

Training Language Models to Generate Text with Citations via Fine-grained Rewards

Chengyu Huang

National University of Singapore
e0376956@nus.edu.sg

Yushi Hu

University of Washington
yushihu@uw.edu

Zequ Wu

University of Washington
zequwu1@uw.edu

Wenya Wang

Nanyang Technological University
wangwy@ntu.edu.sg

Abstract

While recent Large Language Models (LLMs) have proven useful in answering user queries, they are prone to hallucination, and their responses often lack credibility due to missing references to reliable sources. An intuitive solution to these issues would be to include in-text citations referring to external documents as evidence. While previous works have directly prompted LLMs to generate in-text citations, their performances are far from satisfactory, especially when it comes to smaller LLMs. In this work, we propose an effective training framework using *fine-grained rewards* to teach LLMs to generate highly supportive and relevant citations, while ensuring the correctness of their responses. We also conduct a systematic analysis of applying these fine-grained rewards to common LLM training strategies, demonstrating its advantage over conventional practices. We conduct extensive experiments on Question Answering (QA) datasets taken from the ALCE benchmark and validate the model’s generalizability using EXPERTQA. On LLaMA-2-7B, the incorporation of fine-grained rewards achieves the best performance among the baselines, even surpassing that of GPT-3.5-turbo.¹

1 Introduction

Despite the recent success of Large Language Models (LLMs), their outputs often contain incorrect or unverifiable claims (Ji et al., 2022). This reduces the usefulness of their generated content, and more importantly, hurts the user trust in LLMs. Previous work (Guu et al., 2020; Lewis et al., 2020; Izacard et al., 2022) attempts to address this issue with retrieval-augmented generation (i.e., RAG), which makes the LLM condition its generation on information retrieved from an external knowledge source for the task input. However, RAG cannot fully solve the issue because there is no guarantee

¹Our code and data is available at <https://github.com/HCY123902/atg-w-fg-rw>.

User question

What is the population of Greenville county South Carolina?

Retrieved corpus

Document [1]
Title: ...
...estimated population of 68,219 as of 2017...was 400,492...

Document [2]
Title: ...
... As of the census of 2010, the re were 58,409 people...

Document [3]
Title: ...
six cities...the population was 451,225...was 506,837...

Document [4]
Title: ...
Greenville...the population was 451,225...was 506,837...

Document [5]
Title: ...
...an estimated population of 68,219 as of 2017...it is the...

ChatGPT response (In-Context Learning)

As of 2017, the estimated population of Greenville, South Carolina is 68,219, making it the sixth-largest city in the state. [Missing citation]

The population of the surrounding area was 400,492 as of 2010, making it the third-largest urban area in South Carolina as well as the fastest growing. [Missing citation]

Greenville County, where the city is located, had a population of 506,837 as of 2017, making it the most populous county in the state. [1][3][4] [Irrelevant citation]

Gold answer: 451,225 [ChatGPT does NOT capture the gold answer]

Figure 1: An example of ChatGPT performing the task of attributable generation. The model takes a question, retrieved passages, and the task instruction (omitted due to space limit) as the input, and generates a response with in-text citations. The response has 3 sentences, 2 of which do not have supporting citations. The third one has an irrelevant citation [1]. Moreover, ChatGPT does not capture the correct answer (451,225) mentioned in passages [3] and [4].

that the retrieved knowledge is relevant nor that the LM-generated output is always consistent with the knowledge.

Recent studies (Nakano et al., 2021; Gao et al., 2023b) propose to have LLMs generate answers with in-text citations that provide evidence for the generated content, often referred to as *attributable* text generation. The cited references provide easy access for end users to perform fact-checking on the model generations. Most existing work (Malaviya et al., 2023; Sun et al., 2023; Li et al., 2023b; Kamaloo et al., 2023) simply prompts LLMs (e.g., ChatGPT) to include citations in their generation. Gao et al. (2023b) conduct systematic evaluations to show that the model performance can

often be far from satisfactory with such prompting approaches. In this paper, we aim to study how we can *train* an LM to perform better at attributable generation.

One challenge in achieving this goal comes from the expensive cost of collecting supervised training data that contains gold generations interleaved with citations. Another challenge stems from the fine-grained task objective. Instead of a holistic goal of generating the correct response to the task input, attributable generation usually targets on two critical aspects: citation quality and the information correctness of the model response. The citation quality can be further assessed by two dimensions. The first dimension checks if a generated sentence is related to and can be supported by the corresponding cited reference(s). Another dimension inspects if any cited reference in the generation is redundant.

To address these challenges, we resort to two training algorithms, namely rejection sampling (RS) and reinforcement learning (RL), and propose to use fine-grained rewards as training signals to teach an LM (e.g., LLaMA 2 (Touvron et al., 2023)) to generate attributable responses. We leverage the fine-grained automatic evaluation functions from Gao et al. (2023b) to provide localized (sentence-level) and specialized (citation quality or response correctness) reward signals as fine-grained rewards. For training, we adapt from the fine-grained RL algorithm in Wu et al. (2023) and the tree-decoding mechanism in Asai et al. (2023) to perform reinforcement learning and rejection sampling, respectively, using the fine-grained rewards. To ensure a relatively good starting point before RS or RL training, we initialize the LM with distillation from ChatGPT on a handful of sampled training prompts from ALCE (Gao et al., 2023b).

When evaluating our trained LMs on the test set of ALCE and an out-of-domain dataset EXPERTQA (Malaviya et al., 2023), our main findings are as follows:

- Training an LM with fine-grained rewards for attributable generation is much more effective than a holistic reward. This conclusion holds for both RS and RL. To the best of our knowledge, this is the first work that proposes to *train* an LM with *fine-grained* rewards for attributable generation.
- Training with rejection sampling is more efficient and effective than reinforcement learning, while combining them leads to the best

model performance that significantly surpasses ChatGPT.

- We also show the generalizability of RS and RL that they can produce attributable answers on EXPERTQA, a separate dataset that requires intensive domain knowledge.

2 Problem Definition and Methods

We follow Gao et al. (2023b) to define **attributable text generation** as: Given a query x and a passage corpus \mathcal{D} , generate a textual output y that consists of N sentences s_1, s_2, \dots, s_N and citations (such as [1], [1][2], etc.) that refer to passages $\mathcal{C}_i \subseteq \mathcal{D}$ to provide evidence for each sentence s_i .

In this section, we introduce how to train an LM with fine-grained rewards to perform the task. As the corpus \mathcal{D} usually contains millions of passages, we first retrieve a small set (e.g., top-5) of relevant passages as the first step before training an LM to generate the output, conditioning on the query and retrieved passages. The details of the retriever for each dataset are in Appendix B.2. Figure 2 illustrates our training process. As supervision for in-text citations is hard to obtain, we first initialize an LM with distillation from a proprietary LM such as ChatGPT, as described in § 2.1. We then introduce how fine-grained rewards are computed and used to train the LM with RL or RS in § 2.2. Finally, we describe how to train the LM with holistic rewards for comparison purposes in § 2.3.

