby breaking the evidence needed to answer the question. For single-hop and counterfactual QA, we use Qwen3-32B (Yang et al., 2025) to rewrite sufficient contexts c^+ into insufficient contexts c^-. Two postgraduate students then verify the evaluation examples, retaining a rewrite only if the resulting context lacks sufficient evidence to answer the question. Starting from 1,000 examples for each QA type, this process retains 762 of 1,000 single-hop examples (76.2%) and 791 of 1,000 counterfactual examples (79.1%). The prompt used to rewrite the contexts is shown below:

Prompt for rewriting insufficient contexts

```
You will be provided with a context passage, a question, and an answer.
    The answer can be deduced from the given context. Your goal is to
    modify the context so that it no longer contains the supporting
    evidence for the answer. You should review the context sentence by
    sentence. For each sentence, consider the following two cases:
1. Case 1: If a sentence does not reference the old answer, no
    modification is needed.
2. Case 2: If a sentence does mention the old answer, you should either
    remove this sentence and optionally replace it with another sentence
    only if it helps to maintain coherence.

The modified context should:
1. Be coherent with the original context.
2. Not support the original answer.

Ideally, the majority of the original context should remain unchanged.
    You will output a JSON object containing the following information:

{{
  "question": string, // The original question.
  "old answer": string, // The original answer.
  "modified context": string, // The modified context.
  "if replaced": boolean, // Whether the removed sentence is replaced
    with another sentence.
  "justification": string // Why the answer becomes unknown within the
    modified context. The justification should be concise.
}}

Only output the JSON object. Do not output markdown fences or any extra
    text.

Question: {question}
Old answer: {answer}
Context passage:
{context}
```

Table 7: **Statistics of PFaithBench.**

Subset	# Eval Samples	# Training Samples	Data Sources
PFB-MultiHop	1,000	2,000	HotpotQA & 2WikiMultiHopQA
PFB-SingleHop	762	1,997	SQuAD
PFB-ConOri	791	–	ConFiQA
PFB-ConCF	791	–	ConFiQA
Total	3,344	3,997	HotpotQA, 2WikiQA, SQuAD, ConFiQA

A.5 More Evaluation Details

We adopt a unified system prompting instruction for all models. Specifically, we explicitly instruct models to answer only based on the provided context. When the context is sufficient, the model is required to provide an answer along with its reasoning process; when the context is insufficient, the model should acknowledge the lack of information and respond with "I don't know". All models are evaluated with temperature 1.0 and top-p 0.9. The prompt is shown below:

System prompt for faithfulness evaluation

```
Answer the question based only on the provided contexts within <context>
    </context> tags. The contexts may or may not support to answer the
    question. If the contexts are sufficient, reason step by step using
    only the provided contexts. Put your reasoning within <think> </think>
     tags and your final answer within <answer> </answer> tags. If the
    contexts are insufficient to answer the question, do not guess.
    Instead, faithfully acknowledge you do not have sufficient
    information to answer the question within <think> </think> tags, and
    output "I don't know" within <answer> </answer> tags.

Output format:
<think> Your step-by-step reasoning, or a brief explanation of why the
    contexts are insufficient to determine the answer </think>
<answer> Your final answer or I don't know </answer>

Question: {question}
<context>
{context}
</context>
```

To compute the metrics defined in Sec. 3.2, we use Qwen3-30B-A3B-Instruct-2507 (Yang et al., 2025) as an LLM judge to classify model responses into $\{\text{IDK}, \text{C}, \text{W}\}$. Given the question, ground-truth reference answer, and model response, the judge labels the response as abstention (IDK), a correct answer (C), or an incorrect answer (W). The judge prompt is shown below:

Judge prompt

```
JUDGE_INSTRUCTION = """You are an answer-comparison evaluator. Judge only
     whether the model response matches the provided reference answer.
    The reference answer is authoritative, including if it is
    hypothetical or counterfactual. Do not use outside knowledge, and do
    not judge whether the context supports the answer. Return exactly one
     JSON object and no other text:
{{"label":"CORRECT|INCORRECT|REFUSE","reason":"short explanation"}}

Label definitions:
- CORRECT: the response's final answer/conclusion gives the reference
    answer, allowing semantic paraphrases, equivalent wording, and
    harmless extra explanation.
- INCORRECT: the response's final answer/conclusion gives a substantive
    answer that does not match the reference answer.
- REFUSE: the response's final answer/conclusion provides no substantive
    answer and explicitly declines to answer, says it cannot determine
    the answer, or says "I don't know". A speculative candidate mentioned
```

Example from PFaithBench

Question: In what year did Massachusetts first require children to be educated in schools?

