

Advancing Large Language Model Attribution through Self-Improving

Lei Huang¹, Xiaocheng Feng^{1,2*}, Weitao Ma¹, Liang Zhao¹, Yuchun Fan³,
Weihong Zhong¹, Dongliang Xu⁴, Qing Yang⁴, Hongtao Liu⁴, Bing Qin^{1,2}

¹ Harbin Institute of Technology, Harbin, China

² Peng Cheng Laboratory, Shenzhen, China ³ Northeastern University, Shenyang, China

⁴ Du Xiaoman Science Technology Co., Ltd., Beijing, China

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

yuchunfan_neu@outlook.com

{xudongliang, yangqing, liuhongtao01}@duxiaoman.com

Abstract

Teaching large language models (LLMs) to generate text with citations to evidence sources can mitigate hallucinations and enhance verifiability in information-seeking systems. However, improving this capability requires high-quality attribution data, which is costly and labor-intensive. Inspired by recent advances in self-improvement that enhance LLMs without manual annotation, we present START, a **Self-Taught AttRibuTion** framework for iteratively improving the attribution capability of LLMs. First, to prevent models from stagnating due to initially insufficient supervision signals, START leverages the model to self-construct synthetic training data for warming up. To further improve the model’s attribution ability, START iteratively utilizes fine-grained preference supervision signals constructed from its sampled responses to encourage robust, comprehensive, and attributable generation. Experiments on three open-domain question-answering datasets, covering long-form QA and multi-step reasoning, demonstrate significant performance gains of 25.13% on average without relying on human annotations and more advanced models. Further analysis reveals that START excels in aggregating information across multiple sources.

1 Introduction

The rapid development of large language models (LLMs) (OpenAI, 2023; Zhao et al., 2023) has led to their prosperity as indispensable tools for information seeking. Despite their remarkable capability to generate fluent and informative responses to user queries, LLMs also struggle with hallucinations (Huang et al., 2023). To facilitate factuality verification, recent research (Bohnet et al., 2022) has explored attributed text generation, a paradigm that enables LLMs to generate responses with citations. By attributing models’ output to verifiable

sources, it can improve the explainability and credibility of LLM-generated content (Li et al., 2023).

While beneficial, the ability to attribute contextual sources is not inherent in LLMs. Most work induces LLMs to generate text with citations via in-context learning (Gao et al., 2023), which is far from satisfactory (Liu et al., 2023). The current winning recipe for accurate attribution involves fine-tuning on high-quality attribution responses¹ (Li et al., 2024). However, acquiring such data typically requires either manual curation (Malaviya et al., 2023), or distilled from the most advanced LLMs (Huang et al., 2024a,b), both of which are costly and not scalable, thus limiting the growth of models’ attribution capability. One promising solution is self-improvement (Yuan et al., 2023), which has demonstrated the potential to boost model performance by learning from self-generated high-quality samples.

Inspired by this, we aim to explore the potential of self-improvement in bootstrapping the attribution ability of LLMs. However, achieving this goal presents several challenges. One significant challenge lies in the risk of *model stagnation* during the self-improvement process, primarily due to the insufficient supervision signals obtained in the early stage. Concretely, considering the inferior performance of LLMs in handling the attribution task (Gao et al., 2023), generating sufficient high-quality attribution responses solely through sampling proves difficult. This scarcity of high-quality samples limits the opportunities for LLMs to self-improve effectively. Another challenge stems from the limitation of *weak supervision signals*. Current self-improvement approaches (Yuan et al., 2023) primarily involve supervised fine-tuning on high-quality samples while discarding low-quality ones. When applied to LLM attribu-

*Corresponding Author

¹Attribution responses refers to “responses with in-line citations, e.g., [1][2]”.

tion, these high-quality samples provide only weak supervision signals, mainly teaching LLMs on the surface form of attribution (*e.g.*, proper citation format) (Li et al., 2024). Such practice may neglect the potential of exploring fine-grained signals from low-quality samples to learn what constitutes a desirable attribution response.

To address these challenges, we present **START**, a **Self-Taught AttRibuTion** framework designed to bootstrap the attribution capabilities of LLMs. To prevent models from stagnating early due to insufficient supervision signals, we first leverage the model to self-construct high-quality synthetic attribution data (§3.1). The data synthesis process follows *reverse attribution thinking*: the model initially generates a response to a given query, then breaks it into atomic claims, and finally randomly combines them to create synthetic documents. This process not only simulates multi-source information-seeking scenarios but also ensures precise attribution, as each document can be directly traced back to the specific claim it originated from. These high-quality synthetic data are then utilized for warming up, providing a good starting point for LLMs to self-improve. Furthermore, to better explore fine-grained supervision signals for LLM attribution, we introduce an iterative self-improving recipe (§3.2). Specifically, the framework meticulously designs fine-grained rewards tailored for LLM attribution, covering robustness, comprehensiveness, and attributability. By scoring multiple candidates through sampling and selecting those with the highest holistic rewards for supervised fine-tuning, the framework subsequently utilizes low-quality samples to construct fine-grained preference pairs with diverse optimization rewards for preference optimization. This iterative process further fosters the self-improvement of attribution capabilities.

We conduct extensive experiments across three open-domain question-answering datasets, covering long-form QA and multi-step reasoning. Results indicate that **START** achieves significant performance gains of 25.13% on average in citation quality. Moreover, **START** successfully achieves self-improvement in LLM attribution, showing progressive improvements across iterations. Ablation studies confirm that each component significantly contributes to the improvement. Further analysis shows that **START** not only excels in generating superior attributable responses but also in effectively aggregating information across multiple sources.

2 Related Work

2.1 Large Language Model Attribution

Attribution has gained significant attention for enhancing the interpretability and verifiability of LLMs (Gao et al., 2023; Li et al., 2023). Recent studies have focused on improving LLM attribution in a supervised way. Asai et al. (2023) first distill GPT-4 to collect high-quality attribution data, aiming to teach the model to generate grounded answers with citations through self-reflecting. Similarly, Huang et al. (2024a) develop a training framework starting with distilling ChatGPT, followed by designing reward models to teach the LLM to generate highly supportive and relevant citations. Additionally, Li et al. (2024) model the attribution task from a preference learning perspective, where they first fine-tune the model on human-labeled attribution datasets and then perform preference optimization using synthesized preference data. Furthermore, Huang et al. (2024b) take this further by extending the attribution format to a fine-grained citation level, primarily distilled from ChatGPT. It enables the model to first ground the fine-grained quotes within the context and then condition the generation process on them. In contrast to these methods, **START** aims to bootstrap attribution capability without relying on human-labeled data or distilling from more capable LLMs.

2.2 Self-Improvement for LLMs

High-quality data either human-crafted or distilled from advanced LLMs has proven effective in enhancing the performance of LLMs. However, acquiring such high-quality data can be prohibitively expensive. Recently, self-improvement approaches (Gülçehre et al., 2023; Yuan et al., 2024), where LLMs learn from self-generated samples have emerged as a viable solution to compensate for the scarcity of high-quality data. These methods typically involve employing heuristic rules (Zelikman et al., 2022), self-critique (Tian et al., 2024), or training additional verifiers (Hosseini et al., 2024) to assess the quality of model-generated samples. Such practices are particularly effective in reasoning tasks, *e.g.*, mathematical reasoning, where LLMs already demonstrate capable abilities and can receive precise feedback on correctness. However, these advantages are absent in the attribution task, due to its challenging nature. To bridge the gap, we take an initial step towards exploring the potential of self-improvement in LLM attribution.

