

# Improving Contextual Faithfulness of Large Language Models via Retrieval Heads-Induced Optimization

Lei Huang<sup>1</sup>, Xiaocheng Feng<sup>1,2\*</sup>, Weitao Ma<sup>1</sup>, Yuchun Fan<sup>3</sup>, Xiachong Feng<sup>4</sup>, Yangfan Ye<sup>1</sup>, Weihong Zhong<sup>1</sup>, Yuxuan Gu<sup>1</sup>, Baoxin Wang<sup>5</sup>, Dayong Wu<sup>5</sup>, Guoping Hu<sup>5</sup>, Bing Qin<sup>1,2</sup>

<sup>1</sup> Harbin Institute of Technology, China <sup>2</sup> Peng Cheng Laboratory, China

<sup>3</sup> Northeastern University, China <sup>4</sup> The University of Hong Kong, China

<sup>5</sup> iFLYTEK Research, China

{lhuang, xcfeng, wtma, yfye whzhong, yxgu, qinb}@ir.hit.edu.cn

yuchunfan\_neu@outlook.com fengxc@hku.hk {bxwang2, dywu2, gphu}@iflytek.com

## Abstract

Ensuring contextual faithfulness in retrieval-augmented large language models (LLMs) is crucial for building trustworthy information-seeking systems, particularly in long-form question-answering (LFQA) scenarios. In this work, we identify a salient correlation between LFQA faithfulness and *retrieval heads*, a set of attention heads responsible for retrieving contextual information. Leveraging this insight, we propose RHIO<sup>1</sup>, a framework designed to teach LLMs to explicitly discriminate between faithful and unfaithful generations. RHIO first augments unfaithful samples that simulate realistic model-intrinsic errors by selectively masking retrieval heads. Then, these samples are incorporated into joint training, enabling the model to distinguish unfaithful outputs from faithful ones conditioned on *control tokens*. Furthermore, these control tokens are leveraged to self-induce contrastive outputs, amplifying their difference through contrastive decoding. Additionally, to facilitate the evaluation of contextual faithfulness, we also introduce GroundBench, a comprehensive benchmark compiled from five existing LFQA datasets. Extensive experimental results on GroundBench demonstrate that RHIO significantly improves faithfulness, even outperforming GPT-4o<sup>2</sup>.

## 1 Introduction

Large language models (LLMs) have achieved remarkable advancements across various domains due to their powerful generative capabilities (Touvron et al., 2023; OpenAI, 2023). By integrating retrieval techniques, retrieval-augmented LLMs further enhance their performance in information-seeking tasks, producing highly coherent and informative responses (Gao et al., 2023b; Asai et al., 2024b). Particularly, long-form question answering

Model Input
<b>Question:</b> Difference between Russian Blue and British Blue cat? <b>Document [1]:</b> The Russian Blue is a naturally occurring breed that may have originated in the port of Arkhangelsk in Russia. They come in colors varying from a light shimmering silver to a darker, slate grey. (...) They are also considered to be hypoallergenic. <b>Document [2]:</b> The coat is known as a “double coat”, with the undercoat being soft, downy, and equal in length to the guard hairs, which are an even blue with silver tips. (...) The coat is described as thick, plush, and soft to the touch. The feeling is softer than the softest silk. <b>Document [3]:</b> The Russian Blue has bright green eyes, pinkish lavender or mauve paws, two layers of short thick fur, and a blue-grey coat. The color is a bluish-gray which is the dilute expression of the black gene. (...) Its eyes are almost always a dark and vivid green. Any white patches of fur or yellow eyes in adulthood are seen as flaws in show cats.
Model Output
<b>Unfaithful Response:</b> The Russian Blue and the British Blue cat differ in several distinct characteristics. One key difference is the length of their coats. While the Russian Blue has a “double coat” with a soft, downy undercoat and a plush, soft outer coat, <b>the British Blue typically has a single coat</b> . In terms of eye color, Russian Blues typically have dark green eyes, while British Blues may have yellow eyes or white patches of fur in adulthood, which are considered flaws in show cats.

Table 1: An example of the unfaithful response generated by GPT-4o-mini. We highlight the different error types in the response. The content marked in red is a *fabricated hallucination*, and the content marked in orange indicates *inaccurate information synthesis*.

(LFQA) (Fan et al., 2019; Stelmakh et al., 2022), aiming at synthesizing information from disparate retrieval sources to provide in-depth, paragraph-length responses, has emerged as a research hotspot due to its wide-ranging application. However, recent studies (Chen et al., 2023; Stolfo, 2024) reveal that these LLMs suffer from *faithfulness hallucination* in LFQA, where the generated response is not grounded in provided sources (Huang et al., 2023), as shown in Figure 1. Such unsupported information can significantly erode user trust and undermine the reliability of LLMs.

Recent efforts to enhance faithfulness in LFQA have primarily focused on improving context quality via explicit denoising (Xu et al., 2023a), self-critiquing its generation with reflection tokens (Asai et al., 2024a) and context-aware decoding to amplify contextual information (Shi et al., 2024). While effective, these approaches are more of a compensatory way rather than allowing the model

\*Corresponding Author

<sup>1</sup>Acronym for Retrieval Heads-Induced Optimization

<sup>2</sup><https://github.com/LuckyyySTA/faithful-LFQA>

to learn to avoid generating unfaithful responses.

In this work, we investigate the idea of explicitly teaching LLMs to discriminate between faithful and unfaithful outputs to improve contextual faithfulness. One of the key challenges lies in generating realistic unfaithful samples. Inspired by recent research on *retrieval heads* (Wu et al., 2024), a special type of attention head for retrieving information from context, we delve into the role of these retrieval heads in the contextual faithfulness of LFQA. Our pilot study reveals that the activation of retrieval heads potentially explains faithfulness in LFQA: masking out retrieval heads leads to error patterns similar to real unfaithfulness from the model (§2.2). This enables us to augment more diverse and realistic unfaithful samples by simply masking out retrieval heads in LLMs (§3.1). Given faithful and augmented unfaithful samples, RHIO introduces two special *control tokens*, [POS] and [NEG], to signal the generation towards either faithful or unfaithful responses. This process supervises the model to explicitly distinguish unfaithful responses from faithful ones (§3.2). Furthermore, these control codes are further utilized to induce contrastive generations and amplify the difference between them via contrastive decoding to further enhance faithfulness (§3.3).

Additionally, to reliably evaluate the faithfulness of LLMs in LFQA, we also introduce GroundBench, a comprehensive benchmark compiled from five existing LFQA datasets. GroundBench is designed to ensure that retrieved documents contain sufficient information to answer questions, thus providing a controlled evaluation setting. Extensive experiments on GroundBench demonstrate that RHIO achieves significant improvements in faithfulness, with average gains of 12.84% and 12.59% in 7B and 13B models, respectively, outperforming even the state-of-the-art GPT-4o. Human evaluation and further analysis reveal additional insights into the efficacy of RHIO.

## 2 Preliminaries

In this section, we introduce the task formulation of LFQA and our investigation of retrieval heads, which motivates the proposed of RHIO.

### 2.1 Task Formulation

The task of LFQA can be described as follows. Given a collection of questions  $\mathcal{Q}$  and a corpus of documents  $\mathcal{D}$ , for a question  $q \in \mathcal{Q}$ , a retriever




Figure 1: Impact of masking different numbers of masked retrievals heads on model faithfulness.

$R$  first retrieves  $k$  documents, denoted as  $R(q) = \{d_1, d_2, \dots, d_k\}$ , where  $d_i \in \mathcal{D}$ . Subsequently, the question  $q$  and the retrieved documents  $R(q)$  are combined as the input to the LLM  $M$  to generate a paragraph-length response  $\mathcal{S} = \langle s_1, s_2, \dots \rangle$  consisting of multiple sentences. In this work, we primarily focus on the faithfulness issue in LFQA. Notably, a model-generated response  $\mathcal{S}$  is considered unfaithful if it contains at least one sentence  $s_i \in \mathcal{S}$  that either contradicts or cannot be verified using the retrieved document  $R(q)$ .

### 2.2 Retrieval Head

Recent interpretability research (Olsson et al., 2022; Wu et al., 2024) has identified a special type of attention head, termed *retrieval heads*, which is largely responsible for retrieving relevant information from contextual sources. When these heads are activated, the model performs a copy-paste operation from the provided context. This motivates us to link these retrieval heads with contextual faithfulness in LFQA, which largely relies on the model’s ability to synthesize information from contextual cues. To delve into the implications of retrieval heads in contextual faithfulness, we start with a preliminary study on a long-form QA dataset.