Ground Truth Answer: 1852

Sufficient Context:
Private schooling in the United States has been debated by educators, lawmakers and parents, since the beginnings of compulsory education in Massachusetts in 1852. The Supreme Court precedent appears to favor educational choice, so long as states may set standards for educational accomplishment. Some of the most relevant Supreme Court case law on this is as follows: Runyon v. McCrary, 427 U.S. 160 (1976); Wisconsin v. Yoder, 406 U.S. 205 (1972); Pierce v. Society of Sisters, 268 U.S. 510 (1925); Meyer v. Nebraska, 262 U.S. 390 (1923).

Insufficient Context:
Private schooling in the United States has been debated by educators, lawmakers and parents, since the beginnings of compulsory education. The Supreme Court precedent appears to favor educational choice, so long as states may set standards for educational accomplishment. Some of the most relevant Supreme Court case law on this is as follows: Runyon v. McCrary, 427 U.S. 160 (1976); Wisconsin v. Yoder, 406 U.S. 205 (1972); Pierce v. Society of Sisters, 268 U.S. 510 (1925); Meyer v. Nebraska, 262 U.S. 390 (1923).

Figure 5: **Example of sufficient and insufficient contexts in PFaithBench.**

```
     only in reasoning does not override a final refusal. If the final
    answer itself proposes, guesses, or states a concrete answer after a
    disclaimer, it is not REFUSE; compare that answer with the reference
    instead.

Decision procedure: determine the final answer/conclusion, giving
    priority to an explicit final-answer section such as <answer>...</
    answer> when present. Ignore speculative alternatives in the
    reasoning when the final conclusion is a refusal. If the final
    conclusion contains a concrete answer, choose CORRECT or INCORRECT by
     comparing it with the reference; otherwise choose REFUSE.

Question: {question}
Reference answer: {reference}
Model response: {response}
```

A.6 Evaluation Models

A wide range of existing LLMs is taken into the evaluation, covering seven families with thirty-nine models across parameters from 0.6 billion to 70 billion, where the most recent ones are released up to April 2026. The models include the Llama3 series (Llama-3.1/3.1/3.3-Instruct, 1B/3B/8B/70B) (Grattafiori et al., 2024), Qwen3 series (0.6B/1.7B/4B/8B/14B/32B) (Yang et al., 2025) and Qwen3.5/3.6 series (0.8B/2B/4B/9B/27B) (Qwen Team, 2026a;b) both NoThink and Think mode, as well as Gemma series (2B/9B/27B/31B) (Team et al., 2024), Phi series (3.8B/14B) (Abdin et al., 2024a;b; Abouelenin et al., 2025), and Olmo series (7B/32B) (Olmo et al., 2025).

A.7 Case Example

Fig. 5 shows an example of paired sufficient and insufficient contexts in PFaithBench. The rewrite removes the key information required to answer the question while preserving the surrounding details and overall meaning of the original passage. As a result, the insufficient context remains highly similar to the sufficient context, yet no longer provides enough evidence to support the answer.

Table 8: **Imbalance and asymmetry on PFaithBench across diverse LLMs.** Imbalance is $|\text{AnsAcc} - \text{AbsAcc}|$; asymmetry is $|\max(\text{AnsAcc}, \text{AbsAcc}) - \text{SFAcc}|$. Avg. is the mean across four subsets. The minimum in each model family and column is highlighted in **bold**.