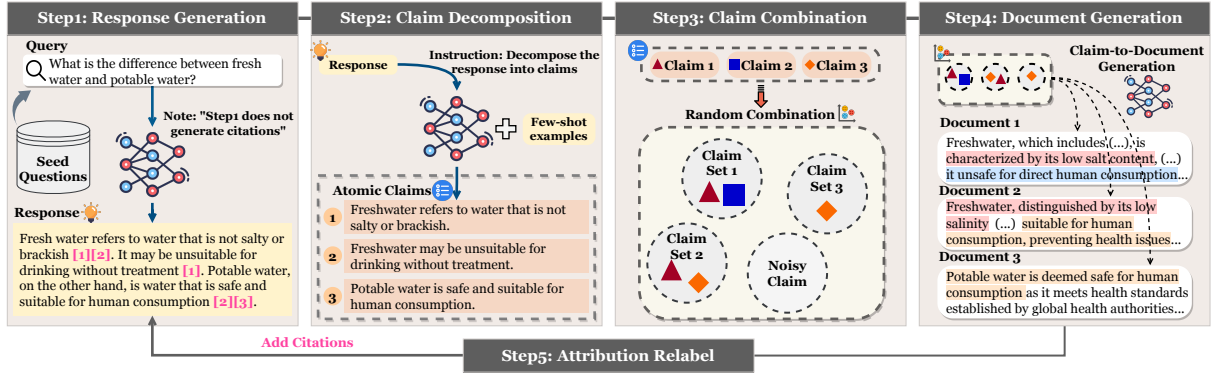


Figure 1: The data synthesis pipeline consists of five steps: given a user query, the LLM first generates an informative response **without citations** in a closed-book setting. Subsequently, the LLM decomposes this response into atomic claims. These claims are then **randomly** grouped into specific sets, which serve as the basis for generating documents that cover all included claims. Finally, we trace back to the initial response to relabel the citations.

3 Problem Formulation and Methodology

We follow a formulation of attributed text generation as described in Gao et al. (2023). This task involves processing a user query q for information-seeking, given a corpus of retrieved documents \mathcal{D} , to generate a response \mathcal{S} with in-line citations. We assume the response \mathcal{S} as consisting of n statements, such that $\mathcal{S} = \{s_1, s_2, \dots, s_n\}$. Each statement $s_i \in \mathcal{S}$ cites a list of passage $\mathcal{C}_i = \{c_{i1}, c_{i2}, \dots\}$, where $c_{ij} \in \mathcal{D}$. Citations are presented in the form of [1][2], which represent the attribution to specific documents in \mathcal{D} .

Next, we present an overview of START, a training framework designed to teach LLMs to self-improve their attribution ability, as illustrated in Figure 2. START consists of two essential stages: synthetic data warm-up (§3.1) and self-improving for LLM attribution (§3.2).

3.1 Synthetic Data Warm-Up

The core of self-improvement lies in generating *high-quality* samples and iteratively learning from them. Intuitively, a *high-quality* attribution response should not be distracted by irrelevant documents (*robustness*) and capture high coverage of viewpoints across multiple documents (*comprehensiveness*) while maintaining high citation quality (*attributability*). However, existing LLMs typically show inferior performance in the attribution task, significantly hindering their ability to generate such high-quality samples. This limitation poses substantial challenges to enhancing their attribution capabilities through self-improvement.

In this stage, we propose utilizing the model to self-construct high-quality synthetic data for warm-

ing up, enabling the model to have the basic ability to generate robust, comprehensive, and attributable responses across multiple sources. The pipeline consists of the following steps, shown in Figure 1. More details can be found in Appendix A.

Step 1: Response Generation Given an arbitrary model, we first sample a query q from seed questions Q and then generate a long-form answer S utilizing the parametric knowledge of the model itself. The model is required to produce informative answers that cover multiple perspectives.

Step 2: Claim Decomposition Prior work (Min et al., 2023) has explored using atomic claims as a fundamental unit in long-form text generation. Thus, for the response S , we ask the model to decompose it into atomic claims. Each atomic claim represents a distinct piece of information.

Step 3: Claim Combination To ensure that the response behaves as an aggregation of information from multiple documents, we randomly combine different claims into one claim set. This process helps simulate the natural diversity of viewpoints and sources, thus enhancing the comprehensiveness and realism of the synthesized responses.

Step 4: Document Generation For each claim set, we prompt the model to generate a synthetic document D that provides a comprehensive discussion of the grouped claims. Additionally, to enhance the robustness of the response, we introduce irrelevant documents by uniformly sampling documents generated from other queries.

Step 5: Attribution Relabel The final step involves labeling the response with citations from

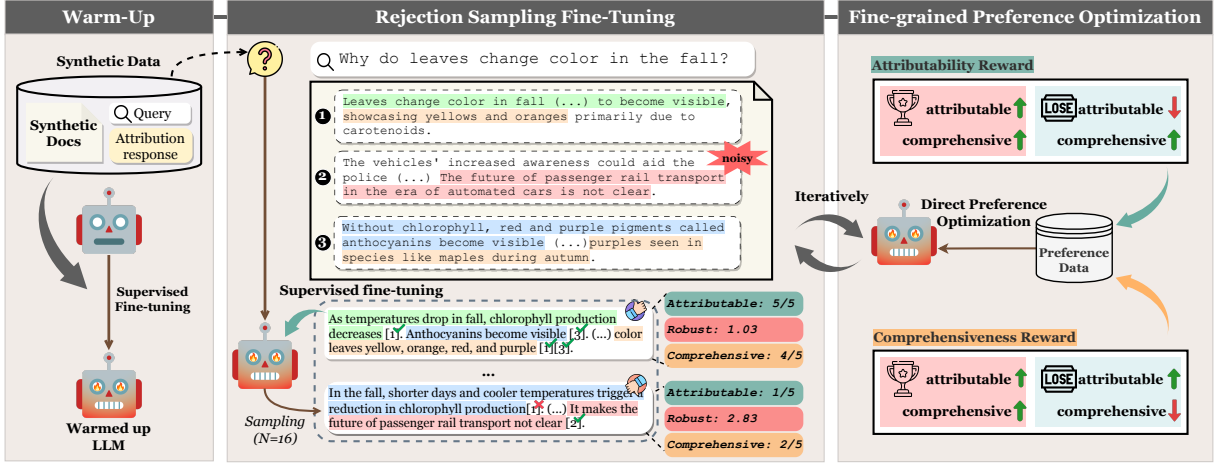


Figure 2: Overview of our self-improving framework, which consists of two stages. The model is first warmed up using synthetic data (§3.1). This provides a good starting point to enable the model to generate high-quality samples in the subsequent iterative training. Next, the model is further trained via rejection sampling fine-tuning and fine-grained preference optimization iteratively (§3.2). This iterative process bootstraps the model’s attribution capability by fully utilizing the supervision signals from its sampled generations.

the generated documents. This process ensures that each claim within the response is explicitly attributed to its source. In this way, for each query q , and documents set D , we can obtain an informative and attributable response while maintaining robustness against irrelevant documents.