2.1 Distillation from ChatGPT

Gao et al. (2023b) observe that the performance of open-sourced LMs like LLaMA-2 for attributable generation is far from satisfactory. In particular, we notice that even with in-context learning, a base LLaMA-2-7B can only generate citations for a very small portion of sentences in its responses. As previous studies (Touvron et al., 2023; Wu et al., 2023) often perform RS or RL training on top of a reasonably good initial LM as a starting point, we first learn a better LM capable of generating in-text citations by distilling knowledge from a powerful proprietary model such as ChatGPT². For each question in the training set, we prompt ChatGPT with in-context demonstrations (Brown et al., 2020) to generate answers that contain citations (See Appendix I). These generated answers serve as weak supervisions to finetune our LM. We use

²We use gpt-3.5-turbo-0301

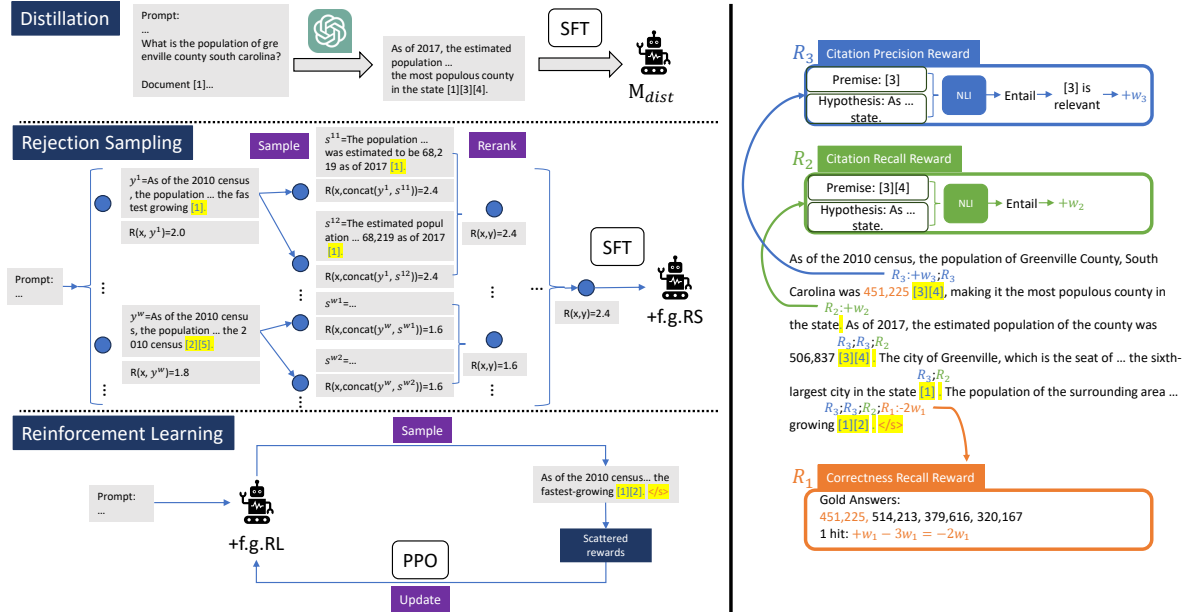


Figure 2: **Right:** The assignment of our fine-grained rewards (R_1 : Answer Correctness, R_2 : Citation Recall, R_3 : Citation Precision). These rewards are assigned to corresponding tokens in the response (citation, EOS Token, etc.; highlighted in **yellow**). **Left:** An overview of our framework. Top: Distillation from ChatGPT (§ 2.1); Middle: Rejection Sampling (§ 2.2.1); Bottom: Reinforcement Learning (§ 2.2.2).

LLaMA-2-7B as our LM and construct its input as the concatenation of the task instruction, the query x , and the passages retrieved for x . We denote the resulting model after distillation as \mathcal{M}_{dist} .

2.2 Training with Fine-grained Rewards

Attributable text generation aims to generate responses that contain correct information and high-quality citations. Its task objective is often decomposed into three aspects (Gao et al., 2023b): information correctness, citation recall, and citation precision. As a result, we use three reward functions to represent these three sub-goals respectively. Our **fine-grained rewards** are detailed below:

Correctness R_1 . Correctness of a model-generated response to a given query is usually measured as the fraction of key information pieces included in the response (Stelmakh et al., 2022). Specifically, in most existing factoid QA datasets whose questions require synthesizing multiple documents to answer (Stelmakh et al., 2022; Rubin et al., 2022; Sun et al., 2022), each question comes with a list of short phrases or important information statements that cover different aspects of the complete answer, referred to as key information list. We can either apply an exact string match (EM) or use a Natural Language Inference (NLI) model to infer whether each item (sub-claim) in the list is

covered in the model response. See more in Appendix A.1. Note that although in-text citations can be hard to obtain for training, the key information lists are often provided in existing datasets. For each item in the list, we assign a reward of $+w_1$ if it is in the response, and $-w_1$ otherwise. We then sum up the rewards for all list items to obtain the reward for the whole response.

Citation Recall R_2 . Citation Recall is the percentage of sentences in the model response that can be supported by the corresponding cited passages. We use an NLI model (Honovich et al., 2022) to determine whether each sentence is entailed by its reference passages. Specifically, the cited passage(s)³ are treated as the premise and the answer sentence is treated as the hypothesis for the NLI model. We assign $+w_2$ for each sentence that can be entailed by the premise, and $-w_2$ otherwise. As a result, R_2 is computed at the sentence level. See more in Appendix A.2.

Citation Precision R_3 . Citation Precision refers to the percentage of citations in the response that can help support their sentences (Gao et al., 2023b). The same NLI model in R_2 is used to infer whether each cited passage is necessary to entail the sentence. We assign $+w_3$ for each citation that is "relevant" (See Appendix A.3), and $-w_3$ otherwise.

³If multiple passages are cited, we concatenate them.

Therefore, we apply R_3 at the citation level.

2.2.1 Rejection Sampling (RS)

RS is a decoding strategy that can be used to generate labels for model fine-tuning. Specifically, the model is prompted with each training example input to generate multiple output samples. Then a *holistic* reward function is used to rank these samples and the top-1 output is used as the label for model fine-tuning.

In previous studies (Touvron et al., 2023; Nakano et al., 2021), RS has proven to be effective in improving model performance. In this work, we apply a novel *fine-grained* RS training approach to improve the LM’s performance on attributable generation after the distillation step in Section 2.1.