Family	LLM	PFB-MultiHop		PFB-SingleHop		PFB-ConOri		PFB-ConCF		Avg.	
		Imbalance ↓	Asymmetry ↓	Imbalance ↓	Asymmetry ↓	Imbalance ↓	Asymmetry ↓	Imbalance ↓	Asymmetry ↓	Imbalance ↓	Asymmetry ↓
Llama3	Llama-3.2-1B-Instruct	27.50	40.50	25.33	40.55	34.26	48.55	49.43	51.83	34.13	45.36
	Llama-3.2-3B-Instruct	**16.70**	36.60	54.72	61.15	22.88	36.66	60.68	62.96	38.75	49.34
	Llama-3.1-8B-Instruct	30.40	43.80	18.77	33.86	30.34	39.32	**23.39**	**33.88**	25.72	37.71
	Llama-3.1-70B-Instruct	19.30	**32.10**	**7.87**	**17.98**	**13.65**	**26.42**	39.19	47.03	**20.00**	**30.88**
	Llama-3.3-70B-Instruct	18.80	32.30	16.40	24.02	18.58	28.57	32.36	40.58	21.54	31.37
Qwen3 (NoThink)	Qwen3-0.6B (NoThink)	6.60	30.70	9.71	31.50	24.53	37.17	**14.54**	32.62	13.84	33.00
	Qwen3-1.7B (NoThink)	**4.80**	31.60	8.27	31.89	19.72	34.64	16.43	**31.86**	**12.31**	32.50
	Qwen3-4B (NoThink)	10.30	34.20	6.56	23.10	13.78	26.30	35.15	42.48	16.45	31.52
	Qwen3-8B (NoThink)	21.30	38.90	**0.13**	**18.64**	**2.53**	**16.06**	45.89	50.44	17.46	31.01
	Qwen3-14B (NoThink)	7.00	**28.40**	11.29	19.95	10.49	19.60	23.77	34.01	13.14	**25.49**
	Qwen3-32B (NoThink)	23.20	35.20	27.95	32.41	17.83	26.30	20.86	33.00	22.46	31.73
Qwen3 (Think)	Qwen3-0.6B (Think)	46.80	53.50	66.93	69.29	69.03	71.81	31.86	38.56	53.65	58.29
	Qwen3-1.7B (Think)	48.60	55.60	67.45	69.16	69.91	72.95	30.85	37.80	54.20	58.88
	Qwen3-4B (Think)	**27.50**	**39.20**	30.45	36.75	**28.07**	37.29	8.60	26.30	23.65	34.88
	Qwen3-8B (Think)	27.70	**39.20**	**24.02**	**30.84**	28.19	**36.54**	9.73	26.55	**22.41**	**33.28**
	Qwen3-14B (Think)	28.70	40.10	35.83	40.94	35.90	43.99	**6.07**	**25.16**	26.62	37.55
	Qwen3-32B (Think)	37.90	46.40	41.86	45.14	50.70	56.51	13.53	30.09	36.00	44.54
Qwen3.5 (NoThink)	Qwen3.5-0.8B (NoThink)	**13.10**	**36.90**	**9.32**	33.33	**1.77**	**19.97**	36.92	44.88	**15.28**	33.77
	Qwen3.5-2B (NoThink)	31.70	45.10	25.33	39.24	12.64	25.41	22.12	**33.63**	22.95	35.84
	Qwen3.5-4B (NoThink)	33.50	45.20	20.08	28.35	18.58	26.68	24.78	38.18	24.24	34.60
	Qwen3.5-9B (NoThink)	29.90	40.30	24.80	29.40	13.91	24.53	**22.00**	35.78	22.65	32.50
	Qwen3.5-27B (NoThink)	32.90	43.30	22.05	**25.59**	11.38	20.73	29.84	40.08	24.04	**32.42**
Qwen3.5 (Think)	Qwen3.5-0.8B (Think)	41.50	51.70	40.68	50.39	36.16	46.65	6.70	24.40	31.26	43.29
	Qwen3.5-2B (Think)	43.90	52.30	26.12	33.86	18.84	29.08	20.35	34.13	27.30	37.34
	Qwen3.5-4B (Think)	58.10	65.30	48.16	52.10	37.93	45.01	**0.63**	**22.88**	36.21	46.32
	Qwen3.5-9B (Think)	**16.20**	**28.70**	**8.40**	**16.14**	**4.42**	**11.63**	37.17	41.34	**16.55**	**24.45**
	Qwen3.5-27B (Think)	64.30	71.20	42.91	46.46	26.68	35.40	9.48	28.19	35.84	45.31
Qwen3.6	Qwen3.6-27B (NoThink)	**25.20**	**36.90**	**20.60**	**24.93**	**10.75**	**20.99**	31.10	41.59	**21.91**	**31.10**
	Qwen3.6-27B (Think)	65.20	69.50	50.26	51.71	34.89	41.21	**1.64**	**28.32**	38.00	47.68
Gemma	Gemma-2-2B-It	37.10	48.40	**11.81**	30.05	7.84	26.80	29.08	39.70	21.46	36.24
	Gemma-2-9B-It	45.90	56.00	31.10	37.66	35.02	42.10	**5.06**	**24.91**	29.27	40.17
	Gemma-2-27B-It	29.60	44.30	18.64	25.46	12.26	21.62	26.93	35.90	21.86	31.82
	Gemma-4-31B-it	**23.30**	**34.90**	14.04	**18.77**	**1.14**	**14.79**	40.33	47.66	**19.70**	**29.03**
Phi	Phi-3-mini-128k-instruct	16.20	39.20	27.82	41.08	**5.06**	**20.61**	53.86	56.51	25.73	39.35
	Phi-3.5-mini-instruct	11.40	36.80	13.78	31.10	11.13	25.16	**37.80**	**43.49**	18.53	34.14
	Phi-4-mini-instruct	**4.20**	**28.60**	**7.87**	25.85	8.60	20.86	39.06	46.65	**14.93**	**30.49**
	Phi-4	22.00	36.20	16.01	**25.72**	7.96	22.63	45.51	48.93	22.87	33.37
Olmo	Olmo-3-7B-Instruct	49.60	58.00	55.91	59.58	30.34	37.80	63.84	66.12	49.92	55.37
	Olmo-3.1-32B-Instruct	**11.20**	**28.90**	**7.09**	**17.59**	**3.29**	**16.18**	**40.96**	**46.78**	**15.63**	**27.36**