Next, the model is fine-tuned for warming up with the MLE objective on the synthesized dataset, which consists of N data entries, each containing a query q_i , a document set \mathcal{D}_i , and a high-quality attributable response y_i :

$$\mathcal{L} = - \sum_{i=1}^N \log P(y_i | q_i, \mathcal{D}_i; \theta) \quad (1)$$

3.2 Self-Improving for LLM Attribution

In this stage, we propose to iteratively boost the model’s attribution capability by exploring more fine-grained supervision signals, rather than solely relying on *golden* responses in synthetic data. This involves leveraging rejection sampling for *data growing* and fine-grained preference optimization for *capability evolution*.

3.2.1 Rejection Sampling Fine-tuning

After warming up, we first sample N candidates for each query in the synthetic dataset and then score each candidate with fine-grained rewards that cover three key dimensions: *robustness*, *comprehensiveness*, and *attributability*.

Attributability serves as the indispensable condition for high-quality attributable generation. It

quantifies the extent to which a response is fully supported by the cited documents. To accurately measure attributability, we employ an off-the-shelf Natural Language Inference (NLI) model² by checking whether each statement in the response is entailed by the corresponding cited documents.

$$\text{AttrScore} = \frac{1}{S} \sum_{i=1}^S \text{Entail}(\text{Docs}, \text{statement}_i) \quad (2)$$

where S is the total number of statements in the response and Entail returns 1 if the statement i is entailed by cited documents, and 0 otherwise.

Robustness measures the degree to which a model-generated response is influenced by irrelevant contexts. Considering that we can identify relevant documents d_r within the document set D for each query q , thus we quantify robustness by calculating the probability difference of the model M to generate the response y under different contexts. The robustness score is defined as follows:

$$\text{RobustScore} = \frac{P_M(y | q \oplus d_r)}{P_M(y | q \oplus D)} \quad (3)$$

Empirically, the closer the score is to 1, the less the response is disturbed by irrelevant documents.

Comprehensiveness measures the extent to which a response captures all relevant information from the source documents. As the *golden* responses in the synthetic data are designed to

²huggingface.co/google/t5-xxl-true-nli-mixture

aggregate and reflect information across multi-documents, thus we quantify comprehensiveness by decomposing them into sub-claims and verifying whether these claims are covered by the sampled generation y . We compute the score as below:

$$\text{CompreScore} = \frac{1}{C} \sum_{i=1}^C \text{Entail}(\text{claim}_i, y) \quad (4)$$

where claim_i represents sub-claims and C is the number of *golden* sub-claims.

Subsequently, we formulate a holistic reward function (Eq. 5) considering the above dimensions. This function is employed to rank generated candidates, with the top-ranked candidate being selected for further supervised fine-tuning.

$$\text{Reward} = \mathbb{I}(\text{AttrScore}) \times \frac{\text{CompreScore}}{\text{RobustScore}} \quad (5)$$

Here, \mathbb{I} is an indicator function that returns 1 if $\text{AttrScore} = 1$, and 0 otherwise.

3.2.2 Fine-grained Preference Optimization

The common way of self-improvement focuses on updating the model with high-quality samples while discarding low-quality ones. For LLM attribution, simply supervised fine-tuning with highly attributable responses only teaches the LLM to learn surface characteristics of attribution, *e.g.*, the correct form of citation. Inspired by human cognition, learning from mistakes provides more fine-grained signals to understand the mechanisms that drive successful attribution than simply imitating correct examples. Thus, we aim to fully unlock the potential of low-quality samples by constructing fine-grained preference pairs with different optimization rewards for preference optimization.

Given the multi-objective nature of LLM attribution, our focus is specifically on *attributability* and *comprehensiveness*, utilizing corresponding rewards functions to construct preference data respectively³. Specifically, we pair samples that exhibit high attributability but low comprehensiveness with the top-ranked sample selected using a holistic reward, and vice versa. These preference pairs, each addressing different optimization objectives, are then aggregated to further train the LLM via DPO (Rafailov et al., 2023):

$$\mathcal{L}_{\text{DPO}} = -\mathbb{E}[\log \sigma(\hat{r}_\theta(x, y^+) - \hat{r}_\theta(x, y^-))]$$

³We do not optimize separately for robustness as the model already shows sufficient robustness after rejection sampling fine-tuning.

$$\hat{r}_\theta(x, y) = \beta \log \frac{\pi_\theta(y | x)}{\pi_{\text{ref}}(y | x)} \quad (6)$$

Here, reference model π_{ref} is initialized with the model after rejection sampling to minimize the distribution shift from the reference distribution.

4 Experiments

4.1 Datasets

Following previous work (Ye et al., 2023; Li et al., 2024), we conduct our experiments using two long-form question-answering datasets: ASQA (Stelmakh et al., 2022) and ELI5 (Fan et al., 2019), as well as a multi-step reasoning dataset, StrategyQA (Geva et al., 2021). Both ASQA and ELI5 feature factoid long-form answers that require synthesizing highly relevant documents in response to a user query. In StrategyQA, answers demand a combination of information-seeking and implicit reasoning. Further details on the data statistics, knowledge corpus used for retrieval, and examples for each dataset are provided in Appendix B.

4.2 Evaluation

Following previous research (Gao et al., 2023), we evaluate model-generated responses mainly on two dimensions: **Citation Quality** and **Correctness**. Our evaluation methodology combines both automated metrics and human evaluation.

Automatic Evaluation. To assess citation quality, we calculate the *citation precision*, *citation recall*, and its harmonic mean *citation F1* based on the definition in Gao et al. (2023). We use TRUE (Honovich et al., 2022), a T5-11B model fine-tuned on a collection of natural language inference (NLI) datasets to examine whether the cited documents entail the generated statement. For correctness, different datasets are measured differently. For ASQA, we report the exact match recall (**EM Rec.**) of correct short answers. For ELI5, we report the claim recall (**Claim**) by checking whether the model output entails the sub-claims generated by text-davinci-003. For StrategyQA, the format of answers begins with yes/no, we evaluate correctness by reporting the accuracy (**Acc.**). See Appendix C for more details.

Human Evaluation. We collected a total of 150 instances from the test sets of ASQA, ELI5, and StrategyQA for human evaluation, with each dataset providing 10 instances from five different systems. The evaluation is divided into two parts:

Model	ASQA				ELI5				StrategyQA			
	Correctness		Citation		Correctness		Citation		Correctness		Citation	
	EM Rec.	Rec.	Prec.	F1.	Claim	Rec.	Prec.	F1	Acc.	Rec.	Prec.	F1
<i>In-context Learning & Post-hoc</i>												
Llama-2-13B (ICL)	35.2	38.4	39.4	38.9	13.4	17.3	15.8	16.5	65.6	20.6	33.1	25.4
Llama-2-13B (PostAttr)	25.0	23.6	23.6	23.6	7.1	5.7	5.8	5.8	64.3	8.7	8.7	8.7
<i>Training-based</i>												
Distill-Llama-3-70B-Instruct	41.1	60.4	53.8	56.9	12.9	28.7	25.2	26.8	70.8	28.4	30.7	29.5
Distill-Mixtral-8x7B-Instruct	40.3	64.9	63.5	64.2	13.8	34.3	35.0	34.6	63.9	38.4	49.2	43.1
Self-RAG (Asai et al., 2023)	31.7	70.3	71.3	70.8	10.7	20.8	22.5	21.6	62.1	31.4	36.5	33.8
AGREE (Ye et al., 2023)	39.4	64.0	66.8	65.4	9.4	21.6	16.0	18.4	64.6	30.2	37.2	33.3
APO (Li et al., 2024)	40.5	72.8	69.6	71.2	13.5	26.0	24.5	25.2	61.8	40.0	39.1	39.6
FGR (Huang et al., 2024a)	38.7	73.5	74.7	74.1	9.8	53.1	55.9	54.5	64.9	29.5	42.4	34.8
START (Warming-up)	39.2	23.2	23.9	23.5	11.9	9.9	10.2	10.0	61.2	9.4	9.6	9.5
START (Iteration 1)	42.2	68.8	75.6	72.0	11.3	47.4	50.5	48.9	73.4	44.4	48.6	46.4
START (Iteration 2)	42.9	76.1	81.0	78.5	10.0	65.6	65.1	65.3	72.7	51.9	54.1	53.0
START (Iteration 3)	44.2	76.2	84.2	80.0	9.6	62.4	69.1	65.6	69.6	60.0	56.6	58.2