We adopt a sentence-level beam search decoding algorithm introduced in Asai et al. (2023) to rank model outputs with our fine-grained reward functions. At each sentence step, the LM generates K different continuation candidates for each sequence in the beam. Then we obtain the top- B generation sequences based on a combined reward, where B is the beam size. We compute the combined reward R as the sum of R_1 , R_2 , R_3 for each generated sequence $y^w = s_{1,\dots,i}$ at a sentence step. More specifically, R is computed as:

$$R(x, y^w) = \sum_{u=1}^3 \sum_{k=1}^{l_u} R_u^k(x, y^w) \quad (1)$$

where x is the query that is inputted together with y^w to compute the rewards of y^w , l_u is the number of corresponding segments in y^w for R_u (e.g., $l_1 = 1$ since R_1 treats the entire generation as a single sequence, and $l_2 = i$ since R_2 is computed at the sentence level). R_u^k denotes the reward score R_u for the k^{th} segment in y^w .

The whole process stops until all sequences in the beam reaches the end of generation or the maximum number of steps H is reached. Then the highest ranked sequence is returned for further LM fine-tuning as in a standard RS process.

2.2.2 Reinforcement Learning (RL)

We also experiment with LM training using RL after the distillation step. Following Wu et al. (2023), we optimize the policy LM using proximal policy optimization (PPO) (Schulman et al., 2017) with fine-grained rewards. We adapt the token-level reward assignment in Wu et al. (2023) to our own fine-grained reward functions by assigning i) the sequence-level R_1 to the EOS token of each generated response; ii) the R_2 score for each sentence

to the last token of that sentence; iii) the R_3 score for each bracketed citation to its end token (i.e., the right bracket). If a token is assigned by multiple reward functions, we take their sum to be the final reward for that token.

Formally, given a (prompt, response) pair (x, y) , we denote the token time step of the EOS token of y as T_1^1 , the time step of the end token of each sentence s_i as T_i^2 , and the time step of the end token for each citation c_j as T_j^3 . We use θ and θ_{init} to represent the current and initial policy respectively and refer to P as the probability distribution over the next token given the generation prefix g_t at time step t . The reward for each token a_t generated at time step t in y is:

$$r_t = \sum_{u=1}^3 \sum_{k=1}^{l_u} (\mathbb{1}(t=T_k^u) R_u^k(x, y)) - \beta \log \frac{P_\theta(a_t|g_t)}{P_{\theta_{init}}(a_t|g_t)} \quad (2)$$

where l_1 , l_2 , and l_3 denote the number of sequences (which is 1), sentences, and citations in y and β is the coefficient for the KL penalty commonly used in PPO. R_u^k denotes R_u for the k^{th} segment in y .

With the above reward assignment, we follow the standard PPO algorithm to optimize both the policy and value models.

2.2.3 Combining RS and RL

Stacking RL on top of RS is shown to be effective (Touvron et al., 2023) for LLM training. Therefore, we also experiment with this option. Starting from \mathcal{M}_{dist} , we apply RS to fine-tune the model, and then apply RL to further train it.

2.3 Training with Holistic Rewards

For comparison purposes, we also analyze model performances when trained with *holistic* rewards. In this setting, each generated response is assigned a single overall reward. Similarly, we experiment with both RS and RL for model fine-tuning. We use R in Eq. 1 when $i = N$ as our holistic reward.

RS We follow the standard RS pipeline to have the LM generate multiple outputs and rank them by the holistic reward R . The best generated output is then used for model fine-tuning. In contrast, our *fine-grained* RS samples and ranks the generated sequences at each sentence step, which allows finer and better control over the sampling process.

RL We assign the holistic reward to the last token of each generated sequence and all other tokens receive a reward of 0. In contrast, *fine-grained* RL assigns denser rewards to tokens in a sequence.

Similar to Eq. 2, we add the KL penalty term to each token’s reward. This RL process is designed in the same way as in most recent work in RL from human preference feedback (Menick et al., 2022; Nakano et al., 2021; Touvron et al., 2023).

3 Experiment Setup

3.1 Datasets

We fine-tune our LM with three datasets in the ALCE (Gao et al., 2023b) benchmark designed for the task of attributable generation: (1) ASQA (Stelmakh et al., 2022), a long-form QA dataset containing ambiguous questions that can have different interpretations and corresponding answers under different contexts; (2) QAMPARI (Rubin et al., 2022), a factoid QA dataset where the answer to each question is a list of entries; (3) ELI5 (Fan et al., 2019), a long-form QA dataset with mostly Why/How/What type of questions. Most questions in these datasets require the model to synthesize knowledge from multiple passages to construct a complete answer. Each dataset has its own passage corpus, and we use GTR (Ni et al., 2022) or BM25 to retrieve the top-5 passages from the corresponding corpus for each question, following Gao et al. (2023b). See more in Appendix B.2.

We train the LM on all of the three QA datasets. We pre-process the training examples to address the imbalance of the three datasets’ training sets. This results in roughly 3k/1k train/dev examples evenly distributed over the three datasets. Furthermore, as ALCE only provides development set examples, we obtain around 3k test set examples with roughly 1k instances from each of the ASQA, QAMPARI, and ELI5 test sets.

To verify the generalizability of our framework, we test our model using a dataset unseen during training, EXPERTQA (Malaviya et al., 2023), which is another long-form QA dataset whose questions requires knowledge in specific domains (e.g., psychology) to answer. See more in Appendix B.1.

3.2 Evaluation Metrics

We use the evaluation metrics in the original ALCE benchmark⁴. **Correctness Recall** measures the percentage of gold answers (key information pieces) captured by the response, and its implementations vary among datasets (**EM Rec;Rec.-5;Claim Rec**).

⁴While ALCE also uses MAUVE (Pillutla et al., 2021) to evaluate fluency, Gao et al. (2023b) show it is not always reliable. We do not include it in the main analysis.

Citation Recall measures the percentage of sentences in the response that are entailed by their cited passages. **Citation Precision** measures the percentage of citations that can help support the output sentences. On QAMPARI, there is an additional **Correctness Precision** metric, which measures the percentage of predicted answer items that are correct, following (Gao et al., 2023b). See Appendix B.3 and Figure 3 for more details.

On EXPERTQA, we use the originally designed metrics as well, including **AutoAIS** (Gao et al., 2023a), which is similar to Citation Recall and measures the percentage of sentences supported by the citations, and **FACTscore** (Min et al., 2023), which measures the percentage of generated claims that are factual.

3.3 Training Details

We fine-tune LLaMA-2-7B (Touvron et al., 2023) in the main experiments. For RS with fine-grained rewards, we use a beam width (B) of 8, and we sample 2 continuations (K) for each sequence in the beam in a single search step. We set a sentence-level maximum search depth (H) of 5 on ASQA and ELI5, and 10 on QAMPARI, since QAMPARI questions have more key information pieces to answer on average. For holistic RS, we set the number of sampling sequences to 16 to make a fair comparison. In terms of reward weights, we set each of w_1 , w_2 , and w_3 to 0.2 throughout the RS and RL process for LLaMA-2-7B. Empirically, we find that this weight combination serves our goal of **optimizing citation quality, while ensuring correctness**. See more details in Appendix C.