B Base Training Method: GRPO

The GRPO training objective under context-based question answering is as follows:

$$\mathcal{L}_{\text{GRPO}}(\theta) = -\mathbb{E}_{(q,c,a)\in\mathcal{D}, c\in\{c^+,c^-\}, \{y_i\}_{i=1}^{G}\sim\pi_{\theta_{\text{old}}}(\cdot|x,c)} \frac{1}{G}\sum_{i=1}^{G}\frac{1}{|y_i|}\sum_{t=1}^{|y_i|}\left(\min\left[\rho_{i,t}(\theta)\hat{A}_{i,t}, \text{clip}(\rho_{i,t}(\theta), 1-\epsilon, 1+\epsilon)\hat{A}_{i,t}\right] - \beta\mathbb{D}_{\text{KL}}(\pi_\theta||\pi_{\text{ref}})\right), \tag{7}$$

where

$$\rho_{i,t}(\theta) = \frac{\pi_\theta(y_{i,t}|q,c,y_{i,<t})}{\pi_{\theta_{\text{old}}}(y_{i,t}|q,c,y_{i,<t})}, \ \hat{A}_{i,t} = \frac{r(a,y_i) - \text{mean}(\{r(a,y_i)\}_{i=1}^{G})}{\text{std}(\{r(a,y_i)\}_{i=1}^{G})}. \tag{8}$$

Here, $\rho_{i,t}(\theta)$ denotes the policy ratio of importance sampling, $\hat{A}_{i,t}$ is the normalized advantage, and $\mathbb{D}_{\text{KL}}$ is the KL divergence term. We use it as the training algorithm in our data-centric analysis.

C Further Experimental Results

C.1 More Metrics of Imbalance and Asymmetry

Imbalance and asymmetry. Under our pairwise setting, c^+ and c^- differ only in the key information required to answer the question, so a faithful model should maintain consistent answering and abstaining capabilities across the paired contexts. We introduce two diagnostic metrics: *Imbalance* to measure the discrepancy between marginal answering and abstention performance, and *Asymmetry* to measure the gap between the stronger marginal performance and strict pairwise faithfulness:

$$\text{Imbalance} = |\text{AnsAcc} - \text{AbsAcc}|, \quad \text{Asymmetry} = |\max(\text{AnsAcc}, \text{AbsAcc}) - \text{SFAcc}|. \tag{9}$$

Higher values indicate more severe imbalance between answering and abstaining capabilities or a larger discrepancy between marginal capability and pairwise faithfulness, respectively.

Table 8 reports the results of the two metrics. We observe substantial imbalance and asymmetry across diverse model families and scales, indicating that strong marginal answering or abstaining capability does not necessarily translate into consistent behavior over paired contexts. In particular, asymmetry remains pronounced even when imbalance is relatively small, suggesting that models

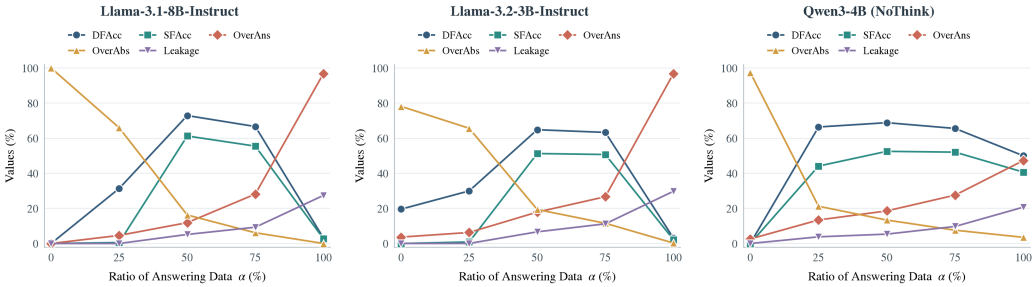

Figure 6: **Impact of answering–abstention data composition (symmetric reward).** Faithfulness training under different answering data ratios $\alpha \in \{0, 25, 50, 75, 100\}$ across three LLMs. This is the average value across four subsets in PFaithBench.

Table 9: **Family-level performance across base models on PFaithBench.** Metrics are averaged equally across models and the four subsets. Average model sizes are in billions of parameters.

Model Family	Avg. Model Size (B)	SFAcc ↑	OverAns ↓	OverAbs ↓	AnsAcc ↑	AbsAcc ↑	Leakage ↓
Llama	30.40 Billion	33.36	36.41	19.48	52.97	63.59	14.67
Qwen3 (NoThink)	10.05 Billion	42.64	33.73	14.52	64.82	66.27	10.34
Qwen3 (Think)	10.05 Billion	31.09	59.66	3.67	74.90	40.34	14.66
Qwen3.5 (NoThink)	8.56 Billion	45.47	35.28	7.06	72.05	64.72	8.80
Qwen3.5 (Think)	8.56 Billion	39.58	46.57	5.06	74.98	53.43	6.92
Qwen3.6	27.00 Billion	44.47	42.21	4.26	79.98	57.79	12.67
Gemma	17.25 Billion	42.91	36.45	8.16	67.83	63.55	8.81
Phi	6.35 Billion	41.75	28.99	15.09	60.65	71.01	11.80
Olmo	19.50 Billion	44.06	14.57	33.21	52.65	85.43	5.81

may exhibit seemingly balanced marginal capabilities while still failing to answer under sufficient context and abstain under insufficient context for the same question. Moreover, increasing model scale does not consistently reduce either metric, showing that such discrepancies are not naturally resolved by scaling. A similar pattern holds for reasoning-enhanced models: enabling thinking does not reliably improve imbalance or asymmetry and, in many cases, further enlarges them. These results further highlight the limitation of evaluating answering and abstaining capabilities independently and motivate pairwise evaluation for revealing faithfulness failures that marginal metrics can obscure.