Table 1: Main result between our method and baselines. Experiments are evaluated on ASQA, ELI5, and StrategyQA datasets. For most baselines, we use the result of previous works (Asai et al., 2023; Ye et al., 2023; Li et al., 2024).

citation quality and overall quality (comprehensiveness and correctness). More details in Appendix D.

4.3 Baselines

We compare START with the following baselines. For more details, please refer to Appendix E.

In-context Learning (ICL). Following Gao et al. (2023), we enable the LLM to generate citations via in-context learning. For each query, we first retrieve five relevant documents and then prompt the LLM with two-shot demonstrations.

Post-hoc Attribution (PostAttr). Following Ye et al. (2023), given a query, we first instruct the LLM to generate an initial response leveraging its parametric knowledge. For each statement in the response, we use the NLI model⁴ to find the maximally supported document and cite accordingly.

Training-based Methods. Training on high-quality data serves as a strong baseline to unlock the attribution ability of LLMs. We consider the following training-based methods.

Knowledge Distillation employs the most capable LLMs, e.g., Llama-3-70B-Instruct and Mixtral-8x7B-Instruct, as teacher models to train a student model on distilled attribution data.

Self-RAG (Asai et al., 2023) first collect *data distilled from GPT-4*, then teach the LLM to retrieve on-demand while reflecting on its generation to improve both generation quality and attributions.

AGREE (Ye et al., 2023) trains the LLM to self-ground its response in retrieved documents using

⁴We use the same NLI model during citation evaluation.

automatically collected data and then leverages test-time adaptation to reinforce unverified statements.

APO (Li et al., 2024) models LLM attribution as a preference learning task, where they first supervised-fine-tuned on *human-labeled high-quality data* and then automatically collect preference data for preference optimization.

FGR (Huang et al., 2024a) first collects *attribution data distilled from ChatGPT* and then designs rewards tailored for LLM attribution to teach the LLM to generate supportive and relevant citations.

4.4 Implementation Details

For a fair comparison, all training-based baselines and START employ Llama-2-13b-base (Touvron et al., 2023). Further details on the implementation of START are presented in Appendix F.

5 Results

5.1 Main Results

We provide the main results and the performance of START across different iterations in Table 1.

START effectively improves performance. As shown in Table 1, START shows superior performance across three datasets and achieves *state-of-the-art* results in citation quality. Specifically, START shows significant improvements over both ICL and Post-hoc approaches, highlighting the benefits of supervised signals in unlocking the attribution ability of LLMs. Notably, compared with methods that rely on distilling from more advanced LLMs or training on human-annotated data, START achieves performance improvement of at

Model	ASQA				ELI5				StrategyQA			
	Correctness		Citation		Correctness		Citation		Correctness		Citation	
	EM Rec.	Rec.	Prec.	F1.	Claim	Rec.	Prec.	F1	Acc.	Rec.	Prec.	F1
START (Iteration 1)	42.2	68.8	75.6	72.0	11.3	47.4	50.5	48.9	73.4	44.4	48.6	46.4
w/o. warm-up	35.7	36.3	32.7	34.4	12.1	15.2	13.7	14.4	65.9	18.0	17.2	17.6
w/o. preference	40.6	42.2	47.2	44.6	12.9	16.5	17.4	16.9	63.7	21.5	24.6	22.9
START (Iteration 2)	42.9	76.1	81.0	78.5	10.0	65.6	65.1	65.3	72.7	51.9	54.1	53.0
w/o. warm-up	33.5	57.4	52.1	54.6	10.0	26.7	23.0	24.7	69.0	32.4	33.2	32.8
w/o. preference	39.8	50.8	53.6	52.2	12.5	22.5	23.3	22.9	65.7	27.2	30.4	28.7
START (Iteration 3)	44.2	76.2	84.2	80.0	9.6	62.4	69.1	65.6	69.6	60.0	56.6	58.2
w/o. warm-up	28.6	67.3	58.2	62.4	6.4	46.8	38.4	42.2	70.4	44.9	39.2	41.9
w/o. preference	40.7	55.7	58.3	57.0	11.9	25.3	26.2	25.7	67.8	31.3	33.5	32.4

Table 2: Ablation study results across three datasets over three iterations. We compare START with two variants: one that does not utilize synthetic data for initial warming-up (w/o warm-up) and another lacking fine-grained preference optimization for self-improvement (w/o preference).

Model	Iteration 1	Iteration 2	Iteration 3
START	42.5%	90.2%	95.9%
w/o. warm-up	3.24%	41.2%	83.8%

Table 3: The pass rate comparison between START and START (w/o. warm-up) across different iterations during the rejection sampling stage.

least 8.0%, 20.4%, and 47.0% in citation quality for ASQA, ELI5, and StrategyQA respectively. Regarding correctness, START also achieves gains of at least 9.1% and 7.2% on both ASQA and StrategyQA, despite a slight decrease on ELI5.

START successfully achieves self-improvement.

We compare the performance of START from iteration 0 to 3 in Table 1, and the results demonstrate consistent improvements across iterations. Initially, at iteration 0 (after warm-up), thanks to the synthetic training data, the model shows decent performance after warm-up. By iteration 1, START exhibits remarkable effectiveness in improving its performance by leveraging its own generated samples (e.g., 23.5 \rightarrow 72.0 on ASQA, 10.0 \rightarrow 48.9 on ELI5, 9.5 \rightarrow 46.4 on StrategyQA). Subsequent iterations continue this trend of incremental improvement, reaching a convergence point at iteration 3.

5.2 Ablation Study and Analysis

We conduct comprehensive ablation studies and analyses to understand how each component in START contributes to the significant improvement.

Effect of synthetic data warming-up. To demonstrate the importance of utilizing synthetic data for initial warm-up in START, we conduct a comparative ablation study employing Llama-2-13b for self-improvement, omitting the initial warm-up stage. Table 2 shows the ablation

results (w/o. warm-up) across three iterations. We observe that omitting the initial warm-up stage can lead to a significant performance drop in the first iteration. Additionally, as the iteration increases, the performance of the model without warm-up shows only modest improvements and remains substantially inferior to the model that underwent warm-up. Moreover, we also calculate the pass rate of sampled response in each iteration as shown in Table 3. The findings indicate that the model with warm-up exhibits a higher pass rate in the first iteration, which allows the model to utilize more supervised signals for self-improvement. These results suggest that warming up effectively facilitates the bootstrapping of supervised data, thus preventing early model stagnation. It’s worth noting that while the warm-up strategy effectively enriches the model with supervision signals at an early stage, it does not lead to noticeable improvements in citation quality, as shown in Table 1. We hypothesize that this limitation stems from the inherent difficulty LLMs face in synthesizing information from multiple sources to generate comprehensive and attributable responses solely through direct supervised fine-tuning.