3.4 Baselines

As mentioned in § 2.1, we initialize the language model (LLaMA-2-7B) with \mathcal{M}_{dist} using the training examples generated from ChatGPT. After this step, we leverage fine-grained rewards to fine-tune the model with rejection sampling (**f.g. RS**) or reinforcement learning (**f.g. RL**). We also experiment with combining RS and RL for model fine-tuning (**f.g. RS+RL**). As discussed in § 2.3, we also report results of using holistic rewards for RS (**h. RS**), RL (**h. RL**), and RS+RL (**h. RS+RL**).

We compare the above *fine-tuned* LMs with results from in-context learning (**ICL**) (Brown et al., 2020). Specifically, we prompt the LM with 2 in-context demonstrations where each contains the question, 5 retrieved passages, and an answer with

citations (See Appendix I). We report ICL results with ChatGPT (GPT-3.5-turbo) or LLaMA-2-7B.

4 Results and Analysis

We first compare all the systems (§ 4.1), and then conduct ablation study on the training effect of each reward type (§ 4.2). We also analyze the effect of retrieval on the final model performance (§ 4.3), and the citation errors that we observe (§ 4.4). Lastly, we test our models’ generalization capability on EXPERTQA that is unseen during training (§ 4.5).

4.1 Main Results

The main results are summarized in Table 1.

Training with fine-grained rewards greatly boosts performance. Adding f.g. RL training to the initial model (\mathcal{M}_{dist}), we observe the performance increase on every metric and dataset, except for Citation Recall on QAMPARI and Claim Recall on ELI5. Compared to RL, training with RS generally leads to more performance improvement across the three datasets. However, we observe that combining RS and RL leads to the best performance.

In addition, there is a much larger performance boost on ASQA and ELI5 with fine-tuning than on QAMPARI. We suspect the reason behind this to be the fact that gold responses of QAMPARI are lists of entities, while the other two datasets have natural language responses, making training biased towards the latter format. Appendix D shows that in the *separate* training setting (i.e., training LMs on each *individual* dataset), fine-tuning with either RS, RL or RS+RL can give a significant performance boost for QAMPARI.

Across different f.g. models, we observe a consistent increase in citation recall and precision, while correctness at least remains on the same level if not increase, which matches our objective.

Fine-grained rewards help smaller LLMs surpass ChatGPT. Although the performance gap remains between LLaMA-2 and ChatGPT after distillation, LLaMA-2 trained with fine-grained rewards can close the gap or even outperform ChatGPT. Training with a combined RS and RL leads to an average of 4.0%, 0.9%, 10.6% improvements over ChatGPT among all evaluation metrics on ASQA, QAMPARI, and ELI5 respectively.

Fine-grained rewards are better than holistic rewards. We compare fine-grained rewards with

holistic ones for model training. Table 1 shows the outstanding effectiveness of fine-grained rewards in all training setups (RS, RL, RS+RL). We observe higher performance gains with fine-grained rewards in almost all metrics and all datasets. The only exception is that the citation recall on ELI5 is similar for h. RL and f.g. RL. The performance gap between using fine-grained and holistic rewards is wider for RS compared to RL.

4.2 Ablation of Reward Models

We conduct our ablation study to understand how learning to optimize one single task objective (answer correctness or attribution quality) affects the model performance. Since we show in § 4.1 that RS generally gives better results than RL, we conduct this ablation analysis on f.g. RS. We train the LM on each *individual* QA dataset by using either the rewards associated with answer correctness (R_1) or citation quality (R_2 and R_3). See results in Table 2.

To show the effect of only optimizing towards answer correctness, we remove R_2 and R_3 during RS. As indicated by the number of tokens (#Tokens) metric in Table 2, if the model is only given the correctness reward, it will increase its response length in an attempt to capture more gold answers. Citation quality is ignored and decreases quickly. As a result, the model achieves the highest correctness recall compared to other systems.

We also attempt to focus only on citation quality and remove R_1 during RS. The models produce better results on citation quality on the long-form QA datasets, while their correctness recall drops.

4.3 Retrieval Analysis

Since the retrieved passages may not cover all gold answers (Gao et al., 2023b), we calibrate the model’s correctness metrics, so that we can directly inspect the model’s ability to capture answers in the given passages. Specifically, we ignore gold answers that are not recalled by the top 5 retrieved passages when computing correctness metrics. We observe that our fine-grained methods (f.g. RL; f.g. RS; f.g.(RS+RL)) can capture around 70% of whatever is already captured by the passages on ASQA and QAMPARI, and this number becomes more than 90% on ELI5. In other words, **our trained LLMs are effective in locating answers within retrieved passages**, although there is still space to improve the models’ correctness recall by capturing the remaining answers in the passages. On the

	ASQA			QAMPARI				ELI5		
	Correctness		Citation	Correctness		Citation		Correctness		Citation
	EM Rec	Rec	Prec	Rec.-5	Prec	Rec	Prec	Claim Rec	Rec	Prec
ICL ChatGPT	39.96	74.72	70.97	18.34	18.57	17.65	20.60	13.47	50.94	47.58
ICL LLaMA-2-7B	34.15	14.12	15.26	8.24	9.23	7.51	7.65	7.83	14.44	8.92
\mathcal{M}_{dist}	35.56	74.80	67.99	17.26	16.18	18.69	18.94	12.03	49.69	45.71
<i>RL/RS only</i>										
+h.RL	34.33	75.77	70.12	17.30	16.44	16.39	18.92	11.52	51.77	49.32
+f.g.RL	35.99	76.30	72.38	18.39	18.81	17.82	19.07	11.60	51.29	51.09
+h.RS	37.96	74.86	68.48	14.62	15.21	16.71	17.98	11.60	54.10	48.95
+f.g.RS	40.07	76.71	74.35	16.14	18.95	18.56	19.50	11.67	58.75	55.03
<i>RS + RL</i>										
+h.(RS+RL)	37.33	74.86	69.37	15.02	15.67	16.82	18.09	11.21	55.62	50.58
+f.g.(RS+RL)	40.05	77.83	76.33	16.65	19.54	19.50	20.03	11.54	60.86	60.23

Table 1: Main result on the test set of three QA datasets.

	ASQA				QAMPARI					ELI5			
	Correctness		Citation		Correctness		Citation			Correctness		Citation	
	EM Rec	Rec	Prec	#Tokens	Rec.-5	Prec	Rec	Prec	#Tokens	Claim Rec	Rec	Prec	#Tokens
ICL ChatGPT	39.96	74.72	70.97	37.51	18.34	18.57	17.65	20.60	13.70	13.47	50.94	47.58	86.06
+f.g.RS _{sep}	40.24	77.65	74.96	51.34	17.48	20.67	20.62	21.65	11.24	11.87	61.27	56.45	83.01
w/o. R_2+R_3	41.29	49.51	67.54	56.34	23.12	18.13	15.99	16.73	15.29	13.60	41.06	43.38	88.38
w/o. R_1	39.79	79.42	75.69	55.89	12.62	17.63	20.31	21.45	10.91	11.17	62.92	58.51	84.26

Table 2: Ablation study results. Our models here are trained on each *individual* dataset.

other hand, to further improve answer correctness, it is important to improve the retrieval performance. It might also be possible that a correctness recall gets high when the retrieved passages don't contain gold answers, as is in the case of ELI5 (the Passage Recall is only 9.13, while ChatGPT's Correctness Recall is 13.47), since LLMs have parametric knowledge embedded during pretraining. However, such generated answers are not attributable. We include more results in Appendix E.