#### 2.2.1 Experimental Setup for Pilot Study

We experiment with the CLAPNQ dataset (Rosen-thal et al., 2024), a component of our GroundBench, to explore how retrieval heads affect contextual faithfulness in LFQA (for details about CLAPNQ and evaluation of faithfulness, see Section §4). Specifically, we utilize models from Llama-2 family (Touvron et al., 2023), with sizes ranging from 7B to 70B. We first apply the retrieval head detection algorithm proposed by Wu et al. (2024) to detect retrieval heads within each model. Subsequently, We gradually mask out the top retrieval heads, from 0 to 100, to observe the corresponding changes in the faithfulness of the model’s response.

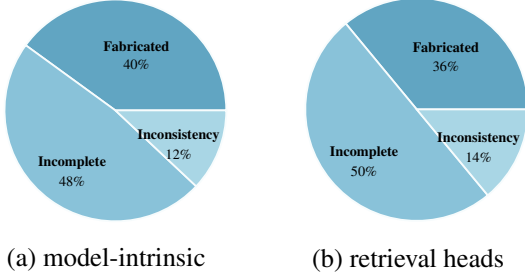


Figure 2: Comparison of two sets of error patterns.

### 2.2.2 Observation

Next, we present our main findings as follows.

#### Findings 1: Unfaithfulness in LFQA is exacerbated when masking out more retrieval heads.

As shown in Figure 1, our results indicate a direct correlation between the number of masked retrieval heads and the severity of unfaithfulness. As the number of masked retrieval heads increases, faithfulness in long-form responses leads to a substantial decrease. In contrast, randomly masking non-retrieval heads shows no significant impact on faithfulness (results available in Appendix A.1). This observation suggests that retrieval heads are crucial for ensuring contextual faithfulness in LFQA.

#### Findings 2: Error patterns of unfaithfulness induced by masking retrieval heads mirror real error patterns.

We further manually analyzed the error types triggered by masking retrieval heads and compared them to unfaithfulness error types generated from the model itself (detailed illustration of different error types available in Appendix A.2). As shown in Figure 2, the results show a notable similarity between these two sets of error patterns, with incomplete hallucinations being the most common, followed by fabricated hallucinations. This similarity suggests that retrieval heads potentially explain contextual faithfulness in LFQA.

## 3 Methodology

In this section, we introduce RHIO, a framework that enhances the contextual faithfulness of LLMs by explicitly teaching them to discriminate between faithful and unfaithful outputs, as illustrated in Figure 3. To achieve this, our approach involves augmenting realistic unfaithful samples, which enables the model to learn to distinguish faithful and unfaithful outputs conditioned on different control tokens. These control tokens are further leveraged to induce contrastive outputs and amplify the difference between them via contrastive decoding.

### 3.1 Unfaithful Data Augmentation

When it comes to negative data augmentation, the key challenge is how to augment samples that simulate the actual error types from the LLM itself. Prior approaches (Mishra et al., 2024) typically focus on entity-centric permutations, which can result in incoherence and low error coverage, thus limiting their utility in mimicking actual error types. Inspired by our pilot study in Section (§2.2), we discovered a salient correlation between LFQA faithfulness and retrieval heads. This finding motivates us to generate more realistic unfaithful outputs by simply masking out retrieval heads of LLMs.

Formally, we define the attention matrix for a given layer as  $\mathbf{A} \in \mathbb{R}^{H \times L \times L}$ , where  $H$  is the number of attention heads, and  $L$  is the sequence length. Let  $\mathcal{R}$  represent the indices of the top- $N$  retrieval heads identified by the detection algorithm in Wu et al. (2024). The masking operation is defined as:

$$\mathbf{A}'_{hij} = \begin{cases} 0 & \text{if } h \in \mathcal{R} \\ \mathbf{A}_{hij} & \text{otherwise} \end{cases}$$

where  $\mathbf{A}'$  is the modified attention matrix with weights of the top- $N$  retrieval heads set to zero.  $h$  indexes the attention heads. Here, we set  $N = 100$ .

### 3.2 Faithfulness-Aware Tuning

Beyond simply fine-tuning LLMs on faithful data, which encourages them to imitate positive behaviors, we argue that the key to improving faithfulness lies in enhancing the model’s ability to perceive and distinguish different types of unfaithfulness errors. To this end, we propose Faithfulness-Aware Tuning (FAT), which aims at teaching LLMs to discriminate faithful responses from unfaithful ones.

To explicitly distinguish the contrastive outputs, we draw inspiration from controllable text generation (Lu et al., 2022) and introduce two special *control codes*, [POS] and [NEG]. These special control tokens signal the generation of either faithful or unfaithful output. Specifically, we prepend a prefix at the beginning of the generation, *e.g.*, [POS] for the faithful output  $y^+$  and [NEG] for the unfaithful output  $y^-$  and then supervise the model to generate corresponding outputs conditioned on the control codes. Our training objectives are twofold:

$$\begin{aligned} \mathcal{L}(\theta) = & - \underbrace{\mathbb{E}_{(\mathbf{x}, \mathbf{c}, y^+)} [\log p_\theta(y^+ | [\text{POS}] \oplus \mathbf{x}, \mathbf{c})]}_{\text{enhancing faithfulness perception}} \\ & - \underbrace{\mathbb{E}_{(\mathbf{x}, \mathbf{c}, y^-)} [\log p_\theta(y^- | [\text{NEG}] \oplus \mathbf{x}, \mathbf{c})]}_{\text{enhancing unfaithfulness perception}} \end{aligned} \quad (1)$$

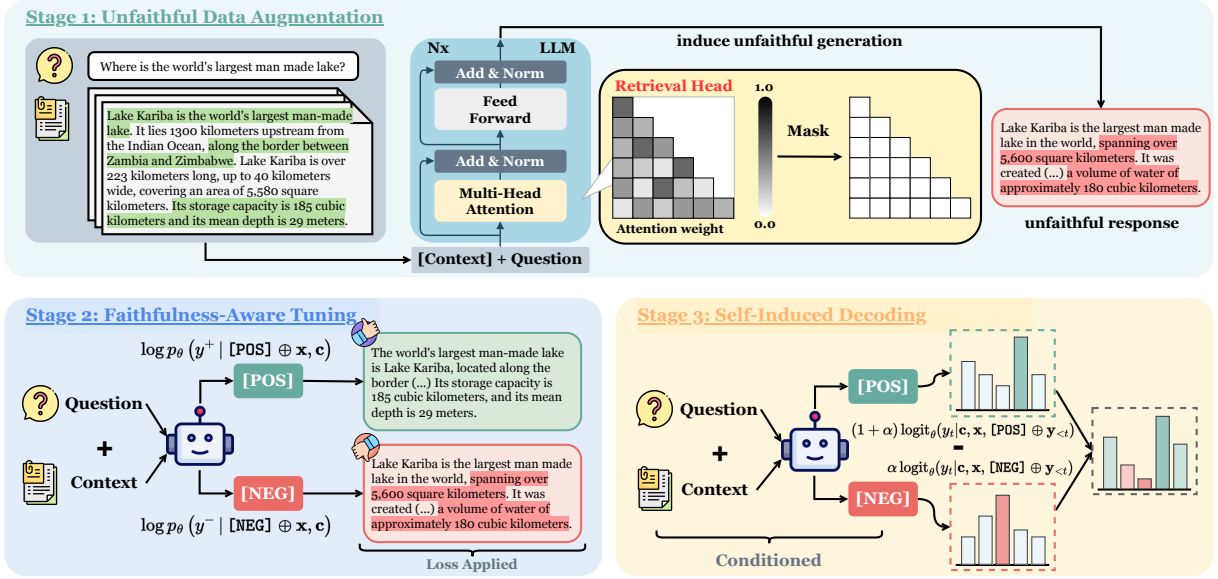


Figure 3: An overview of RHIO: (1) unfaithful data augmentation (§3.1), which augments unfaithful output by masking out attention heads responsible for contextual faithfulness; (2) faithfulness-aware tuning (§3.2), which teaches LLMs to explicitly discriminate between faithful and unfaithful outputs; (3) self-induced decoding (§3.3), which further enhances faithfulness by amplifying the differences between induced contrastive outputs.

where  $\mathbf{x}$  and  $\mathbf{c}$  represent the question and context, respectively, and  $\oplus$  denotes the concatenation of the control code with the question and context.

### 3.3 Self-Induced Decoding

During the inference phase, these trained control tokens enable the model to tailor its behavior. When prepending [POS], it steers the model to generate faithful responses. Conversely, it can also induce a hallucinated unfaithful response when prepending [NEG]. Inspired by studies on contrastive decoding (Li et al., 2023b; Shi et al., 2024), we introduce self-induced decoding (SID), which aims to further enhance faithfulness by amplifying the differences between induced contrastive outputs.