Effect of fine-grained preference optimization.

To further understand the significance of fine-grained preference optimization, we compare an ablation of START that solely relies on high-quality samples for iteratively supervised fine-tuning, discarding low-quality samples for fine-grained preference optimization. As shown in Table 2, there is a significant decline in performance when fine-grained preference optimization is removed. This highlights the effectiveness of START in fully unlocking the potential of low-quality samples to en-




Figure 3: The impact of supervision signals from different stages (synthetic data v.s. self-improvement) on attribution performance across ASQA, ELI5, and StrategyQA. The blue line represents the model that undergoes only supervised fine-tuning use synthetic data at iteration 0. The red line represents the model that first trains for two epochs with synthetic data at iteration 0, followed by one iteration of self-improvement.




Figure 4: Ablation study on the effect of synthetic data size on attribution and correctness performance. We sample 1k, 3k, and 5k user queries for data synthesis. hance attribution performance.

Effect of synthetic data size. We investigate the effect of varying synthetic data sizes on the performance of START. Figure 4 demonstrates their effect on citation quality and correctness after three iterations of self-improving. Specifically, we sample 1k, 3k, and 5k unlabeled queries to generate synthetic training data accordingly, which provides different levels of supervision signals. As shown in Figure 4, even with 1k synthetic data points, START demonstrates comparable performance. Moreover, as the training size increases, START achieves notable improvement in citation quality and exhibits stability in correctness.

Supervision signals from synthetic data v.s. iterative self-improvement. We further investigate the differential impact of supervision signals derived from data synthesis versus those from the iterative self-improvement stage. We utilize synthetic training data to train the model for multiple epochs, extending up to 10 epochs, and compare its performance to that of a model that undergoes only the first iteration of self-improvement. As depicted in Figure 3, training with synthetic data during the initial iteration yields minimal performance gains. The attribution performance climbs slowly

	Attribution			Overall Quality	
	Full	Partial	No	Corr.	Comp.
ChatGPT (ICL)	68.5%	22.1%	9.4%	3.6	4.4
Distill-Llama-3-70B-Instruct	54.6%	32.4%	13.0%	2.9	3.2
Self-RAG (Asai et al., 2023)	45.7%	27.5%	26.8%	2.4	2.1
FGR (Huang et al., 2024a)	58.4%	28.7%	12.9%	2.5	2.8
START (Ours)	76.2%	18.3%	5.5%	<u>3.5</u>	4.6

Table 4: Human evaluation results on attribution, correctness (**Corr.**), and comprehensiveness (**Comp.**). **Bold** numbers indicate the best performance, while “ ” indicates the second-best performance.

as training epochs increase and fails to surpass the performance of the model after just one iteration of self-improvement. This observation reveals the importance of the supervision signals provided by the model itself during self-improvement.

6 Human Evaluation

Human evaluation results, detailed in Table 4, indicate that START generates significantly more attributable responses compared to all baselines, even surpassing ChatGPT⁵. Specifically, 76.2% of the statements generated by START are fully supported by the cited documents, which outperforms ChatGPT by 11.24%. Additionally, 18.3% of the statements are partially supported, with only 5.5% unsupported. In terms of factuality, START outperforms all training-based baselines, slightly inferior to ChatGPT. Moreover, START achieves the highest score in comprehensiveness, demonstrating its exceptional ability to generate responses that extensively cover information from multiple sources. Overall, these findings are in line with the automatic evaluation results in Table 1.

7 Conclusion

We propose START, a self-improvement framework to push the frontier of LLM attribution. We iden-

⁵We utilize gpt-3.5-turbo-0125 version.

tify two key limitations for LLM attribution self-improvement. To address these, START first leverages self-constructed synthetic data for warming up, aiming to prevent models from early stagnation due to insufficient supervision signals. To explore more fine-grained supervision signals, START constructs fine-grained preference supervision signals from low-quality samples for preference optimization. Both automatic and human evaluations demonstrate significant improvement in attribution without relying on human annotations and more advanced LLMs.

Limitations

Despite significant performance improvements, our work presents several limitations worth noting. **Firstly**, while our data synthesis process provides a good starting point for the model to self-improve and demonstrate some generalization on existing benchmarks, it may not cover all scenarios encountered in user information-seeking. This limitation raises concerns regarding the generalizability of synthetic data in a more complex information-seeking environment. **Secondly**, the iterative training pipeline of our self-improving framework is time-consuming, presenting a significant trade-off between performance and training duration. **Thirdly**, although our self-improving framework does not rely on human annotations and more advanced LLMs, it still necessitates the integration of off-the-shelf NLI models to guarantee the quality of attribution in the generated samples. The performance of the NLI model significantly impacts the quality of our outputs to a certain extent. To move towards a fully self-improving framework that does not rely on external judgment, future research could investigate the use of intrinsic attribution signals derived directly from the LLM itself.

Acknowledgements

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, the International Cooperation Project of PCL, PCL2022D01 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. 2023. [Self-rag: Learning to retrieve, generate, and critique through self-reflection](#). *CoRR*, abs/2310.11511.
- Bernd Bohnet, Vinh Q. Tran, Pat Verga, Roei Aharoni, Daniel Andor, Livio Baldini Soares, Jacob Eisenstein, Kuzman Ganchev, Jonathan Herzig, Kai Hui, Tom Kwiatkowski, Ji Ma, Jianmo Ni, Tal Schuster, William W. Cohen, Michael Collins, Dipanjan Das, Donald Metzler, Slav Petrov, and Kellie Webster. 2022. [Attributed question answering: Evaluation and modeling for attributed large language models](#). *CoRR*, abs/2212.08037.
- 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.
- Tianyu Gao, Howard Yen, Jiatong Yu, and Danqi Chen. 2023. [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.
- Mor Geva, Daniel Khashabi, Elad Segal, Tushar Khot, Dan Roth, and Jonathan Berant. 2021. [Did aristotle use a laptop? A question answering benchmark with implicit reasoning strategies](#). *Trans. Assoc. Comput. Linguistics*, 9:346–361.
- Çağlar Gülçehre, Tom Le Paine, Srivatsan Srinivasan, Ksenia Konyushkova, Lotte Weerts, Abhishek Sharma, Aditya Siddhant, Alex Ahern, Miaosen Wang, Chenjie Gu, Wolfgang Macherey, Arnaud Doucet, Orhan Firat, and Nando de Freitas. 2023. [Reinforced self-training \(rest\) for language modeling](#). *CoRR*, abs/2308.08998.
- Or Honovich, Roei Aharoni, Jonathan Herzig, Hagai Taitelbaum, Doron Kukliansy, Vered Cohen, Thomas Scialom, Idan Szpektor, Avinatan Hassidim, and Yossi Matias. 2022. [TRUE: re-evaluating factual consistency evaluation](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL 2022, Seattle, WA, United States, July 10-15, 2022*, pages 3905–3920. Association for Computational Linguistics.
- Arian Hosseini, Xingdi Yuan, Nikolay Malkin, Aaron C. Courville, Alessandro Sordani, and Rishabh Agarwal. 2024. [V-star: Training verifiers for self-taught reasoners](#). *CoRR*, abs/2402.06457.
- Chengyu Huang, Zeqiu Wu, Yushi Hu, and Wenya Wang. 2024a. [Training language models to generate](#)