4.4 Citation Error Analysis

We randomly select and inspect responses generated by f.g. RS+RL for 30 test examples (10 from ASQA/QAMPARI/ELI5 each), and observe the following citation errors. We include examples for each of the error categories in Appendix G.

Mixing up of passage ids (5.26%). This refers to the case where the claim is *unsupported* as a result of the model mixing up the passage ids. The model sometimes *correctly* extracts and paraphrases the content from one passage, but attributes them to another passage.

Redundant citations (31.58%). This refers to the case where passages that are *irrelevant* to the response are cited. An example case is when no retrieved passage is relevant to the question. In this case, the model may generate the response based on its own parametric knowledge, but still cite irrelevant passages. This error is common on

the long-form QA datasets⁵ ASQA and ELI5, and it directly impacts Citation Precision.

Misinterpretation of cited passages (63.16%). We notice that our LLMs's responses are more *extractive* after training. However, the model may still interpret and paraphrase the facts mentioned in the passages *incorrectly*. There are two subcategories. The first is **misinterpretation of a single passage (52.63%)**, where facts in a single passage are distorted. This error is common on QAMPARI, where the questions usually involve multi-hop relations between multiple entities and the model can easily make mistakes in the reasoning chain. The second is **incorrect synthesis of multiple passages (10.53%)**, where the model makes logical errors (e.g., hallucinating relations between entities) when connecting information from different passages.

4.5 Analysis on Generalizability

To demonstrate the generalizability of our model, we evaluate it on a separate dataset, EXPERTQA (Malaviya et al., 2023). We adopt the metrics from its paper, which measure citation recall and factuality, and the results are shown in Table 3.

The capability of attributable generation is generalizable. Similar to the previous three datasets, the model trained with fine-grained rewards gives

⁵We explicitly ask the model to attach only 1 citation to each item in the response on QAMPARI following Gao et al. (2023b). As a result, this error is less common on QAMPARI.

	EXPERTQA		
	AutoAIS	FActscore	#Sents
ICL ChatGPT	56.98	85.83	8145
ICL LLaMA-2-7B	19.63	82.77	10053
\mathcal{M}_{dist}	51.33	83.46	7293
<i>RL/RS only</i>			
+h.RL	53.64	83.82	8139
+f.g.RL	56.15	83.89	7322
+h.RS	57.88	83.47	8436
+f.g.RS	63.49	83.85	7450
<i>RS+RL</i>			
+h.(RS+RL)	59.42	83.71	8012
+f.g.(RS+RL)	66.12	83.78	6256

Table 3: Result on EXPERTQA. #Sents: Total number of sentences in the responses to the 2,169 samples.

the highest AutoAIS scores (Gao et al., 2023a), signifying the highest percentage of sentences that are supported by citations. This indicates the robustness and transferability of our model’s attribution capability when trained with fine-grained rewards.

Trained LLMs have high factual precision.

While we do not use factuality as a reward model, the FActscore (Min et al., 2023) metric shows that more than 80% of sentences generated by our methods are factual, which is close to that of ChatGPT.

5 Related Work

Retrieval-Augmented Language Models (RALM). There exists prior work on RALMs for tasks such as question-answering and fact-checking. Aside from the works on kNN-LM (Khandelwal et al., 2020; Zhong et al., 2022), most of them involve (1) document selection and (2) document reading. Many studies attempt to improve document selection by training a stronger retriever or applying additional filtering mechanisms. Methods such as REALM (Guu et al., 2020), RAG (Lewis et al., 2020), and Atlas (Izacard et al., 2022) train the retrievers and LMs jointly in an end-to-end setting. Lin et al. (2022) finetune the LMs and retrievers in two separate stages that complement each other. Others made efforts to train an extra NLI model (Yoran et al., 2023) or only the LLMs themselves (Luo et al., 2023; Asai et al., 2023) to filter the retrieved documents. Other lines of work explore effective retrieve-on-the-fly methods (Parisi et al., 2022; Schick et al., 2023; Jiang et al., 2023; Ram et al., 2023; Gao et al., 2023b) and post-hoc retrieval for answer attribution or revision (Gao et al., 2023b; Malaviya et al., 2023; He et al., 2022).

Most previous work does not include explicit citations or quotes in their models’ responses. However, generating citations can both help reduce hallucination and increase verifiability (Huang and Chang, 2023; Li et al., 2023a). Among those prior methods that do provide citations, they either require intensive human annotations such as demonstrations and preferences (Nakano et al., 2021), only consider simple questions requiring a single source (Menick et al., 2022; Bohnet et al., 2023), do not achieve satisfactory results on smaller LLMs compared with ChatGPT (Gao et al., 2023b; Malaviya et al., 2023; Asai et al., 2023), or use sophisticated mechanism that involves repeated prompting of LLMs and revision of their responses during inference (Sun et al., 2023; Ye et al., 2023). The performance of existing commercial systems like BingChat is neither satisfactory in terms of their citation recall (Malaviya et al., 2023; Liu et al., 2023).

Training with Reward Models. Recent work uses reward models to train LLMs for various tasks. In the context of RALMs, Menick et al. (2022) trains a preference-based reward model and then uses Reinforcement Learning from Human Preference Feedback (RLHF) to train its model to quote knowledge sources. Nakano et al. (2021) train similar reward models, but they also use them for best-of-n sampling during inference on top of RLHF. Zhou et al. (2023) use tree search for QA tasks and utilize an LM value function to generate scores for each solution trajectory. A common characteristic of these reward models is that they only measure the overall quality of a response. A more natural and informative reward model for citations is one that is fine-grained since citations can be scattered throughout a piece of text. Inspired by Wu et al. (2023), we divide a holistic reward into more specific ones using both reward density and category.

6 Conclusion and Future Directions

In this work, we systematically analyze common LLM training methods for attributable text generation. We show that training with fine-grained rewards gives the strongest performance boost and it also enables small LLMs to outperform ChatGPT.

We propose a number of directions to explore. Firstly, our retrieval analysis shows room to further improve LLMs’ correctness recall on QA datasets, and more work can be done to enhance LLMs’ reading comprehension and synthesis ability given the

retrieved passages. Secondly, according to [Touvron et al. \(2023\)](#), iteratively applying imitation learning and reinforcement learning for multiple rounds can continuously increase the model’s performance. While this is not the main focus of our work, it will be intriguing to see how far a smaller LLM can go in our task setting using this approach.