Concretely, consider an LLM parameterized by  $\theta$ , the model takes a question  $\mathbf{x}$  and source documents  $\mathbf{c}$  as input to generate a long-form response  $\mathbf{y}$ . We prepend the two control codes before generation, the next-token probability distribution is determined by amplifying the predictions from the faithful prediction and downplaying the induced unfaithful one, which can be formed as:

$$y_t \sim \text{Softmax} \left[ \left( (1 + \alpha) \log p_\theta(y_t | \mathbf{c}, \mathbf{x}, [\text{POS}] \oplus \mathbf{y}_{<t}) - \alpha \log p_\theta(y_t | \mathbf{c}, \mathbf{x}, [\text{NEG}] \oplus \mathbf{y}_{<t}) \right) / \tau \right] \quad (2)$$

Here,  $\tau$  is the temperature in the sampling decoding strategy. A larger  $\alpha$  indicates more weight focus on the context  $\mathbf{c}$  and  $\alpha = 0$  reduces to vanilla decoding. We set  $\alpha = 0.2$  in our experiments.

## 4 Benchmark: GroundBench

In this section, we introduce GroundBench, a benchmark tailored for evaluating the faithfulness of LFQA, and present the data collection pipeline as well as the manual effort for constructing GroundBench.

### 4.1 Data Collection

Several datasets (Fan et al., 2019; Stelmakh et al., 2022) have been introduced to assess the performance of LLMs on retrieval-augmented LFQA across various domains and evaluation dimensions. However, no existing dataset is specifically designed to measure faithfulness. This gap largely stems from the lack of guarantee that the retrieved documents contain sufficient information to answer the questions, a primary cause of unfaithfulness in LFQA (Chen et al., 2023). This limitation complicates the evaluation of LLMs' faithfulness.

To bridge the gap, we introduce GroundBench, an aggregated benchmark composed of five LFQA datasets, including ELI5-WebGPT (Nakano et al., 2021), ExpertQA (Malaviya et al., 2023), HA-GRID (Kamalloo et al., 2023), CLAPNQ (Rosenthal et al., 2024), and QuoteSum (Schuster et al., 2024). These datasets feature a wide range of queries, retrieval sources, and varying levels of difficulty, providing a comprehensive evaluation testbed. For detailed statistics, see Appendix B.



## 4.2 Dataset Specific Settings

To facilitate the evaluation of faithfulness in LFQA, we implement a controlled setting for each dataset to ensure the provided documents contain sufficient information to answer the questions.

**ELI5-WebGPT** consists of questions sourced from the “Explain Like I’m Five” subreddit. We use golden documents collected by human annotators via commercial search engines, which are deemed relevant and sufficiently informative to answer the questions. Each question is also equipped with human-labeled answers.

**ExpertQA** contains information-seeking questions formulated by experts spanning 32 fields, each paired with relevant documents and expert-verified answers. Considering that some of the annotated data have missing textual forms of evidence, we manually curated data with expert-revised evidence for GroundBench.

**HAGRID** includes questions designed for information-seeking scenarios, each accompanied by a set of manually labeled relevant documents and an answer generated by LLMs. We manually select those entries the answers are considered both informative and well-grounded by human evaluators for GroundBench.

**CLAPNQ** is built on real web search queries sampled from Natural Questions (Kwiatkowski et al., 2019). It features a gold document for each question with a corresponding grounded long-form answer. Additionally, we retrieve four additional documents from Wikipedia to create a multi-source information synthesis scenario.

**QuoteSum** is a semi-extractive LFQA dataset that involves questions with relevant documents and human-written answers derived explicitly from extracted spans across multiple sources. Specifically, we include the test subset in GroundBench.

## 5 Experimental Setup

### 5.1 Evaluation

All methods are evaluated on GroundBench, combining both automatic and human evaluation.

#### 5.1.1 Automatic Metrics

We evaluate long-form responses mainly on two dimensions: **Overall Quality** and **Faithfulness**.

**Overall Quality** measures the alignment of the model response with the human-labeled response. For CLAPNQ, ExpertQA, and HAGRID, we employ ROUGE-L to evaluate overall quality. For ELI5-WebGPT, we calculate the claim recall between the model response and golden sub-claims, following Chen et al. (2023). Regarding QuoteSum, we employ SEMQA (Schuster et al., 2024) to measure the overall quality. For further details about the automatic metrics, please refer to Appendix C.

**Faithfulness** measures the extent to which the model response is grounded in the provided context. We use MiniCheck (Tang et al., 2024)<sup>3</sup> to assess the consistency between each statement in the response and the documents. The average for all statements in each dataset is reported as the FaithScore.

#### 5.1.2 Human Evaluation.

Recent research (Xu et al., 2023b) highlighted the challenges of evaluating LFQA. To this end, we engage human annotators to manually analyze 200 generations from selected methods across the following dimensions. Each generation is rated on a 5-point Likert scale for each dimension. For more details, please refer to Appendix D.

**Faithfulness** evaluates whether the answer is *fully supported*, *partially supported*, or *not supported* by the provided documents.

**Completeness** measures whether all relevant information in the context is included to generate an informative answer to the question.

### 5.2 Baselines

We compare RHIO with the following baselines.

**Prompting-based Method** simply prompts LLMs to generate long-form responses that are faithful to the provided context for each question. Specifically, we evaluate several SOTA proprietary LLMs with carefully designed prompts, including GPT-4o and GPT4o mini<sup>4</sup>, as well as open-source models of varying sizes, covering models from the Llama-3.1 family and Mistral-NeMo-12B-Instruct.

**Supervised Fine-Tuning (SFT)** directly fine-tunes the model with high-quality LFQA data, aiming at teaching the model to utilize contextual information to generate faithful responses.

<sup>3</sup><https://huggingface.co/bespokelabs/Bespoke-MiniCheck-7B>

<sup>4</sup>Specifically, we utilize gpt-4o-2024-08-06 and gpt-4o-mini-2024-07-18 version for evaluation.

Method	CLAPNQ		ExpertQA		HAGRID		ELI5-WebGPT		QuoteSum		Avg. Faith.
	RL.	Faith.	RL.	Faith.	RL.	Faith.	Claim.	Faith.	SEM.	Faith.	
GPT-4o	40.53	91.81	<b>46.34</b>	69.48	<u>57.76</u>	<b>90.86</b>	<u>59.04</u>	81.00	42.56	78.51	82.33
GPT-4o-mini	37.72	90.35	45.30	66.53	54.87	87.94	56.09	81.89	40.74	78.16	80.97
Llama-3.1-70B-Instruct	39.44	88.64	<b>43.02</b>	69.35	49.21	79.08	<u>51.66</u>	74.87	41.24	67.42	75.87
Mistral-NeMo-12B-Instruct	<u>35.28</u>	78.71	42.76	54.19	<u>53.05</u>	80.16	<b>53.84</b>	65.06	39.50	69.85	69.59
Llama-3.1-8B-Instruct	17.14	58.47	31.67	51.22	<u>16.47</u>	55.80	47.11	55.74	25.96	41.70	52.59
<b>Llama-2-7B</b>											
+ SFT	37.43	82.46	42.47	68.14	50.89	78.14	47.85	73.25	40.81	62.93	72.98
+ RECOMP	30.88	63.52	39.72	51.95	38.30	62.72	40.22	49.45	37.61	54.98	56.52
+ Self-RAG	36.22	86.11	26.31	58.53	43.14	71.77	19.19	63.22	40.64	63.36	68.60
+ RHIO	<u>38.71</u>	<u>90.34</u>	39.34	<u>75.49</u>	49.59	<u>86.43</u>	42.07	<u>82.51</u>	<u>42.81</u>	<u>76.98</u>	<u>82.35</u>
<b>Llama-2-13B</b>											
+ SFT	37.18	84.72	<u>42.85</u>	70.87	<b>53.88</b>	78.07	47.72	71.61	40.64	66.72	74.40
+ RECOMP	30.62	64.91	<u>40.18</u>	52.18	37.65	59.85	39.36	51.96	37.62	51.75	56.13
+ Self-RAG	31.34	61.62	21.72	56.05	42.62	56.27	10.82	36.87	40.69	67.40	55.64
+ RHIO	<b>41.14</b>	<b>92.14</b>	39.65	<b>77.86</b>	50.64	<b>87.18</b>	42.84	<b>82.94</b>	<b>42.82</b>	<b>78.73</b>	<b>83.77</b>

Table 2: Experimental results on GroundBench. **Bold** and underline numbers indicate the best performance and second performance among non-proprietary models. And **gray-colored bold** text indicates the best proprietary model when it outperforms all non-proprietary models. We also include the experimental results of the Mistral series model (Jiang et al., 2023) in Appendix G.