- text with citations via fine-grained rewards. *CoRR*, abs/2402.04315.
- Lei Huang, Xiaocheng Feng, Weitao Ma, Yuxuan Gu, Weihong Zhong, Xiachong Feng, Weijiang Yu, Weihua Peng, Duyu Tang, Dandan Tu, and Bing Qin. 2024b. [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, 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.
- 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*.
- Dongfang Li, Zetian Sun, Baotian Hu, Zhenyu Liu, Xinshuo Hu, Xuebo Liu, and Min Zhang. 2024. [Improving attributed text generation of large language models via preference learning](#). *CoRR*, abs/2403.18381.
- Dongfang Li, Zetian Sun, Xinshuo Hu, Zhenyu Liu, Ziyang Chen, Baotian Hu, Aiguo Wu, and Min Zhang. 2023. [A survey of large language models attribution](#). *CoRR*, abs/2311.03731.
- Nelson F. Liu, Tianyi Zhang, and Percy Liang. 2023. [Evaluating verifiability in generative search engines](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023, Singapore, December 6-10, 2023*, pages 7001–7025. Association for Computational Linguistics.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Chaitanya Malaviya, Subin Lee, Sihao Chen, Elizabeth Sieber, Mark Yatskar, and Dan Roth. 2023. [Expertqa: Expert-curated questions and attributed answers](#). *CoRR*, abs/2309.07852.
- Sewon Min, Kalpesh Krishna, Xinxi Lyu, Mike Lewis, Wen-tau Yih, Pang Wei Koh, Mohit Iyyer, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2023. [Factscore: Fine-grained atomic evaluation of factual precision in long form text generation](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023*, pages 12076–12100. Association for Computational Linguistics.
- Jianmo Ni, Chen Qu, Jing Lu, Zhuyun Dai, Gustavo Hernández Ábrego, Ji Ma, Vincent Y. Zhao, Yi Luan, Keith B. Hall, Ming-Wei Chang, and Yinfei Yang. 2022. [Large dual encoders are generalizable retrievers](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022*, pages 9844–9855. Association for Computational Linguistics.
- OpenAI. 2023. [GPT-4 technical report](#). *CoRR*, abs/2303.08774.
- Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Dmytro Okhonko, Samuel Broscheit, Gautier Izacard, Patrick S. H. Lewis, Barlas Oguz, Edouard Grave, Wen-tau Yih, and Sebastian Riedel. 2021. [The web is your oyster - knowledge-intensive NLP against a very large web corpus](#). *CoRR*, abs/2112.09924.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D. Manning, Stefano Ermon, and Chelsea Finn. 2023. [Direct preference optimization: Your language model is secretly a reward model](#). In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*.
- Jeff Rasley, Samyam Rajbhandari, Olatunji Ruwase, and Yuxiong He. 2020. [Deepspeed: System optimizations enable training deep learning models with over 100 billion parameters](#). In *KDD '20: The 26th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Virtual Event, CA, USA, August 23-27, 2020*, pages 3505–3506. ACM.
- Ivan Stelmakh, Yi Luan, Bhuwan Dhingra, and Ming-Wei Chang. 2022. [ASQA: factoid questions meet long-form answers](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022*, pages 8273–8288. Association for Computational Linguistics.
- Ye Tian, Baolin Peng, Linfeng Song, Lifeng Jin, Dian Yu, Haitao Mi, and Dong Yu. 2024. [Toward self-improvement of llms via imagination, searching, and criticizing](#). *CoRR*, abs/2404.12253.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton-Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten,

Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, 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.

Xi Ye, Ruoxi Sun, Sercan Ö. Arik, and Tomas Pfister. 2023. [Effective large language model adaptation for improved grounding](#). *CoRR*, abs/2311.09533.

Asaf Yehudai, Boaz Carmeli, Yosi Mass, Ofir Arviv, Nathaniel Mills, Assaf Toledo, Eyal Shnarch, and Leshem Choshen. 2024. [Genie: Achieving human parity in content-grounded datasets generation](#). *CoRR*, abs/2401.14367.

Weizhe Yuan, Richard Yuanzhe Pang, Kyunghyun Cho, Sainbayar Sukhbaatar, Jing Xu, and Jason Weston. 2024. [Self-rewarding language models](#). *CoRR*, abs/2401.10020.

Zheng Yuan, Hongyi Yuan, Chengpeng Li, Guanting Dong, Chuanqi Tan, and Chang Zhou. 2023. [Scaling relationship on learning mathematical reasoning with large language models](#). *CoRR*, abs/2308.01825.

Eric Zelikman, Yuhuai Wu, Jesse Mu, and Noah D. Goodman. 2022. [Star: Bootstrapping reasoning with reasoning](#). In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*.

Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, Yifan Du, Chen Yang, Yushuo Chen, Zhipeng Chen, Jinhao Jiang, Ruiyang Ren, Yifan Li, Xinyu Tang, Zikang Liu, Peiyu Liu, Jian-Yun Nie, and Ji-Rong Wen. 2023. [A survey of large language models](#). *CoRR*, abs/2303.18223.

A Data Synthesis

A.1 Data Sources

The queries employed for data synthesis are sourced from the Wish-QA (Yehudai et al., 2024), which provides high-quality grounded data suitable for content-grounded generation tasks such as long-form question-answering and summarization. Specifically, we utilize the ELI5 subset of the WishQA, noted for its high lexical diversity, comprising a total of 8,413 queries. Notably, we randomly sample 5,000 user queries for our data synthesis, resulting in the creation of 5,000 synthetic data points.

A.2 Prompts for Data Synthesis

We detail the prompts employed in the synthetic data generation stage, covering response generation, claim decomposition, and document generation, shown in Figure 5.

A.3 Implementation Details

In our work, we use Llama-2-13b-base for data synthesis, as our goal is to realize self-improving for the attribution ability of LLMs, the models used in the data synthesis stage and the subsequent main experiment need to be consistent without introducing additional more powerful models. To enhance the LLM’s ability to accurately follow instructions at each step, we utilize in-context learning, incorporating two demonstrations for response generation, claim decomposition, and document generation.

A.4 Quality of Synthetic Data

We focus on evaluating the attributability of the final response. Specifically, we employ an off-the-shelf Natural Language Inference (NLI) model, TRUE (Honovich et al., 2022), to verify whether each statement in the response is fully supported by the cited documents and to check for the presence of any irrelevant citations. The results indicate that the synthetic data are of significantly high quality: 92.3% of the statements are fully supported by the cited documents, and 94.1% are free from irrelevant citations.