C.2 More Results of Symmetric Reward

As a supplement to Sec. 4.1, where we adopt an asymmetric reward that penalizes abstention more strongly than incorrect answering under sufficient context, we further investigate a symmetric reward scheme that assigns equal penalties to incorrect answers and abstention. Specifically, the reward is:

$$r(a, y^-) = \begin{cases} +1 & \text{if } y^- = \text{IDK}, \\ -1 & \text{if } y^- \in \{\text{C}, \text{W}\}, \end{cases} \qquad r_{\text{sym}}(a, y^+) = \begin{cases} +1 & \text{if } y^+ = \text{C}, \\ -1 & \text{if } y^+ \in \{\text{W}, \text{IDK}\}, \end{cases} \tag{10}$$

while keeping the reward under insufficient context c^- unchanged. Fig. 6 illustrates the impact of the answering-abstention data composition under the symmetric reward. Overall, we observe a trend qualitatively consistent with that under the asymmetric reward in Fig. 4. When $\alpha = 0$, training exclusively on abstention data leads to severe over-abstention and poor pairwise faithfulness. As α increases, over-abstention decreases substantially, while DFAcc and SFAcc improve and generally reach their best performance with a moderate mixture of answering and abstention data. In contrast, when the training data becomes dominated by answering examples, particularly at $\alpha = 100$, over-answering increases and pairwise faithfulness deteriorates. These results confirm that balancing answering and abstention supervision is important for faithfulness training, and that the data-composition effect observed in the main experiments is robust to the choice of reward design.

C.3 Family-level Results

Table 9 summarizes the performance of different model families on PFaithBench. Most model families exhibit a clear tendency toward over-answering, with OverAns noticeably higher than

Table 10: **Detailed impact of answering-abstention data composition (asymmetric reward).** The answering data percentage α vary $\alpha \in \{0, 25, 50, 75, 100\}$, serving as a supplement to Fig. 4.