**RECOMP** (Xu et al., 2023a) train an abstractive summarization model to filter irrelevant information in the context, avoiding the impact of noisy documents on faithful generation.

**Self-RAG** (Asai et al., 2024a) utilize special reflection tokens to teach the model to assess the retrieval quality and self-reflect with its generation, thereby improving faithfulness.

### 5.3 Implementation Details

To ensure a fair comparison, we employ the Llama-2-7B and Llama-2-13B as backbones for all the training-based baselines. We train models using the long-form split of FRONT dataset (Huang et al., 2024a), which consists of diverse real-world user queries for information-seeking. In the decoding stage, to avoid low-quality outputs during long-form generation, we uniformly adopt a sampling decoding strategy with a temperature of 1 and top-p of 0.95. For more implementation details of baselines and RHIO, please refer to Appendix E.

## 6 Results

### 6.1 Main Results

We present the results of GroundBench in Table 2.

**Existing LLMs are struggling to generate faithful responses.** We can observe that GPT-4o achieves the best performance across five datasets, with an average faithfulness score of 82.33%, while

GPT-4o-mini ranks in the second tier. The performance of open-source models shows a clear downward trend as the model size decreases. Even for the most powerful LLMs, achieving full contextual faithfulness remains a significant challenge. Crucially, when it comes to the smaller LLM, *e.g.*, Llama-3.1-8B-Instruct, it exhibits a substantial performance gap of 23.28% compared to its larger counterpart. This indicates the challenge for smaller LMs to maintain faithfulness to the context.

**LLMs fine-tuned to utilize contextual information achieve better faithfulness.** Compared with other training-based baselines, simply fine-tuning LLMs to utilize contextual information (SFT) results in more faithful long-form responses, leading to an average improvement of 30.84% and 20.05% compared with RECOMP and Self-RAG, respectively. This finding highlights the effectiveness of training models with attribution in mind to enhance faithfulness, which also aligns with prior work (Chen et al., 2023).

**RHIO significantly improves contextual faithfulness, even surpassing GPT-4o.** For both the 7B and 13B models, RHIO significantly improves faithfulness compared to all training-based baselines in GroundBench, leading to an average improvement of 12.84% and 12.59% in faithfulness for the 7B and 13B models, respectively. Notably, RHIO even outperforms the state-of-the-art GPT-4o by 1.74%.

Method	CLAPNQ		ExpertQA		HAGRID		ELI5-WebGPT		QuoteSum		Avg. Faith.
	RL.	Faith.	RL.	Faith.	RL.	Faith.	Claim.	Faith.	SEM.	Faith.	
<i>Llama-2-7B</i>											
+ RHIO	<b>38.71</b>	<b>90.34</b>	39.34	<b>75.49</b>	49.59	<b>86.43</b>	42.07	<b>82.51</b>	<b>42.81</b>	<b>76.98</b>	<b>82.35</b>
w/o SID	38.03	88.57	41.17	73.37	48.96	84.54	43.05	79.00	41.81	74.68	80.03
w/o FAT	37.43	82.46	<b>42.47</b>	68.14	<b>50.89</b>	78.14	<b>47.85</b>	73.25	40.81	62.93	72.98
<i>Llama-2-13B</i>											
+ RHIO	<b>41.14</b>	<b>92.14</b>	39.65	<b>77.86</b>	50.64	<b>87.18</b>	42.84	<b>82.94</b>	<b>42.82</b>	<b>78.73</b>	<b>83.77</b>
w/o SID	37.74	88.33	40.28	73.23	49.55	85.40	<b>48.13</b>	79.54	42.43	75.58	80.42
w/o FAT	37.18	84.72	<b>42.85</b>	70.87	<b>53.88</b>	78.07	47.72	71.61	40.64	66.72	74.40

Table 3: Ablation experimental results on GroundBench. The best results for each dataset are in **bold**.

This indicates that teaching the model to distinguish between faithfulness and unfaithfulness can effectively establish a mechanism for faithful generation.

## 6.2 Ablation Studies

We conduct ablation studies to verify the effectiveness of RHIO, and results are shown in Table 3.

**Effect of faithfulness-aware tuning.** To understand the significance of faithfulness-aware tuning (FAT), we compare RHIO with models that are solely fine-tuned using faithful samples. As shown in Table 3, omitting FAT leads to a substantial decrease in faithfulness (82.35%  $\rightarrow$  72.98% in 7B and 83.77%  $\rightarrow$  74.40% in 13B, respectively). This highlights the effectiveness of unfaithful samples in enhancing the perception of faithfulness in LLMs.

**Effect of self-induced decoding.** To demonstrate the effectiveness of self-induced decoding (SID), we compare RHIO with models that use the vanilla decoding strategy ( $\alpha = 0$ ). As indicated in Table 3, SID further improves average faithfulness by 2.90% and 4.17% in 7B and 13B models, respectively. Besides, we further investigate the effect of varying  $\alpha$  on faithfulness by adjusting  $\alpha$  from 0.1 to 0.5. As shown in Figure 4 (a), as  $\alpha$  increases, the model’s performance on faithfulness first and gradually decreases, reaching its best performance at  $\alpha = 0.2$ . We further compare SID with context-aware decoding (CAD), which amplifies the difference between output probabilities with and without context. As shown in Figure 4 (b), SID slightly outperforms CAD in faithfulness. This indicates that SID, by leveraging more diverse model-intrinsic error types induced by [NEG], more effectively improves model faithfulness than merely amplifying context-aware contrasts.




Figure 4: Ablation study on hyperparameter  $\alpha$  and the decoding strategy in self-induced decoding.




Figure 5: Ablation study on different negative sample augmentation strategies.

## 6.3 Further Analysis

**Does masking out retrieval heads outperform other augmentation strategies?** To demonstrate the effectiveness of utilizing negative samples triggered by masking out retrieval heads in improving faithfulness, we conduct a comparative ablation study by employing various negative sample augmentation strategies, including entity replacement, relation corruption, and direct prompting. Detailed implements for these augmentation strategies are presented in Appendix F. As shown in Figure 5, entity and relationship perturbations yield only marginal improvement. Meanwhile, compared to prompting methods, unfaithful samples induced by retrieval heads lead to more significant enhancements in faithfulness.

**Does masking out retrieval heads facilitate preference learning for improving contextual faithfulness?** Preference learning has demonstrated considerable effectiveness in aligning LLMs with human preferences through the utilization of *chosen-rejected* preference pairs. In this study, we explore the potential of leveraging masking retrieval heads to augment unfaithful samples for aligning LLMs in contextual faithfulness. Specifically, we directly employ Direct Preference Optimization (DPO) (Rafailov et al., 2023) on top of the SFT baseline. During the alignment stage, we further train the SFT model for two epochs, with a learning rate of  $5e-7$  and set the  $\beta$  to 0.1. The experimental results, presented in Table 4, indicate that DPO significantly enhances the contextual faithfulness of the SFT model across five datasets on GroundBench. This also suggests that our negative data augmentation strategy, by leveraging the intrinsic mechanism of existing LLMs’ retrieval heads, successfully generates high-quality preference data without requiring manual annotations or external models. This capability further confirms the effectiveness and broad applicability of our approach in improving contextual faithfulness through preference alignment.

Method	CLAPNQ	ExpertQA	HAGRID	ELI5-WebGPT	QuoteSum
<b>Llama-2-7B</b>					
+ SFT	<b>82.46</b>	69.14	78.14	73.25	62.93
+ DPO	82.13	<b>75.34</b>	<b>80.00</b>	<b>77.61</b>	<b>69.09</b>
<b>Llama-2-13B</b>					
+ SFT	84.72	70.87	78.07	71.61	66.72
+ DPO	<b>85.40</b>	<b>76.48</b>	<b>83.25</b>	<b>80.10</b>	<b>73.72</b>

Table 4: Preference learning results on GroundBench. The best results for each dataset are in **bold**.

**Does self-induced unfaithful samples perform better?** To analyze this, we separately trained two models using unfaithful outputs generated by masking retrieval heads from different models. As shown in Figure 9, self-induced negative samples lead to more improvement in faithfulness. We posit that self-induced samples, embedded with intrinsic model-specific errors, provide a more effective learning environment for the models to distinguish between faithful and unfaithful outputs. In addition, we explore the impact of the number of masked retrieval heads on performance in Appendix G.