Limitations

There are several limitations of our work. Firstly, our correctness reward model on ELI5 may not reflect the true correctness recall fully. The sub-claims inferred by text-davinci-003 are shown to be mostly correct, where more than 90 percent of them are faithful to the original long-form answers from which they are derived ([Gao et al., 2023b](#)). However, these sub-claims can be incomplete, meaning that they may not cover every aspect of their sources. This can cause inconsistencies between the training objectives, evaluation metrics, and actual correctness. Inspection of the sub-claims completeness requires intensive human annotations and is deferred to future work. Secondly, our method still requires an initial distillation step with ChatGPT. This can possibly hinder accessibility when larger, more capable LLMs are not available. A possible resolution to this is to iteratively use In-Context Learning and beam search sampling to bootstrap responses of high quality, and then apply behavioral cloning and reinforcement learning to update the model weights, which can be a future direction to explore.

Ethics Statement

The original datasets from which we curate our own train/dev/test splits are open-sourced and publicly available, and thus the risk of sensitive information leakage is kept at a minimum. However, we do acknowledge that our proposed method, especially the corpus that we use for passage retrieval, may introduce certain noises and biases into LLMs’ responses. For instance, the corpus that we use on ELI5 and EXPERTQA, Sphere, contains many noisy and biased contents due to its sheer size, and can potentially harm the LLMs’ generation. While this is not the main focus of our work, investigating noise-free and socially responsible retrieval mechanisms for attributable text generation is an important future direction.

References

- Akari Asai, Zeqiu Wu, Yizhong Wang, Avirup Sil, and Hannaneh Hajishirzi. 2023. [Self-rag: Learning to retrieve, generate, and critique through self-reflection](#). *arXiv preprint arXiv:2310.11511*, 2023.
- Bernd Bohnet, Vinh Q. Tran, Pat Verga, Roe Aharoni, Daniel Andor, Livio Baldini Soares, Massimiliano Ciaramita, Jacob Eisenstein, Kuzman Ganchev, Jonathan Herzig, Kai Hui, Tom Kwiatkowski, Ji Ma, Jianmo Ni, Lierni Sestorain Saralegui, Tal Schuster, William W. Cohen, Michael Collins, Dipanjan Das, Donald Metzler, Slav Petrov, and Kellie Webster. 2023. [Attributed question answering: Evaluation and modeling for attributed large language models](#). *arXiv preprint arXiv:2212.08037*.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Daniel Ziegler, Aditya Ramesh, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*.
- Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. 2023. [Qlora: Efficient finetuning of quantized llms](#). *arXiv preprint arXiv:2305.14314*.
- Angela Fan, Yacine Jernite, Ethan Perez, David Grangier, Jason Weston, and Michael Auli. 2019. [Eli5: Long form question answering](#). In *Association for Computational Linguistics (ACL)*, pages 3558—3567.
- Luyu Gao, Zhuyun Dai, Panupong Pasupat, Anthony Chen, Arun Tejasvi Chaganty, Yicheng Fan, Vincent Zhao, Ni Lao, Hongrae Lee, Da-Cheng Juan, and Kelvin Guu. 2023a. [Rarr: Researching and revising what language models say, using language models](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, page 16477–16508.
- Tianyu Gao, Howard Yen, Jiatong Yu, and Danqi Chen. 2023b. [Enabling large language models to generate text with citations](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 6465—6488.
- Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Mingwei Chang. 2020. [Realm: Retrieval augmented language model pre-training](#). In *Proceedings of International Conference on Machine Learning, 2020*, pages 33–40.
- Hangfeng He, Hongming Zhang, and Dan Roth. 2022. [Rethinking with retrieval: Faithful large language model inference](#). *arXiv preprint arXiv:2301.00303*, 2022.

- 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 *North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, pages 3905–3920.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. [Lora: Low-rank adaptation of large language models](#). *arXiv preprint arXiv:2106.09685*, 2021.
- Jie Huang and Kevin Chen-Chuan Chang. 2023. [Citation: A key to building responsible and accountable large language models](#). *arXiv preprint arXiv:2307.02185*.
- Gautier Izacard, Patrick Lewis, Maria Lomeli, Lucas Hosseini, Fabio Petroni, Timo Schick, Jane Dwivedi-Yu, Armand Joulin, Sebastian Riedel, and Edouard Grave. 2022. [Atlas: Few-shot learning with retrieval augmented language models](#). *arXiv preprint arXiv:2208.03299*, 2022.
- Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. 2022. [Survey of hallucination in natural language generation](#). *ACM Computing Surveys*, 55(12):1–38.
- Zhengbao Jiang, Frank F Xu, Luyu Gao, Zhiqing Sun, Qian Liu, Jane Dwivedi-Yu, and Yiming Yang. 2023. [Active retrieval augmented generation](#). *arXiv preprint arXiv:2305.06983*, 2023.
- Ehsan Kamalloo, Aref Jafari, Xinyu Zhang, Nandan Thakur, and Jimmy Lin. 2023. [HAGRID: A human-llm collaborative dataset for generative information-seeking with attribution](#). *arXiv:2307.16883*.
- Urvashi Khandelwal, Omer Levy, Dan Jurafsky, Luke Zettlemoyer, and Mike Lewis. 2020. [Generalization through memorization: Nearest neighbor language models](#). In *International Conference on Learning Representations*.
- 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*.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Kuttler, Mike Lewis, Wen tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. [Retrieval-augmented generation for knowledge-intensive nlp tasks](#). In *Advances in Neural Information Processing Systems*.
- Dongfang Li, Zetian Sun, Xinshuo Hu, Zhenyu Liu, Ziyang Chen, Baotian Hu, Aiguo Wu, and Min Zhang. 2023a. [A survey of large language models attribution](#). *arXiv:2311.03731*.
- Xinze Li, Yixin Cao, Liangming Pan, Yubo Ma, and Aixin Sun. 2023b. [Towards verifiable generation: A benchmark for knowledge-aware language model attribution](#). *arXiv preprint arXiv:2310.05634*.
- Xi Victoria Lin, Xilun Chen, Mingda Chen, Weijia Shi, Maria Lomeli, Rich James, Pedro Rodriguez, Jacob Kahn, Gergely Szilvassy, Mike Lewis, Luke Zettlemoyer, and Scott Yih. 2022. [Ra-dit: Retrieval-augmented dual instruction tuning](#). *arXiv preprint arXiv:2310.01352*, 2022.
- Nelson F. Liu, Tianyi Zhang, and Percy Liang. 2023. [Evaluating verifiability in generative search engines](#). *arXiv preprint arXiv:2304.09848*.
- Hongyin Luo, Yung-Sung Chuang, Yuan Gong, Tianhua Zhang, Yoon Kim, Xixin Wu, Danny Fox, Helen Meng, and James Glass. 2023. [Sail: Search-augmented instruction learning](#). *arXiv preprint arXiv:2305.15225*, 2023.
- Chaitanya Malaviya, Subin Lee, Sihao Chen, Elizabeth Sieber, Mark Yatskar, and Dan Roth. 2023. [Expertqa: Expert-curated questions and attributed answers](#). *arXiv preprint arXiv:2309.07852*, 2023.
- Jacob Menick, Maja Trebacz, Vladimir Mikulik, John Aslanides, Francis Song, Martin Chadwick, Mia Glaese, Susannah Young, Lucy Campbell-Gillingam, Geoffrey Irving, and Nat McAleese. 2022. [Teaching language models to support answers with verified quotes](#). *arXiv preprint arXiv:2203.11147*, 2022.
- Sewon Min, Kalpesh Krishna, Xixi 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](#). *arXiv preprint arXiv:2305.14251v1*.
- Reiichiro Nakano, Jacob Hilton, Suchir Balaji, Jeff Wu, Long Ouyang, Christina Kim, Christopher Hesse, Shantanu Jain, Vineet Kosaraju, William Saunders, Xu Jiang, Karl Cobbe, Tyna Eloundou, Gretchen Krueger, Kevin Button, Matthew Knight, Benjamin Chess, and John Schulman. 2021. [Webgpt: Browser-assisted question answering with human feedback](#). *arXiv preprint arXiv:2112.09332*, 2021.
- Jianmo Ni, Chen Qu, Jing Lu, Zhuyun Dai, Gustavo Hernandez Abrego, Ji Ma, Vincent Zhao, Yi Luan, Keith 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*, pages 9844–9855.
- Aaron Parisi, Yao Zhao, and Noah Fiedel. 2022. [Talm: Tool augmented language models](#). *arXiv preprint arXiv:2205.12255*, 2022.

- Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Dmytro Okhonko, Samuel Broscheit, Gautier Izacard, Patrick Lewis, Barlas Oğuz, Edouard Grave, Wen tau Yih, and et al. 2021. [The web is your oyster - knowledge-intensive nlp against a very large web corpus.](#) *arXiv preprint arXiv:2112.09924*, 2021.
- Krishna Pillutla, Swabha Swayamdipta, Rowan Zellers, John Thickstun, Sean Welleck, Yejin Choi, and Zaid Harchaoui. 2021. [Mauve: Measuring the gap between neural text and human text using divergence frontiers.](#) In *Advances in Neural Information Processing Systems*.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer.](#) *The Journal of Machine Learning Research (JMLR)*, 21(140).
- Ori Ram, Yoav Levine, Itay Dalmedigos, Dor Muhlgay, Amnon Shashua, Kevin Leyton-Brown, and Yoav Shoham. 2023. [In-context retrieval-augmented language models.](#) In *Transactions of the Association for Computational Linguistics*, 2023.
- Samuel Joseph Amouyal Ohad Rubin, Ori Yoran, Tomer Wolfson, Jonathan Herzig, and Jonathan Berant. 2022. [Qampari: An open-domain question answering benchmark for questions with many answers from multiple paragraphs.](#) *arXiv preprint arXiv:2205.12665*, 2022.
- Michael Santacroce, Yadong Lu, Han Yu, Yuanzhi Li, and Yelong Shen. 2023. [Efficient rlhf: Reducing the memory usage of ppo.](#) *arXiv preprint arXiv:2309.00754*, 2023.
- Timo Schick, Jane Dwivedi-Yu, Roberto Dessì, Roberta Raileanu, Maria Lomeli, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. 2023. [Toolformer: Language models can teach themselves to use tools.](#) *arXiv preprint arXiv:2302.04761*, 2023.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. [Proximal policy optimization algorithms.](#) *arXiv preprint arXiv:1707.06347*, 2017.
- Ivan Stelmakh, Yi Luan, Bhuwan Dhingra, and Ming-Wei Chang. 2022. [Asqa: Factoid questions meet long-form answers.](#) *arXiv preprint arXiv:2204.06092*, 2022.
- Haitian Sun, William Cohen, and Ruslan Salakhutdinov. 2022. [Conditionalqa: A complex reading comprehension dataset with conditional answers.](#) In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics*, page 3627–3637.
- Hao Sun, Hengyi Cai, Bo Wang, Yingyan Hou, Xiaochi Wei, Shuaiqiang Wang, Yan Zhang, and Dawei Yin. 2023. [Towards verifiable text generation with evolving memory and self-reflection.](#) *arXiv preprint arXiv:2312.09075*.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, and et al. 2023. [Llama 2: Open foundation and fine-tuned chat models.](#) *arXiv preprint arXiv:2307.09288*, 2023.
- Zequi Wu, Yushi Hu, Weijia Shi, Nouha Dziri, Alane Suhr, Prithviraj Ammanabrolu, Noah A Smith, Mari Ostendorf, and Hannaneh Hajishirzi. 2023. [Fine-grained human feedback gives better rewards for language model training.](#) In *Advances in Neural Information Processing Systems*.
- Xi Ye, Ruoxi Sun, Sercan Ö. Arik, and Tomas Pfister. 2023. [Effective large language model adaptation for improved grounding.](#) *arXiv preprint arXiv:2311.09533*.
- Ori Yoran, Tomer Wolfson, Ori Ram, and Jonathan Berant. 2023. [Making retrieval-augmented language models robust to irrelevant context.](#) *arXiv preprint arXiv:2310.01558*, 2023.
- Xiang Yue, Boshi Wang, Kai Zhang, Ziru Chen, Yu Su, and Huan Sun. 2023. [Automatic evaluation of attribution by large language models.](#) *arXiv preprint arXiv:2305.06311*.
- Zexuan Zhong, Tao Lei, and Danqi Chen. 2022. [Training language models with memory augmentation.](#) In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 5657–5673.
- Andy Zhou, Kai Yan, Michal Shlapentokh-Rothman, Haohan Wang, and Yu-Xiong Wang. 2023. [Language agent tree search unifies reasoning acting and planning in language models.](#) *arXiv preprint arXiv:2310.04406*, 2023.

A Details on Fine-grained Rewards

A.1 Correctness Recall Reward

In ASQA, each answer (key information piece) corresponds to a separate interpretation of the question and is usually a single phrase. In QAMPARI, each question has a list of gold answers, and each answer is also a single phrase. For these two datasets, we check how many answers from the gold answer set (key information list) are exact substrings of the model response. In ELI5 and EXPERTQA, since only a single long-form answer is available for each question, we infer three sub-claims from the gold long-form answers with text-davinci-003, and use these claims as the key information pieces. We then use an NLI model to check how many sub-claims are entailed by the model response.

In QAMPARI, since users usually may not want the entire list, we do not penalize a missing answer if at least 5 answer entries have been captured.

As there is no clear boundary for an answer unit in long-form response, we adopt a sequence-level granularity for R_1 , meaning that the reward is computed once and assigned to a single token for the entire response sequence.