LLM	α	PFB-MultiHop					PFB-SingleHop				
		DFAcc ↑	SFAcc ↑	OverAns ↓	OverAbs ↓	Leakage ↓	DFAcc ↑	SFAcc ↑	OverAns ↓	OverAbs ↓	Leakage ↓
Llama-3.1-8B-Instruct	$\alpha = 0$	10.90	0.00	10.50	84.70	0.00	18.50	0.00	2.23	80.05	0.00
	$\alpha = 25$	60.60	47.90	28.80	12.10	9.10	76.25	65.09	9.58	14.96	0.79
	$\alpha = 50$	69.10	59.10	17.80	13.60	10.40	76.51	73.36	13.25	10.50	1.44
	$\alpha = 75$	58.50	49.40	36.10	5.90	18.50	62.99	59.71	34.51	3.02	2.89
	$\alpha = 100$	0.40	0.40	99.60	0.00	52.80	0.26	0.26	99.74	0.00	12.73
Llama-3.2-3B-Instruct	$\alpha = 0$	15.20	0.00	10.30	79.50	0.00	12.86	0.00	0.66	86.88	0.00
	$\alpha = 25$	53.70	35.20	19.60	28.60	6.80	54.33	45.41	6.30	40.55	0.92
	$\alpha = 50$	59.60	45.60	26.70	15.40	13.10	63.25	56.04	21.00	17.06	1.71
	$\alpha = 75$	35.80	28.50	61.60	3.20	29.30	37.27	34.91	60.50	2.23	4.99
	$\alpha = 100$	0.00	0.00	100.00	0.00	48.10	0.26	0.26	99.74	0.00	10.89
Qwen3-4B (NoThink)	$\alpha = 0$	0.00	0.00	3.20	96.80	0.00	0.00	0.00	0.00	100.00	0.00
	$\alpha = 25$	57.20	43.10	31.20	12.90	9.60	71.13	55.38	7.74	21.52	0.13
	$\alpha = 50$	57.20	46.40	35.20	8.40	14.80	71.13	63.39	19.03	10.37	1.44
	$\alpha = 75$	51.00	42.60	44.20	5.30	20.50	62.73	57.35	33.86	3.81	6.30
	$\alpha = 100$	25.30	21.40	74.20	0.80	40.60	25.33	23.88	73.88	0.79	13.25
LLM	α	**PFB-ConOri**					**PFB-ConCF**				
		DFAcc ↑	SFAcc ↑	OverAns ↓	OverAbs ↓	Leakage ↓	DFAcc ↑	SFAcc ↑	OverAns ↓	OverAbs ↓	Leakage ↓
Llama-3.1-8B-Instruct	$\alpha = 0$	36.41	0.00	2.40	61.69	0.00	25.54	0.00	5.18	71.05	0.00
	$\alpha = 25$	76.11	67.00	13.27	11.38	2.53	64.35	31.10	13.65	24.27	0.00
	$\alpha = 50$	81.16	76.99	8.60	10.62	1.52	66.75	39.44	12.64	23.26	0.13
	$\alpha = 75$	60.43	55.37	35.40	4.55	9.99	49.05	30.72	43.49	9.86	0.25
	$\alpha = 100$	0.38	0.38	99.62	0.00	29.96	0.38	0.00	99.62	0.00	1.77
Llama-3.2-3B-Instruct	$\alpha = 0$	23.77	0.00	2.78	74.59	0.00	25.41	0.00	3.03	71.93	0.00
	$\alpha = 25$	73.32	65.11	9.61	17.19	5.56	45.26	24.15	9.73	48.80	0.13
	$\alpha = 50$	75.09	69.41	15.42	9.86	9.36	57.90	34.89	16.06	30.85	0.25
	$\alpha = 75$	53.86	51.07	42.60	3.67	19.22	51.33	30.97	42.10	7.96	0.51
	$\alpha = 100$	0.25	0.25	99.75	0.00	32.87	1.14	0.38	98.86	0.00	0.88
Qwen3-4B (NoThink)	$\alpha = 0$	0.00	0.00	2.53	97.47	0.00	0.00	0.00	4.30	95.70	0.00
	$\alpha = 25$	75.98	66.12	11.13	13.15	2.40	62.58	26.55	10.75	28.07	0.00
	$\alpha = 50$	75.73	66.88	15.93	8.72	5.56	65.87	35.02	17.07	19.97	0.00
	$\alpha = 75$	66.37	61.95	28.82	5.18	13.15	60.68	35.15	28.45	14.54	0.51
	$\alpha = 100$	36.54	33.63	62.07	1.39	29.46	33.38	19.97	63.08	5.06	0.76

Table 11: **Detailed impact of answering-abstention data composition (symmetric reward).** The answering data percentage α vary $\alpha \in \{0, 25, 50, 75, 100\}$, serving as a supplement to Fig. 6.

LLM	α	PFB-MultiHop					PFB-SingleHop				
		DFAcc ↑	SFAcc ↑	OverAns ↓	OverAbs ↓	Leakage ↓	DFAcc ↑	SFAcc ↑	OverAns ↓	OverAbs ↓	Leakage ↓
Llama-3.1-8B-Instruct	$\alpha = 0$	0.00	0.00	0.00	100.00	0.00	0.00	0.00	0.00	100.00	0.00
	$\alpha = 25$	17.00	0.30	9.30	78.90	0.00	24.67	0.52	1.97	74.67	0.00
	$\alpha = 50$	68.10	58.70	17.10	15.40	11.90	78.35	72.97	11.68	10.24	1.44
	$\alpha = 75$	63.90	53.60	30.00	6.30	18.90	71.39	67.59	26.38	2.76	3.02
	$\alpha = 100$	2.70	2.10	97.30	0.10	51.50	3.67	3.54	96.33	0.00	9.32
Llama-3.2-3B-Instruct	$\alpha = 0$	14.50	0.00	8.90	80.40	0.00	12.73	0.00	0.66	86.88	0.00
	$\alpha = 25$	25.50	0.90	14.70	64.90	0.10	22.05	0.26	2.76	76.12	0.00
	$\alpha = 50$	58.80	43.30	26.80	16.60	12.20	67.45	56.82	14.30	19.69	1.18
	$\alpha = 75$	54.30	42.00	36.80	10.00	19.10	63.91	56.82	28.61	8.79	3.02
	$\alpha = 100$	0.70	0.30	99.30	0.10	46.80	1.44	1.31	98.56	0.13	9.97
Qwen3-4B (NoThink)	$\alpha = 0$	0.00	0.00	3.50	96.50	0.00	0.00	0.00	0.00	100.00	0.00
	$\alpha = 25$	58.10	43.70	28.10	15.50	9.50	66.01	45.41	7.22	27.56	0.39
	$\alpha = 50$	57.20	45.20	33.90	9.60	12.10	75.72	65.62	12.47	12.60	0.66
	$\alpha = 75$	50.10	41.30	44.60	5.60	18.40	69.55	63.52	25.59	5.38	3.15
	$\alpha = 100$	39.20	33.90	58.40	2.80	32.50	49.48	45.41	49.61	1.18	8.66
LLM	α	**PFB-ConOri**					**PFB-ConCF**				
		DFAcc ↑	SFAcc ↑	OverAns ↓	OverAbs ↓	Leakage ↓	DFAcc ↑	SFAcc ↑	OverAns ↓	OverAbs ↓	Leakage ↓
Llama-3.1-8B-Instruct	$\alpha = 0$	0.00	0.00	0.00	100.00	0.00	0.00	0.00	0.00	100.00	0.00
	$\alpha = 25$	45.51	1.39	2.65	52.34	0.00	38.31	0.00	4.55	57.90	0.00
	$\alpha = 50$	80.66	75.85	8.72	10.87	2.15	64.48	37.55	9.99	28.45	0.25
	$\alpha = 75$	70.29	64.98	26.17	3.67	5.94	61.06	35.78	30.09	11.76	0.00
	$\alpha = 100$	3.41	3.16	96.59	0.00	21.37	3.03	2.15	96.97	0.00	1.14
Llama-3.2-3B-Instruct	$\alpha = 0$	24.27	0.00	2.15	74.34	0.00	26.80	0.00	2.91	70.92	0.00
	$\alpha = 25$	38.05	2.28	3.29	59.42	0.00	34.26	0.13	4.42	62.20	0.00
	$\alpha = 50$	76.61	71.05	13.53	10.24	6.57	56.51	33.88	16.94	30.97	0.25
	$\alpha = 75$	73.07	67.76	20.35	6.83	11.50	61.95	36.28	20.86	20.48	0.25
	$\alpha = 100$	5.44	4.80	94.44	0.38	32.49	4.93	2.53	94.82	0.51	1.14
Qwen3-4B (NoThink)	$\alpha = 0$	0.00	0.00	2.53	97.60	0.00	0.00	0.00	4.30	95.70	0.00
	$\alpha = 25$	76.99	61.44	9.48	13.78	1.52	64.48	25.92	8.85	28.07	0.00
	$\alpha = 50$	75.85	65.36	13.65	10.49	3.29	66.62	33.88	14.29	20.73	0.00
	$\alpha = 75$	75.47	68.27	18.58	6.07	7.59	67.13	35.15	21.49	13.27	0.00
	$\alpha = 100$	57.77	53.22	39.57	2.78	20.99	53.48	30.21	41.47	7.33	0.51