## 7 Human Evaluation

Human evaluation results, detailed in Table 5, indicate that RHIO generates significantly more

	Faithfulness			Completeness
	Full	Partial	No	Rate
GPT-4o	86.5%	8.1%	5.4%	<b>4.2</b>
Llama-3.1-70B-Instruct	82.6%	10.2%	7.2%	3.7
SFT-13B	76.8%	13.8%	9.4%	3.2
RHIO-13B (Ours)	<b>87.5%</b>	<b>7.6%</b>	<b>4.9%</b>	<u>3.8</u>

Table 5: Human evaluation results on faithfulness and completeness. **Bold** numbers indicate the best performance. “\_” indicates the second-best performance.

grounded responses compared to all baselines. We show some case studies of RHIO in Appendix H.

## 8 Related Work

Recently, the demand for contextual LLMs has continued to grow, particularly in retrieval-augmented generation (RAG) scenarios. Despite its significant progress, hallucinations (Huang et al., 2023) remains a critical challenge. Generally, hallucination in LLMs can be categorized into two types: *factuality hallucination*, where generated content deviates from established world knowledge, and *faithfulness hallucination*, where the generated response is inconsistent with the provided context. In this work, we focus on faithfulness hallucination, particularly in long-form question-answering (LFQA) (Fan et al., 2019; Han et al., 2024).

Unlike factoid QA (Kwiatkowski et al., 2019; Mallen et al., 2023), where the answer units often appear as short-form entities, LFQA is a challenging task, which requires synthesizing relevant information from retrieved documents to produce a complex, paragraph-length answer. Contextual faithfulness (Nguyen et al., 2024) in LFQA measures whether the model generation is supported by the provided context (Huang et al., 2024b). Recent study (Chen et al., 2023) reveals that unfaithfulness in LFQA mainly comes from retrieval failure, hallucinated facts, and incorrect synthesis. Moreover, Stolfo (2024) found that the unfaithfulness phenomenon is prevalent across different model sizes and highlighted the propensity of LLMs to blend correct information with hallucinated content.

Numerous efforts have been made to improve contextual faithfulness in LFQA. RECOMP (Xu et al., 2023a) leveraged a summarization model to filter out the irrelevant context in retrieved documents before the generation process. Self-RAG (Asai et al., 2024a) designed special reflection tokens to teach models self-reflection abilities. This enables the model to adaptively retrieve documents on-demand, and criticize its own output to improve the faithfulness and overall quality. Shi



et al. (2024); Wang et al. (2024) combine pointwise mutual information with contrastive decoding (Li et al., 2023b) to let the model focus more on contextual information to reduce unfaithfulness. Despite these efforts, these approaches are more of a compensatory way to improve faithfulness, rather than developing models with a *built-in* mechanism for faithfulness generation.

In contrast to these methods, our work draws inspiration from retrieval heads (Wu et al., 2024), which play a fundamental role in contextual faithfulness. This allows us to augment realistic unfaithful samples and improve the perception of faithfulness by teaching the model to learn to distinguish between faithful and unfaithful responses.

## 9 Conclusion

This work proposes RHIO, a framework designed to enhance the contextual faithfulness of LLMs. RHIO first augments realistic unfaithful outputs by selectively masking out attention heads responsible for information retrieval. Given both faithful and unfaithful samples, RHIO trains the model to explicitly distinguish between them, conditioned on faithfulness control codes. RHIO further enhances faithfulness with contrastive decoding, aimed at amplifying the differences between contrastive outputs induced by control codes. Additionally, we also introduce GroundBench, a comprehensive benchmark compiled from five diverse LFQA datasets, providing a controllable evaluation of faithfulness in LLMs. Extensive evaluations on GroundBench demonstrate that RHIO significantly improves faithfulness, even surpassing GPT-4o.

## Limitations

This work exhibits several limitations worth noting. Firstly, to ensure a fair comparison, our experiments primarily employ the Llama-2 series models, which are utilized in the baseline models, *e.g.*, Self-RAG and RECOMP. Future work will involve extending our method to encompass additional series models. Secondly, to facilitate the evaluation of faithfulness in retrieval-augmented LFQA, GroundBench is constructed within a controlled setting, where the context is assured to contain sufficient information to answer the question. Consequently, faithfulness hallucination caused by retrieval failures is not discussed in GroundBench. We aim to extend the benchmark to cover a diverse range of task scenarios in future work, thus providing a more

comprehensive evaluation of contextual faithfulness in LLMs. Thirdly, while our findings demonstrate that retrieval heads can produce a diverse array of realistic and model-intrinsic unfaithful errors, the current study does not control the generation of specific types of unfaithful errors. Future research could explore augmenting specific error types by delving into the distinct contributions of various retrieval heads.

## Ethics Statement

In this work, all data used for constructing GroundBench derive from open-source public datasets, and no additional collection of sensitive information was conducted. Throughout the experimental process, all data and models were strictly utilized following their intended purposes and respective licenses. Our methodology aims to enhance the contextual faithfulness of retrieval-augmented LLMs, which positively impacts information-seeking systems by improving the transparency of LLMs in real-world applications. However, when deployed, our approach still carries inherent issues associated with LLMs, such as the potential for generating biased, harmful, or offensive output. Aside from this, to the best of our knowledge, there are no additional ethical issues associated with this paper.

## Acknowledgments

Xiaocheng Feng is the corresponding author of this work. We thank the anonymous reviewers for their insightful comments. This work was supported by the National Natural Science Foundation of China (NSFC) (grant 62276078, U22B2059), the Key R&D Program of Heilongjiang via grant 2022ZX01A32, and the Fundamental Research Funds for the Central Universities (Grant No.HIT.OCEF.2023018).

## References

- Akari Asai, Zeqiu Wu, Yizhong Wang, Avirup Sil, and Hannaneh Hajishirzi. 2024a. [Self-rag: Learning to retrieve, generate, and critique through self-reflection](#). In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.
- Akari Asai, Zexuan Zhong, Danqi Chen, Pang Wei Koh, Luke Zettlemoyer, Hannaneh Hajishirzi, and Wen-tau Yih. 2024b. [Reliable, adaptable, and attributable language models with retrieval](#). *CoRR*, abs/2403.03187.

- Hung-Ting Chen, Fangyuan Xu, Shane A. Arora, and Eunsol Choi. 2023. [Understanding retrieval augmentation for long-form question answering](#). *CoRR*, abs/2310.12150.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, Albert Webson, Shixiang Shane Gu, Zhuyun Dai, Mirac Suzgun, Xinyun Chen, Aakanksha Chowdhery, Alex Castro-Ros, Marie Pellat, Kevin Robinson, Dasha Valter, Sharan Narang, Gaurav Mishra, Adams Yu, Vincent Y. Zhao, Yanping Huang, Andrew M. Dai, Hongkun Yu, Slav Petrov, Ed H. Chi, Jeff Dean, Jacob Devlin, Adam Roberts, Denny Zhou, Quoc V. Le, and Jason Wei. 2024. [Scaling instruction-finetuned language models](#). *J. Mach. Learn. Res.*, 25:70:1–70:53.
- Angela Fan, Yacine Jernite, Ethan Perez, David Grangier, Jason Weston, and Michael Auli. 2019. [ELI5: long form question answering](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 3558–3567. Association for Computational Linguistics.
- Yuchun Fan, Yongyu Mu, Yilin Wang, Lei Huang, Junhao Ruan, Bei Li, Tong Xiao, Shujian Huang, Xiaocheng Feng, and Jingbo Zhu. 2025. [SLAM: towards efficient multilingual reasoning via selective language alignment](#). In *Proceedings of the 31st International Conference on Computational Linguistics, COLING 2025, Abu Dhabi, UAE, January 19-24, 2025*, pages 9499–9515. Association for Computational Linguistics.
- Tianyu Gao, Howard Yen, Jiatong Yu, and Danqi Chen. 2023a. [Enabling large language models to generate text with citations](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023*, pages 6465–6488. Association for Computational Linguistics.
- Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, Qianyu Guo, Meng Wang, and Haofen Wang. 2023b. [Retrieval-augmented generation for large language models: A survey](#). *CoRR*, abs/2312.10997.
- Rujun Han, Yuhao Zhang, Peng Qi, Yumo Xu, Jinyuan Wang, Lan Liu, William Yang Wang, Bonan Min, and Vittorio Castelli. 2024. [Rag-qa arena: Evaluating domain robustness for long-form retrieval augmented question answering](#). *Preprint*, arXiv:2407.13998.
- Lei Huang, Xiaocheng Feng, Weitao Ma, Yuxuan Gu, Weihong Zhong, Xiachong Feng, Weijiang Yu, Weihua Peng, Duyu Tang, Dandan Tu, and Bing Qin. 2024a. [Learning fine-grained grounded citations for attributed large language models](#). In *Findings of the Association for Computational Linguistics, ACL 2024, Bangkok, Thailand and virtual meeting, August 11-16, 2024*, pages 14095–14113. Association for Computational Linguistics.
- Lei Huang, Xiaocheng Feng, Weitao Ma, Liang Zhao, Yuchun Fan, Weihong Zhong, Dongliang Xu, Qing Yang, Hongtao Liu, and Bing Qin. 2024b. [Advancing large language model attribution through self-improving](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing, EMNLP 2024, Miami, FL, USA, November 12-16, 2024*, pages 3822–3836. Association for Computational Linguistics.
- Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, and Ting Liu. 2023. [A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions](#). *CoRR*, abs/2311.05232.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de Las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, L  lio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timoth  e Lacroix, and William El Sayed. 2023. [Mistral 7b](#). *CoRR*, abs/2310.06825.
- Ehsan Kamalloo, Aref Jafari, Xinyu Zhang, Nandan Thakur, and Jimmy Lin. 2023. [HAGRID: A human-llm collaborative dataset for generative information-seeking with attribution](#). *CoRR*, abs/2307.16883.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur P. Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, Kristina Toutanova, Llion Jones, Matthew Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. 2019. [Natural questions: a benchmark for question answering research](#). *Trans. Assoc. Comput. Linguistics*, 7:452–466.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. 2023. [Efficient memory management for large language model serving with pagedattention](#). In *Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles*.
- Junyi Li, Xiaoxue Cheng, Xin Zhao, Jian-Yun Nie, and Ji-Rong Wen. 2023a. [Halueval: A large-scale hallucination evaluation benchmark for large language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023*, pages 6449–6464. Association for Computational Linguistics.
- Xiang Lisa Li, Ari Holtzman, Daniel Fried, Percy Liang, Jason Eisner, Tatsunori Hashimoto, Luke Zettlemoyer, and Mike Lewis. 2023b. [Contrastive decoding: Open-ended text generation as optimization](#). In