B Details of evaluation datasets

Our evaluation utilizes the ASQA, ELI5, and StrategyQA datasets. For both ASQA and StrategyQA, Wikipedia serves as the external knowledge base, specifically employing the Wikipedia snapshot from 2018-12-20. For the ELI5 dataset, the external knowledge source is Sphere (Piktus et al., 2021). Regarding the retrievers, we use the dense retriever GTR (Ni et al., 2022) for Wikipedia and the sparse retriever BM25 for Sphere. Detailed statistics for these datasets are presented in Table 5. In line with previous research by Gao et al. (2023), we use the same evaluation datasets for ASQA and ELI5. Regarding StrategyQA, we adopt the settings of Ye et al. (2023), utilizing a randomly split subset of 490 test instances for evaluation. To further clarify, we provide an example from each dataset in Table 6.

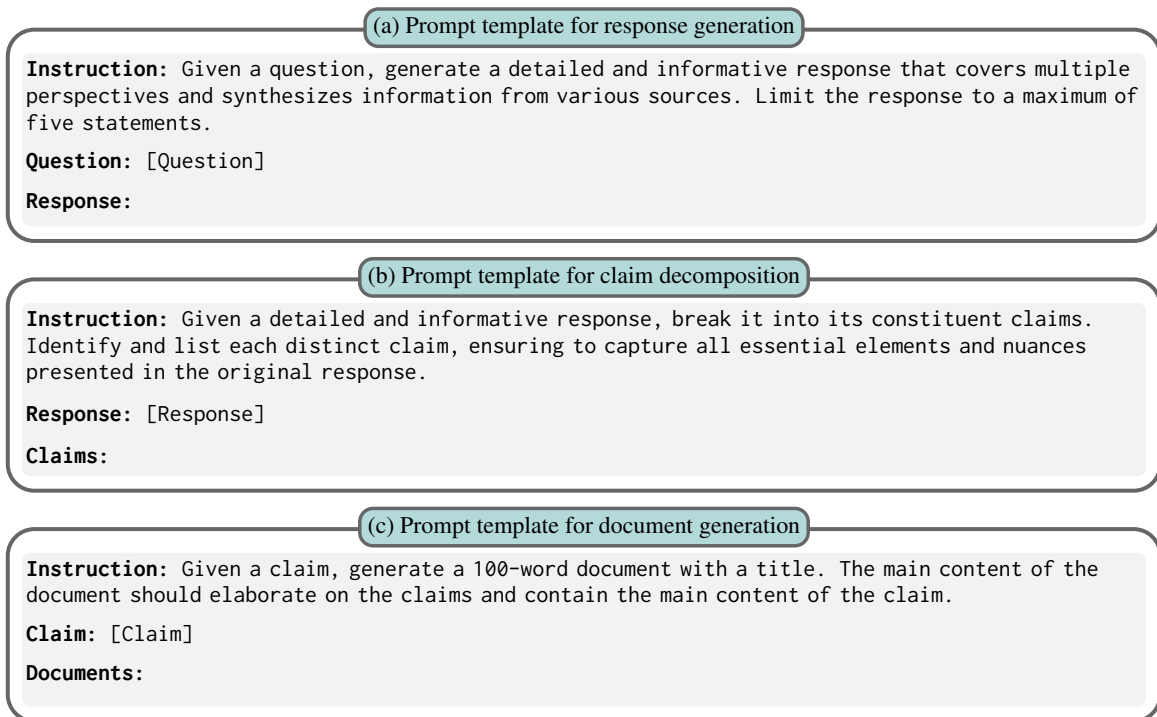


Figure 5: Illustration of the prompting design for the data synthesis pipeline.

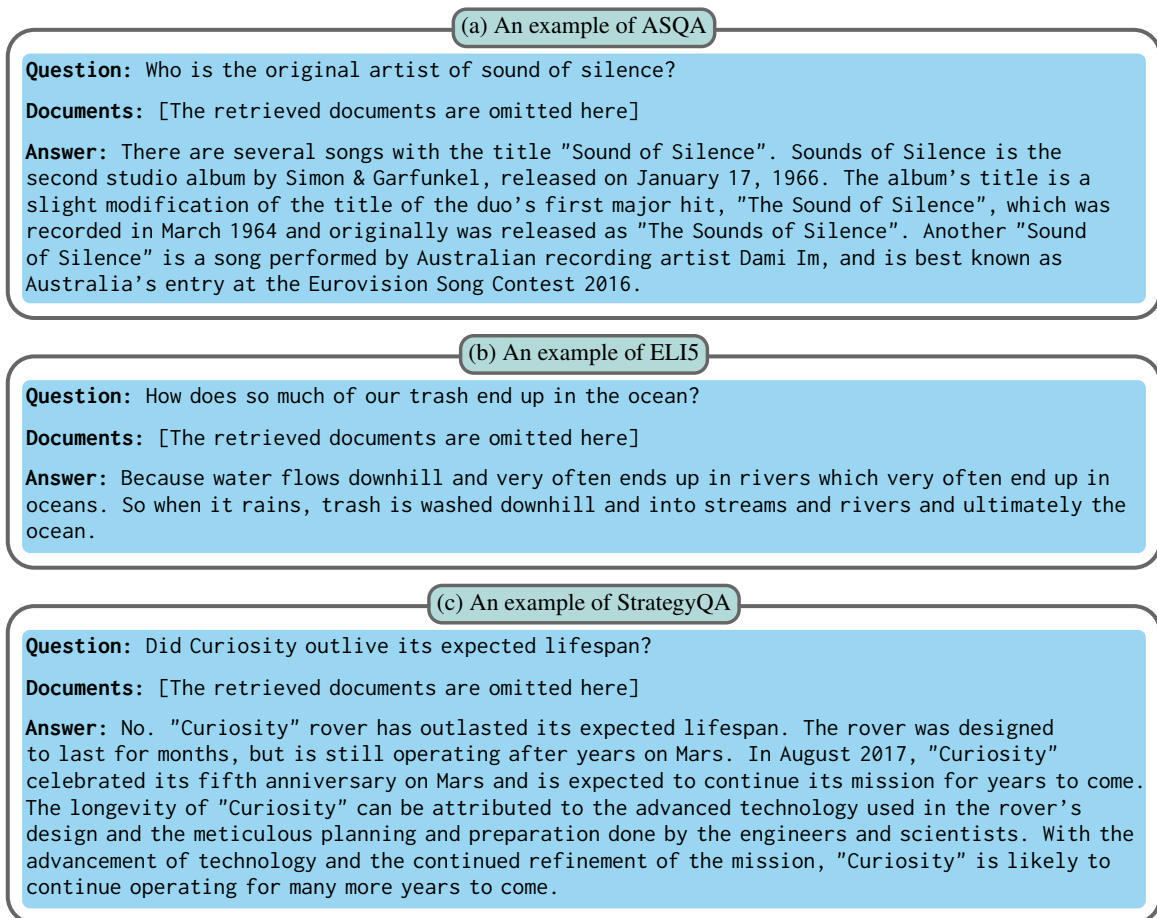


Figure 6: Examples of the ASQA, ELI5, and StrategyQA datasets.

(a) Prompt template of ASQA and ELI5

Instruction: Write an accurate, engaging, and concise answer for the given question using only the provided search results (some of which might be irrelevant) and cite them properly. Use an unbiased and journalistic tone. Always cite for any factual claim. When citing several search results, use [1][2][3]. Cite at least one document and at most three documents in each sentence. If multiple documents support the sentence, only cite a minimum sufficient subset of the documents.