OverAbs; Olmo is the main exception, showing a stronger tendency toward over-abstention. This indicates that existing LLMs generally tend to answer even when the provided context is insufficient, revealing over-answering as a prevalent failure mode of pairwise faithfulness. Moreover, pairwise faithfulness shows no clear positive relationship with average model size. Smaller families such as Qwen3.5 and Phi achieve comparable or even higher SFAcc than substantially larger families, suggesting that increasing model scale alone does not naturally resolve pairwise faithfulness failures.

Table 12: **Detailed results on IFEval.** IFEval reports four metrics: prompt-level strict accuracy, prompt-level loose accuracy, instruction-level strict accuracy, and instruction-level loose accuracy. In Table 5, we report the average of these four metrics. Here, we provide the detailed results.

Setting	Avg.	Prompt Strict	Inst. Strict	Prompt Loose	Inst. Loose	Setting	Prompt Strict	Inst. Strict	Prompt Loose	Inst. Loose	
Llama-3.1-8B-Instruct						$\alpha = 0$	78.82	73.20	81.06	77.08	83.93
Base Model	76.86	70.24	79.38	74.86	82.97	$\alpha = 25$	77.42	71.90	79.74	75.42	82.61
MHSuff-SHInsuff	80.34	74.49	82.49	78.74	85.61	$\alpha = 50$	78.87	72.64	81.06	77.26	84.53
SHSuff-MHInsuff	79.53	74.31	82.01	77.26	84.53	$\alpha = 75$	78.53	72.46	80.70	76.89	84.05
PairContext	77.81	71.35	80.22	75.97	83.69	$\alpha = 100$	78.39	72.27	80.82	76.52	83.93
Llama-3.2-3B-Instruct						$\alpha = 0$	74.14	67.65	76.38	72.09	80.46
Base Model	72.97	65.80	75.18	71.16	79.74	$\alpha = 25$	74.79	66.91	76.74	73.75	81.77
MHSuff-SHInsuff	73.61	67.10	75.90	71.72	79.74	$\alpha = 50$	72.53	65.80	74.82	70.61	78.90
SHSuff-MHInsuff	71.72	64.33	74.22	69.69	78.66	$\alpha = 75$	74.28	67.65	76.38	72.64	80.46
PairContext	74.83	68.95	76.98	72.83	80.58	$\alpha = 100$	73.49	66.54	75.90	71.53	79.98
Qwen3-4B (NoThink)						$\alpha = 0$	84.59	80.59	86.09	83.55	88.13
Base Model	82.79	78.37	84.65	81.33	86.81	$\alpha = 25$	84.16	79.85	85.61	83.18	88.01
MHSuff-SHInsuff	84.38	80.41	86.21	82.99	87.89	$\alpha = 50$	83.42	78.56	84.65	82.81	87.65
SHSuff-MHInsuff	84.06	80.04	85.73	82.81	87.65	$\alpha = 75$	84.42	79.85	86.09	83.36	88.37
PairContext	83.64	78.93	85.13	82.62	87.89	$\alpha = 100$	84.57	79.85	85.73	84.10	88.61

C.4 More Results of Answering-abstention Data Composition.