Melanie Kambadur, Sharan Narang, Aurélien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. [Llama 2: Open foundation and fine-tuned chat models](#). *CoRR*, abs/2307.09288.

Han Wang, Archiki Prasad, Elias Stengel-Eskin, and Mohit Bansal. 2024. [Adacad: Adaptively decoding to balance conflicts between contextual and parametric knowledge](#). *Preprint*, arXiv:2409.07394.

Wenhao Wu, Yizhong Wang, Guangxuan Xiao, Hao Peng, and Yao Fu. 2024. [Retrieval head mechanistically explains long-context factuality](#). *CoRR*, abs/2404.15574.

Fangyuan Xu, Weijia Shi, and Eunsol Choi. 2023a. [RECOMP: improving retrieval-augmented lms with compression and selective augmentation](#). *CoRR*, abs/2310.04408.

Fangyuan Xu, Yixiao Song, Mohit Iyyer, and Eunsol Choi. 2023b. [A critical evaluation of evaluations for long-form question answering](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2023, Toronto, Canada, July 9-14, 2023*, pages 3225–3245. Association for Computational Linguistics.

## A Preliminary Study of Retrieval Heads

### A.1 Experiments for Randomly Masking

We provide the experimental results of random masking in Tabel 6. Specifically, we randomly mask out top-100 non-retrieval heads for each model. We can observe that randomly masking out non-retrieval heads does not have a significant impact on the faithfulness of the model.

Model	# R.(0)	# R.(100)	# Non-R.(100)
Llama-2-7B-Chat	80.14	35.85	77.12
Llama-2-13B-Chat	80.09	41.09	75.13
Llama-2-70B-Chat	86.36	51.49	88.16

Table 6: Experimental results of randomly masking. **R.(0)** represents masking 0 retrieval heads, **R.(100)** represents masking the top-100 retrieval heads, while **Non-R.(100)** represents randomly masking 100 non-retrieval heads. The values in the table represent the average faithfulness of the model response.

### A.2 Analysis of Error Patterns

To analyze the relationship between error types caused by masking retrieval heads and error types of unfaithful responses generated by the model itself, we extracted 50 unfaithful responses produced by the Llama-2-7B-Chat model on the CLAPNQ dataset. We manually identified and categorized several patterns of unfaithfulness as follows.

**Fabricated Hallucination.** This error type refers to LLMs generating completely fabricated facts that are not derived from the provided context. Such unfaithfulness typically arises from retrieval failures where the retriever is unable to retrieve relevant information related to the question, or from the model relying solely on its parametric knowledge to generate an answer. In our GroundBench, the documents within the context contain sufficient information to answer the questions. Therefore, this type of unfaithfulness typically stems from the latter scenario.

**Incomplete Hallucination.** This error type refers to the model generating content that is partially faithful, with the remainder of the content being fabricated. It generally occurs when the model captures only partial relevant information or when the context contains incomplete information. Consequently, the model may utilize its parametric knowledge to complete the response, leading to incomplete hallucinations.

**Inconsistent Hallucination.** This error type refers to the model generating content that contradicts the provided context. Typically, this stems from the model inaccurately synthesizing information from separate documents, resulting in content that is inconsistent with the provided context.

## B Detailed Data Statistics of RHIO

We provide detailed dataset statistics of GroundBench in Table 7, covering the instance count, query sources, document sources, the average length of responses, as well as the average length of documents.

## C Details of Automatic Metrics

We provide a detailed description of the evaluation metrics employed to assess the overall quality and faithfulness of the model-generated responses.

**Overall Quality.** To evaluate the overall quality of model-generated responses, we compare them to human-labeled answers. We utilize different automatic metrics tailored to the specific formats of the human-labeled responses across various datasets. For ExpertQA, CLAPNQ, and HA-GRID, we directly use the ROUGE-L score to measure the n-gram similarity between the model’s responses and the human-labeled answers. For ELI5-WebGPT, which typically features longer responses, we follow the methodology described in



Dataset	Number	Question Sources	Response Sources	Document Sources	Response Avg. Len.	Documents Avg. Len.
ExpertQA	527	Domain Experts	Human & Bing-Chat & GPT-4	Web Search	134.6	617.0
ELI5-WebGPT	271	ELI5	Human	Web Search	106.0	268.7
CLAPNQ	300	NQ	Human	Wikipedia	51.7	632.2
HAGRID	496	MIRACL	Human & GPT-3.5	Wikipedia	41.15	302.2
QuoteSum	1,319	NQ & PAQ	Human	Wikipedia	60.3	319.1

Table 7: Statistics of datasets used in GroundBench.

Gao et al. (2023a) and use ChatGPT<sup>5</sup> to decompose the human answers into sub-claims. We then assess the overall quality by comparing the degree of entailment between the model’s responses and these sub-claims. For QuoteSum, which is designed as a semi-extractive QA task, models are required to synthesize information from multiple sources while explicitly extracting factual spans. The human-written semi-extractive answers predominantly consist of quoted answers. Following the evaluation protocol in Schuster et al. (2024), answers are assessed for fluency and attribution preciseness. Fluency is quantified using the ROUGE-L score. Preciseness is measured by calculating the normalized token-F1 score for each source separately and then averaging these scores across all sources, taking the maximum score per source across reference quoted answers. Finally, we compute the geometric mean of these scores, referred to as SEMQA, to reflect the overall answer quality.

**Faithfulness.** Faithfulness is the most critical evaluation dimension in GroundBench, assessing whether the model-generated response is fully supported by the provided documents. To evaluate faithfulness, we first segment the model-generated long-form response into statements by sentence boundaries<sup>6</sup>. We then assess the entailment of each sentence against the provided context using an NLI-based model. Specifically, we utilize Bespoke-MiniCheck-7B, the most advanced variant of MiniCheck (Tang et al., 2024) during the evaluation. The overall faithfulness score is calculated by the average score across all sentences in the response.

## D Details of Human Evaluation

We recruited two annotators, each holding at least a bachelor’s degree, to participate in the human evaluation. The evaluation focuses on two key aspects: faithfulness and completeness of the responses. For faithfulness, we selected the outputs from four systems’ across the five datasets in

GroundBench: GPT-4o, Llama-3.1-70B-Instruct, SFT-Llama-2-13B, and RHIO-13B, resulting in a total of 200 generations. Unlike the NLI-based model that outputs binary labels, we ask the annotators to perform a more fine-grained evaluation, determining whether each statement in the responses is fully supported, partially supported, or not supported by the provided context. For completeness, annotators assessed whether the responses fully addressed the given questions and captured all relevant information in the provided context. During the evaluation, both dimensions were rated using a 5-point Likert scale, capturing varying levels of faithfulness and completeness.