Question: [Question]

Documents: [Documents]

(c) Prompt template of StrategyQA

Instruction: Answer “yes” or “no” first. Then, write a clear and concise answer that combines reasoning with relevant search results and cite the sources properly, even if some might be irrelevant.

Question: [Question]

Documents: [Documents]

Figure 7: Illustration of the prompting design of evaluation datasets.

Dataset	Source	# Examples
ASQA (Stelmakh et al., 2022)	Wiki	948
ELI5 (Fan et al., 2019)	Sphere	1000
StrategyQA (Geva et al., 2021)	Wiki	490

Table 5: Statistics of datasets used for evaluation.

C Automatic Evaluation Details

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

Citation Quality. Citation Quality is a critical evaluation dimension in attributed text generation, assessing whether the answer is fully supported by the cited documents and that no irrelevant documents are cited. Following Liu et al. (2023) and Gao et al. (2023), the evaluation of citation quality is typically divided into two parts: **Citation Recall** and **Citation Precision**.

Citation Recall evaluates whether all generated statements are fully supported by the cited documents. Specifically, for each statement $s_i \in \mathcal{S}$, its citation recall is scored as 1 if there is at least one valid citation ($\mathcal{C}_i \neq \emptyset$) and the concatenation of cited documents $\text{concat}(\mathcal{C}_i)$ fully support the statement ($\phi(\text{concat}(\mathcal{C}_i), s_i) = 1$), where $\phi(\text{premise}, \text{hypothesis})$ is an NLI model that outputs 1 if the premise entails the hypothesis. The final citation recall is calculated by averaging over all statements in \mathcal{S} .

Citation Precision assesses whether any citations in the response are irrelevant. A citation $c_{i,j}$ is

determined as “irrelevant” if (a) $c_{i,j}$ alone cannot support statement s_i and (b) removing $c_{i,j}$ does not affect the rest of the citations to support s_i .

Citation F1 is a metric that combines citation precision and citation recall by calculating their harmonic mean. In our work, we utilize this metric to evaluate the overall citation quality of the response, where a higher *Citation F1* score indicates a more accurately and comprehensively attributed response.

$$F_1 = 2 \cdot \frac{\text{citation precision} \cdot \text{citation recall}}{\text{citation precision} + \text{citation recall}}, \quad (7)$$

Correctness. Correctness is crucial in long-form QA tasks. Given the ambiguous nature of the ASQA dataset, where each question requires multiple short answers to cover different aspects, we follow Stelmakh et al. (2022) and calculate the recall of correct short answers using exact match.

As for the ELI5 dataset, evaluating the correctness of long-form answers is challenging. Thus, the ALCE benchmark employs InstructGPT (text-davinci-003) to generate three “sub-claims” based on the human-annotated answers. To assess correctness, we use a T5-11B model⁶ that has been fine-tuned on a collection of NLI datasets to check whether the model-generated outputs entail these sub-claims.

⁶https://huggingface.co/google/t5_xx1_true_nli_mixture

D Human Evaluation Details

Considering the open-ended nature of long-form QA tasks, automatic evaluation of correctness may not cover all possible answers. Furthermore, the evaluation of citation quality is constrained by the capabilities of the off-the-shelf NLI model, which may not adequately detect cases of *partial support*. Therefore, we conduct a human evaluation to assess the attribution quality and correctness of START. We recruited two annotators, holding at least a bachelor’s degree to participate in our study.

To evaluate citation quality, annotators are asked to verify whether each statement in the responses is fully supported, partially supported, or not supported by the cited documents and identify error types if the statement is not fully supported.

Next, we evaluate the overall quality of the responses, focusing on comprehensiveness and correctness. Annotators are asked to rate both comprehensiveness and correctness using a 5-point Likert scale, capturing different levels of content coverage and factuality.

E Baselines

Knowledge Distillation: We employ supervised fine-tuning to teach Llama-2-13B to generate responses with citations, utilizing training data distilled from the most advanced LLMs. Specifically, the queries and documents are sourced from our synthetic dataset and the attributed responses are generated by Llama-3-70B-Instruct / Mixtral-8x7B-Instruct.

Self-RAG (Asai et al., 2023): The method involves training the LLM to generate text with reflection tokens, which are categorized into retrieval and critique tokens to indicate the need for retrieval and the attributability of its generation, respectively. Specifically, it first collects over 145,619 supervised data by prompting GPT-4 with specific instructions to generate responses with reflection tokens for knowledge-intensive queries. These data are then used to train the LLM to generate responses with self-reflection via supervised fine-tuning.

AGREE (Ye et al., 2023): The method involves training the LLM to generate grounded claims with citations and to identify unverified claims. Specifically, it first collects 4,500 attribution data via post-hoc attribution with the help of an NLI model.

These data are then used to train the model to generate grounded responses with citations and also clearly state the unsupported statements. An iterative retrieval process is employed to search for additional information for the unsupported statements via a test-time adaptation (TTA) strategy.

APO (Li et al., 2024): This method models the attributed text generation task as a preference learning task. Specifically, the model is first trained using 6,330 human-labeled high-quality attribution data for supervised fine-tuning to learn the basic ability of attribution. It then leverages automatically constructed preference data for preference learning, where a positive response is generated from relevant documents accompanied by a positive prompt, while a negative response is generated using irrelevant documents or a negative prompt.

FGR (Huang et al., 2024a): The method first collects 3,000 in-domain user queries along with retrieved documents and then leverages ChatGPT to generate high-quality attributed responses. These data then serve as training data to teach the model the basic ability of citation generation via supervised fine-tuning. Subsequently, the method designs reward models to teach the model to generate well-supported and accurate responses via fine-grained reinforcement learning.

To ensure a fair comparison, we employ the same base model (Llama-2-13b-base) for evaluating all baselines. For Self-RAG, AGREE, and APO, we directly utilize their published experimental results. In the case of FGR, which does not provide Llama-2-13b-base results, we reproduce the experiments using the official code and the same settings provided by the authors.

F Implement Details

In all experiments, training is conducted using eight A100-80GB GPUs, leveraging Deepspeed stage 3 (Rasley et al., 2020) for multi-GPU distributed training, with training precision Bfloat16 enabled.

During the initial warm-up stage, we employ the AdamW (Loshchilov and Hutter, 2019) optimizer with a warm-up ratio of 0.03. The total batch size is set at 64, and the learning rate is maintained at $2e-5$. The maximum input sequence length is configured to 2048 tokens. The model is trained with only 20% of the synthetic dataset for two epochs in this stage. This strategy is designed to prevent the model from overfitting to the synthetic data during

the warm-up stage, enabling it to generate more diverse samples in the subsequent rejection sampling fine-tuning stage. In the self-improving stage, we conduct rejection-sampling fine-tuning for three epochs at each iteration, maintaining the same training settings as those used during the warming-up stage. To get the highest quality responses during rejection sampling, we set the threshold for attributability reward at 1.0, ensuring that every statement in the response is fully supported by the cited documents. For comprehensive, we set the threshold to 0.8, which means that at least 80% of the statements need to be cited. Subsequently, during the fine-grained preference optimization, the model is further trained for one additional epoch using a learning rate of $1e-5$.

During the evaluation, we utilize the vLLM framework (Kwon et al., 2023) for efficient inference. Without special instructions, the sampling parameters are specifically configured with a temperature of 1.0 and a top-p setting of 0.95. We present detailed prompts used during the evaluation process in Figure 7.