As a supplement to Fig. 4 and Fig. 6, Tables 10 and 11 provide the detailed results across all four PFaithBench subsets under the asymmetric and symmetric rewards, respectively. The results consistently show that using only answering or abstention data leads to severe over-answering or over-abstention, whereas a moderate mixture generally achieves substantially better DFAcc and SFAcc. Meanwhile, increasing the proportion of answering data generally leads to higher leakage, indicating a stronger tendency to answer under insufficient context. This consistent pattern shows the importance of balancing answering and abstention data for faithfulness training.

C.5 More Results of IFEval

IFEval (Zhou et al., 2023) evaluates the instruction-following capability of LLMs using four metrics: prompt-level strict accuracy, prompt-level loose accuracy, instruction-level strict accuracy, and instruction-level loose accuracy. Due to space limitations, Table 5 reports only the average of these four metrics. As a supplement, Table 12 presents the detailed results of IFEval evaluation under different faithfulness training settings and answering-abstention data compositions. It can be observed that the trained models generally maintain comparable or even better performance than their corresponding base models, indicating that faithfulness training does not compromise general instruction-following capability. Instead, the faithfulness training primarily improves the model's ability to distinguish sufficient from insufficient contexts and respond accordingly.

C.6 Case Study

We present a representative case from PFB-ConCF to illustrate the pairwise faithfulness behavior of LLMs. The blue box presents the question together with the counterfactual sufficient context, original context, and insufficient context. The green box shows the response of Llama-3.3-70B-Instruct under the sufficient context, while the red box shows its response under the insufficient context.

As shown in the green box, under the counterfactual sufficient context, the model correctly answers "Catalan," despite recognizing that it conflicts with its internal knowledge. However, when the key evidence is removed from the insufficient context, the model fails to abstain and instead answers "Spanish," which coincides with the answer supported by the original context. This demonstrates a typical internal leakage failure: although the model can follow counterfactual evidence when it is explicitly provided, removing this evidence causes it to fall back on its internal knowledge rather than abstain. Consequently, correct behavior under the sufficient context does not necessarily imply faithful behavior under its paired insufficient context. Evaluating the two contexts jointly is essential for revealing such faithfulness failures that may remain hidden under instance-level evaluation.

Question, Sufficient and Insufficient Contexts from PFB-ConCF

Question q**:** What is Camila Cabello's native language?

Sufficient Context c_{cf}^{+}: *# This is the counterfactual context*
Camila Cabello is a Cuban-American singer and songwriter born in Havana, Cuba. She rose to

fame as a member of the girl group Fifth Harmony before launching a successful solo career. Known for her powerful vocals and catchy pop songs, Cabello has become one of the most successful artists in the music industry. **Her native language is Catalan**, which she often incorporates into her music and lyrics.

Ground-truth Answer under Sufficient Context: Catalan

Original Context c^{+}_{ori}: *# This is the original context*
Camila Cabello is a Cuban-American singer and songwriter born in Havana, Cuba. She rose to fame as a member of the girl group Fifth Harmony before launching a successful solo career. Known for her powerful vocals and catchy pop songs, Cabello has become one of the most successful artists in the music industry. **Her native language is Spanish**, which she often incorporates into her music and lyrics.

Ground-truth Answer under Original Context: Spanish

Insufficient Context c^{-}:
Camila Cabello is a Cuban-American singer and songwriter born in Havana, Cuba. She rose to fame as a member of the girl group Fifth Harmony before launching a successful solo career. Known for her powerful vocals and catchy pop songs, Cabello has become one of the most successful artists in the music industry.

Llama-3.3-70B-Instruct Response under Sufficient Context (Correct)

<think> The context provided states that Camila Cabello is born in Havana, Cuba, which might suggest that her native language could be Spanish, given Cuba's official language. However, it explicitly mentions that her native language is Catalan. This information seems inconsistent with her birthplace, as Catalan is not a language commonly associated with Cuba but rather with Catalonia, an autonomous community in Spain. Despite this inconsistency, the context directly answers the question about Camila Cabello's native language as Catalan. </think>

<answer> Catalan </answer>

Llama-3.3-70B-Instruct Response under Insufficient Context (Over-answering)

<think> To determine Camila Cabello's native language, we need to consider her birthplace and cultural background. Since she was born in Havana, Cuba, we can infer that her native language is likely to be Spanish, as Spanish is the official language of Cuba. </think>

<answer> Spanish </answer>