## E Details of Implementation

### E.1 Training Dataset

As for the training dataset, we utilize the FRONT (Huang et al., 2024a) dataset for training, selected for its diverse and high-quality retrieval documents refined through a rigorous data filtering process to ensure relevance to the questions. The dataset comprises 5,677 diverse user questions from the Natural Questions (NQ) dataset (Kwiatkowski et al., 2019), designed for long-form responses, and 2,431 questions for short-form factoid answers. Specifically, we only use the long-form split for training the baselines and RHIO.

### E.2 Prompts for Training and Evaluation

The prompt template employed in both the training and evaluation phases is shown in Figure 6. The prompt explicitly instructs the model to ensure the generation is well-grounded in the retrieved documents (Fan et al., 2025). We use this template for training the SFT baseline and RHIO, as well as for evaluating the SFT, prompting baseline, and RHIO.

### E.3 Baselines

The detailed baseline implementation is as follows:

**Prompting-based Method:** For each question in GroundBench, the method involves directly prompting LLMs to generate responses based on

<sup>5</sup>We use gpt-3.5-turbo-1106 version.

<sup>6</sup>Specifically, we employ the NLTK library for this segmentation.

(a) Prompt template used for training and evaluation

```
### Instruction:
Write an accurate, engaging, and concise answer to the given question using only the provided
retrieval documents. Ensure the answer is well-grounded in the relevant information,
disregarding irrelevant information in documents.

### Input:
Question: [Question]
Retrieved documents: [Documents]

### Response:
```

Figure 6: Illustration of the prompting used for training and evaluation.

the provided retrieved documents. To comprehensively evaluate GroundBench, we employ both proprietary and open-source LLMs. For proprietary LLMs, we use state-of-the-art models, including GPT-4o and its more efficient variant, GPT-4o-mini. Regarding open-sourced LLMs, we cover a range of model sizes, from the most powerful, such as Llama-3.1-70B-Instruct<sup>7</sup>, to smaller models like Mistral-NeMo-12B-Instruct<sup>8</sup> and Llama-3.1-8B-Instruct<sup>9</sup>. During the evaluation, we utilized the vLLM framework (Kwon et al., 2023) for efficient inference and adopted a sampling decoding strategy for all models, setting the temperature to 1.0 and top-p to 0.95. The maximum generation length was set to 512 tokens.

**Supervised Fine-Tuning:** The method involves fine-tuning the model on retrieved documents to generate long-form responses to the given questions. In this way, models are trained to leverage contextual evidence to produce a well-grounded response. Instead of fine-tuning using the existing responses from the FRONT dataset, which are generated by ChatGPT with specific instructions for attributed text generation (Huang et al., 2024a) and may introduce additional bias, we employ the Llama-3.1-70B-Instruct model to generate faithful responses as the supervised outputs, chosen for its remarkable faithfulness in long-form QA. During training, we use the Llama-2-7B and Llama-2-13B models. Specifically, we conduct full fine-tuning for three epochs using eight A100-80GB GPUs. The total batch size is set to 64, and the learning rate is maintained at  $2e-5$ . The maximum input

sequence length is set to 2,048 tokens. During decoding, we adopt the same decoding configuration as the prompting-based method.

**RECOMP (Xu et al., 2023a):** The method involves training an abstractive summarization model to filter irrelevant information from the context. Specifically, for each question and provided documents in FRONT, we first utilize GPT-3.5-turbo to summarize all the important information needed to answer the question. Following the setting in Xu et al. (2023a), we train an abstractive summarization model, *e.g.*, Flan-T5-Large (Chung et al., 2024) using the distilled summarization dataset above. The summarization model is then employed to summarize all documents in GroundBench, ensuring that irrelevant information is filtered out to avoid distraction during faithful generation. During the evaluation, we prepend the corresponding summarization to each question in GroundBench to guide the models in generating faithful long-form responses. In particular, we use the chat version of Llama-2-7b and Llama-2-13b for evaluation. The prompt templates used for both summarization and evaluation are presented in Figure 7.

**Self-RAG (Asai et al., 2024a):** The method trains an LLM that adaptively retrieves documents on-demand and generates and reflects on retrieved documents and its own generations using reflection tokens. These reflection tokens are categorized into retrieval and critique tokens to indicate the need for retrieval and the attributability of its generation, respectively. Specifically, we directly utilized models provided by the authors, trained using Llama-2-7b<sup>10</sup> and Llama-2-13b<sup>11</sup>. During the evaluation,

<sup>7</sup><https://huggingface.co/meta-llama/Meta-Llama-3.1-70B-Instruct>

<sup>8</sup><https://huggingface.co/nvidia/Mistral-NeMo-12B-Instruct>

<sup>9</sup><https://huggingface.co/meta-llama/Llama-3.1-8B-Instruct>

<sup>10</sup>[https://huggingface.co/selfrag/selfrag\\_llama2\\_7b](https://huggingface.co/selfrag/selfrag_llama2_7b)

<sup>11</sup>[https://huggingface.co/selfrag/selfrag\\_llama2\\_13b](https://huggingface.co/selfrag/selfrag_llama2_13b)

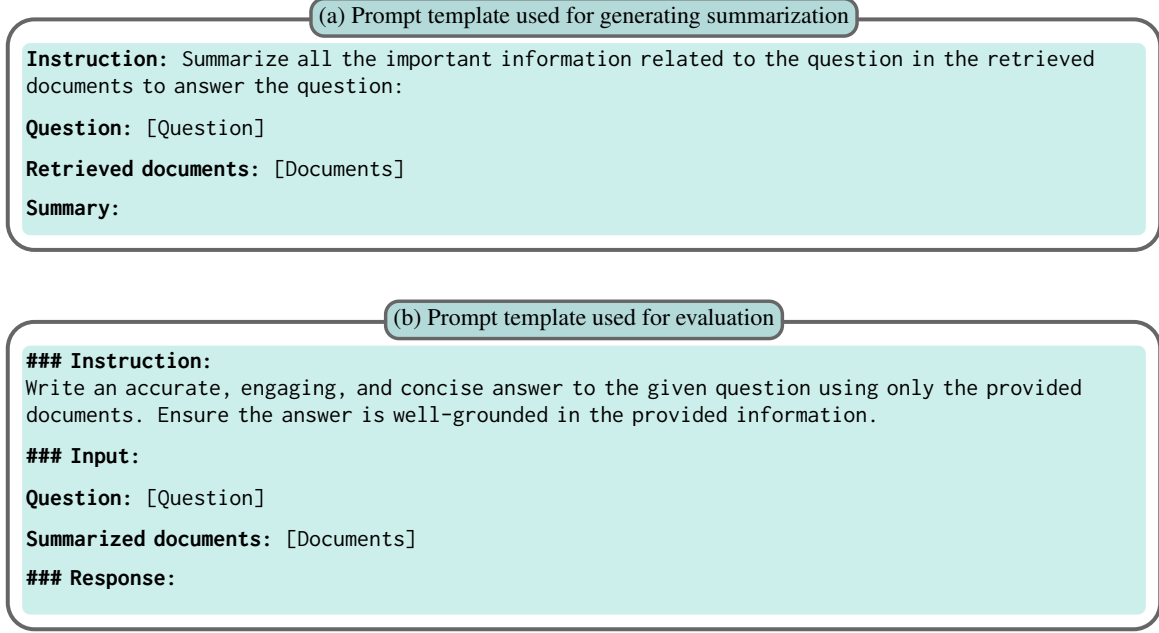


Figure 7: Illustration of the prompting used for RECOMP.

we employ the official prompt template designed for the ASQA task, using a maximum decoding length of 512 tokens. We used the default configuration for long-form generation, with relevance, attribution, and overall completeness critique rewards each weighted at 1.0. During decoding, we adopt a beam width of 7 at each segment level.

#### E.4 Implementantation for RHIO

RHIO consists of three main components: unfaithful data augmentation, faithfulness-aware tuning, and self-induced decoding.

**Unfaithful Data Augmentation:** In this stage, for each question and its corresponding documents in the FRONT dataset, we mask out the top 100 retrieval heads in both Llama-2-7B and 13B models. The responses generated from these models are categorized as unfaithful and are subsequently used for fine-tuning. This process yields 5,677 unfaithful responses for both the Llama-2-7B and Llama-2-13B models, respectively.

**Faithfulness-Aware Tuning:** In this phase, we utilize the unfaithful responses generated during the data augmentation stage along with the faithful responses from the SFT baseline training data for fine-tuning. Training is conducted on eight A100-80GB GPUs using Deepspeed (Rasley et al., 2020) stage 3 for efficient multi-GPU distribution, with training precision set to Bfloat16. We maintain a

total batch size of 64, a learning rate of  $2e-5$ , and a maximum input sequence length of 2,048 tokens. Both the 7B and 13B models are trained over 3 epochs.

**Self-Induced Decoding:** We set  $\alpha = 0.2$  for both RHIO-7B and RHIO-13B models, with a decoding temperature of 1.0 and a top-p setting of 0.95. The same prompt template used for the SFT baseline, as depicted in Figure 6, is employed for both training and evaluation.

#### F Negative Sample Augmentation Strategies

Due to our task scenario being retrieval-augmented LFQA, our goal is to generate unfaithful responses that are not faithful to the provided context. In LFQA, the main types of faithfulness hallucination include context inconsistency where generated responses directly contradict the context, mainly involving inconsistencies in entities and relationships (Mishra et al., 2024), as well as fabric hallucinations that completely make things up. Therefore, we have designed three strategies for generating negative samples based on entity-replacement, relation-replacement, and prompting-based approaches.

**Entity and Relation Replacement.** This strategy involves replacing key entities and relations in the provided faithful response with other plausible but

contextually irrelevant alternatives. We adopt the methodologies outlined in Mishra et al. (2024), instructing powerful LLMs, *e.g.*, Llama-3.1-70B-Instruct, to execute entity or relation replacements. This process is designed to generate a diverse array of unfaithfulness error types, resulting in synthetic responses that are naturally contradictory to the contextual information.

**Prompting-based.** We employ the method in Li et al. (2023a) for generating unfaithful samples. Specifically, we directly feed the complete instruction shown in Figure 8 into the corresponding model (Llama-2-7B-Chat and Llama-2-13B-Chat) to generate a response inconsistent with the context. In addition to instructing the model to produce a response inconsistent with the context, we also provide a faithful answer as a reference. This allows the model to generate diverse and multi-faceted unfaithful responses for each question. Upon manual review, it was found that the primary error type in unfaithful responses generated by this method focuses on fabric hallucination, which involves the model generating responses based on parameterized knowledge that is unrelated to the given context.

## G Additional Analysis Experiments

**Negative Samples from Different Models.** To explore the impact of negative samples induced from different models on faithfulness, we utilized the Llama-2-7B and Llama-2-13B models to generate negative samples by masking out their own top 100 retrieval heads. These self-induced negative samples were then employed to train each model separately. As shown in Figure 9 (a), models trained with their own negative samples, *e.g.*, Llama-2-7B with samples induced by itself (7B-Induced) and Llama-2-13B with samples induced by itself (13B-Induced), demonstrated superior performance compared to models trained with negative samples induced by the other model. This indicates that self-induced negative samples may be more closely aligned with each model’s specific error tendencies, thereby providing more effective training for improving faithfulness.

**Impact of the Number of Masked Retrieval heads.** We investigate the impact of the number of masked retrieval heads on model faithfulness. We conduct an ablation study to mask the top 50, 100, and 150 retrieval heads of Llama-2-7B to

generate respective sets of unfaithful samples. As shown in Figure 9 (b), increasing the number of masked retrieval heads from 50 to 100 led to an improvement in the model’s faithfulness, suggesting that masking more retrieval heads can enhance faithfulness by generating more challenging negative samples. However, further increasing the number of masked heads to 150 resulted in a decline in faithfulness. We hypothesize that excessively masking retrieval heads degrades the quality of the generated unfaithful samples, introducing issues such as reduced coherence. This degradation could hinder the model’s ability to effectively learn from these unfaithful samples, ultimately impacting the training process negatively.

**Extend to Mistral-series models.** In addition to the Llama series models, we also extend RHIO to Mistral series models. The results are shown in Table 8.

Method	CLAPNQ	ExpertQA	HAGRID	ELIS-WebGPT	QuoteSum
<i>Mistral-7B-v0.3</i>					
+ SFT	65.89	55.97	66.85	54.22	50.89
+ RHIO	<b>85.17</b>	<b>77.56</b>	<b>82.85</b>	<b>77.11</b>	<b>73.14</b>

Table 8: Experimental results on GroundBench. **Bold** indicates the best performance.

## H Case Study

We provide a case study in Table 9. It can be seen that compared to the response generated by GPT-4o-mini in Table 1, which contains hallucinated facts and inaccurate information synthesizing, our model demonstrates higher faithfulness and well-grounded in the provided documents.



(a) Prompt template used for training and evaluation

```
### Instruction:
I want you act as a hallucination answer generator. Given a question, provided documents, and
a relatively faithful answer, write a unfaithful answer that is not grounded in the provided
documents.

### Input:
Question: [Question]
Retrieved documents: [Documents]
Faithful answer: [Faithful Answer]

### Response:
```

Figure 8: Illustration of the prompting used for generating unfaithful responses.




Figure 9: Analysis of the impact of self-induced negative samples and the number of masked retrieval heads.

---

**### Instruction:**

Write an accurate, engaging, and concise answer to the given question using only the provided retrieval documents. Ensure the answer is well-grounded in the relevant information, disregarding irrelevant information in documents.

**### Input:**

**Question:** Difference between Russian Blue and British Blue cat

**Retrieved documents:**

Document [1]: The Russian Blue is a cat breed that comes in colors varying from a light shimmering silver to a darker, slate grey. They develop close bonds with their owners and are sought out as pets due to their personalities, beauty and coat. It is their short, dense coat which has been the hallmark of the Russian breed for more than a century. The dense coat stands out from the body and one can draw patterns in the coat that will stay until one smoothes them out again. They are also considered to be hypoallergenic.

Document [2]: The Russian Blue has bright green eyes, pinkish lavender or mauve paws, two layers of short thick fur, and a blue - grey coat. The color is a bluish - gray that is the dilute expression of the black gene. However, as dilute genes are recessive ("d") and each parent will have a set of two recessive genes ("dd") two non-CPC Russian Blues will always produce a blue cat. Due to the breeding with Siamese after World War II, there are colorpoint genes floating around. If two carriers are bred together, then they will produce a litter of mixed colors - solid blue or white with blue like a Siamese. People call these CPC cats "colorpoint" , "whites" or "pointed" Russians. In most registries, one can not register, breed or show a colorpoint Russian.

Document [3]: The coat is known as a "double coat", with the undercoat being soft, downy and equal in length to the guard hairs, which are an even blue with silver tips. However, the tail may have a few very dull, almost unnoticeable stripes. The coat is described as thick, plush and soft to the touch. The feeling is softer than the softest silk. The silver tips give the coat a shimmering appearance. Its eyes are almost always a dark and vivid green. Any white patches of fur or yellow eyes in adulthood are seen as flaws in show cats. Russian Blues should not be confused with British Blues (which are not a distinct breed, but rather a British Shorthair with a blue coat as the British Shorthair breed itself comes in a wide variety of colors and patterns), nor the Chartreux or Korat which are two other naturally occurring breeds of blue cats, although they have similar traits.

Document [4]: Russian Blue 1 year old Russian Blue male (American style) Other names Archangel Blue, Archangel Cat Origin Russia Breed standards CFA standard FIFe standard TICA standard AACE standard ACF standard ACFA / CAA standard CCA - AFC standard GCCF standard Notes The ACF and GCCF also recognize Russian Blues in white and black, however the CFA does not. In addition, ACFA recognizes Russian Shorthairs in white, black, and blue. Domestic cat (Felis catus)

Document [5]: The Russian Blue is a naturally occurring breed that may have originated in the port of Arkhangel'sk in Russia. They are also sometimes called Archangel Blues. It is believed that sailors took Russian Blues from the Archangel Isles to Great Britain and Northern Europe in the 1860s. The first recorded appearance outside of Russia was in 1875 at The Crystal Palace in England as the Archangel Cat. The Russian Blue competed in a class including all other blue cats until 1912, when it was given its own class. The breed was developed mainly in England and Scandinavia until after World War II.

**### Response:**

The Russian Blue cat is known for its bright green eyes, pinkish lavender or mauve paws, two layers of short thick fur, and a blue-grey coat. The color of the Russian Blue is a bluish-gray, which is the dilute expression of the black gene. On the other hand, British Blues are not a distinct breed but rather a British Shorthair with a blue coat. The British Shorthair breed itself comes in a wide variety of colors and patterns, unlike the Russian Blue, which is a specific breed known for its unique coat and characteristics.

---

Table 9: Case study of how RHIO helps the model to generate the faithful response from CLAPNQ dataset.