Specifically, if the response y to a question from dataset d string matches (ASQA; QAMPARI) or entails (ELI5; EXPERTQA) h items out of t gold answer items, then the reward is

$$R_1^1(x, y) = \begin{cases} w_1 h - w_1(t - h), & d \neq \text{qampari} \\ w_1 h - w_1 \max(\min(t, 5) - h, 0), & d = \text{qampari} \end{cases} \quad (3)$$

A.2 Citation Recall Reward

Given a sentence s_i , its set of cited passages C_i , and an NLI model σ that checks whether a set of cited documents entail a sentence, s_i has a Citation Recall reward of $+w_2$ (i.e., $R_2^i(x, y) = +w_2$) if and only if it is entailed by the concatenation of the cited passages, meaning that $\sigma(\text{concat}(C_i), s_i) = 1$. Otherwise, the reward is $-w_2$ (i.e., $R_2^i(x, y) = -w_2$).

On long-form QA datasets (ASQA; ELI5; EXPERTQA), this reward is computed for each sentence in the response. On QAMPARI, since the model response is a comma-separated list, we treat each comma-separated item (formatted as '{question} {item}') in the response as a sentence.

A.3 Citation Precision Reward

We use the definition in (Gao et al., 2023b). Given a sentence s_i , its set of cited passages C_i , and

an NLI model σ , a citation that refers to passage $c_j \in C_i$ has a Citation Precision reward of $+w_3$ (i.e., $R_3^j(x, y) = +w_3$) if and only if $\sigma(\text{concat}(C_i), s_i) = 1$ and either of the following conditions holds

- c_j explicitly entails s with $\sigma(c_j, s) = 1$
- c_j implicitly helps other passages entail s with $\sigma(\text{concat}(C_i - \{c_j\}), s) = 0$

Otherwise, the reward is $-w_3$ (i.e., $R_3^j(x, y) = -w_3$).

B Datasets and Metrics

B.1 Datasets

ASQA The authors of ASQA have not released its test split. Therefore, we randomly shuffle and then divide its original train split samples into two groups so that one group is used for training and the other is used for validation. We use its original dev split as our test split. This gives 3,853/500/948 samples for train/dev/test respectively.

QAMPARI We randomly sample 4,000 instances from the original train split of QAMPARI and take its entire development and test splits, giving 4,000/1,000/1,000 samples for train/dev/test usage.

ELI5 The original ELI5 dataset is much larger than ASQA. To ensure the balance between different datasets, we randomly sample 3,728/1,000/1,000 instances from the train/dev/test splits of ELI5. We also include in the train split 264 samples from WebGPT (Nakano et al., 2021), since these samples' answers contain citations and are of relatively high quality. This gives 3,992/1,000/1,000 samples in the train/dev/test splits.

The above gives around 4k/0.5k/1k samples for train/valid/test splits on ASQA and approximately 4k/1k/1k samples on QAMPARI and ELI5. They are used directly in the *separate* setting (i.e., training LMs on each *individual* dataset). In the *combined* setting (i.e., training LMs on the mixture; This is our main experiment setting), we further sample around 1k/334/1k instances from each filtered dataset to get 2,992/1,002/3,000 samples for joint training and inference.

EXPERTQA We remove 8 samples that do not have human-revised answers, which gives us the remaining 2,169 samples. The reason is that we

need annotated answers to infer sub-claims for correctness recall computation in the additional experiments in Appendix H and we want our testing set in § 4.5 to be consistent with what is used in the additional experiments. Since some groups of samples in EXPERTQA are evaluated in a closed-book setting or use a different retrieval mechanism, we manually retrieve the top 5 passages from Sphere (Piktus et al., 2021) again for each sample to ensure consistency with other datasets. We use EXPERTQA for testing only.

B.2 Retrieval

We use GTR (Ni et al., 2022) to retrieve the top-5 passages from the 2018-12-20 Wikipedia snapshot on ASQA and QAMPARI, and use BM25 to retrieve the top-5 passages from Sphere (Piktus et al., 2021) on ELI5 and EXPERTQA respectively.

B.3 Metrics

We use the metrics from ALCE (Gao et al., 2023b) to evaluate our model generations’ correctness and citation quality. Their computations are illustrated in Figure 3.

We illustrate their computations with a response y with l_2 sentences and l_3 citations to a question with t gold answers (key information pieces). See examples in Figure 3.

Correctness Recall Suppose y captures h out of t key information pieces (i.e., they have a Correctness Recall reward of $+w_1$ (See Appendix A.1)), then its Correctness Recall is $\frac{h}{t}$. The detailed implementations vary among datasets (**EM Rec**; **Rec-5**; **Claim Rec**).

More specifically, since the key information lists are available in ASQA and QAMPARI, we compute an exact string match recall (**EM Rec**) that checks how many key information pieces are an exact substring of generated responses. On top of that, following (Gao et al., 2023b), we compute **Rec-5** (i.e., $\frac{\min(h,5)}{\min(t,5)}$) on QAMPARI, so that the recall will be 100 if at least 5 correct answers are hit. On ELI5, we only have a single long-form answer for each question. Therefore, we follow (Gao et al., 2023b) to infer three sub-claims for each answer with text-davinci-003, and use them as the key information pieces. We then use an NLI model to check how many sub-claims are entailed by the model response (**Claim Rec**).

Additionally, on QAMPARI, we directly ask the model to format its response as a list of items. To

check whether the model is blindly making more guesses to increase Correctness Recall, we include a **Correctness Precision** metric on QAMPARI. It measures the percentage of predicted items that match a key information piece. Suppose there are t' comma-separated entries in the generated response, then the Correctness Precision will be $\frac{h}{t'}$.

Citation Recall Suppose l'_2 out of l_2 sentences are entailed by the cited passages (i.e., they each have a Citation Recall reward of $+w_2$ (See Appendix A.2)), then the Citation Recall is $\frac{l'_2}{l_2}$.

On QAMPARI, since the model generates a comma-separated list instead of a natural language response, we treat each comma-separated item (formatted as '{question} {item}') as a sentence.

Citation Precision Suppose l'_3 out of l_3 citations are "relevant", (i.e., they each have a Citation Precision reward of $+w_3$ (See Appendix A.3)), then the Citation Precision is $\frac{l'_3}{l_3}$.

MAUVE MAUVE (Pillutla et al., 2021) is an additional metric in ALCE that is used to evaluate fluency, which we do not include in the main analysis since it is sensitive to output length and unstable (Gao et al., 2023b). We only report it in Appendix D. It is computed on the model’s generated responses against the provided gold natural language responses. Following (Gao et al., 2023b), we remove citations in the model’s generations, append the question to both the gold and generated responses, and truncate them to the first 100 words before computing MAUVE.

C Additional Training Details

Efficiency For memory efficiency, we fine-tune LoRA (Hu et al., 2021) in the distillation and RS fine-tuning steps, and QLoRA (Dettmers et al., 2023) for RLHF. Moreover, in RLHF training, we use an architecture similar to Hydra-PPO (Santacrose et al., 2023) with a slight modification that separates the weights of the reward model from the main model. We use vLLM (Kwon et al., 2023) to accelerate inference.

Context and generation length We set the maximum source length to be 1,024 for T5-large⁶ and 1,200 for LLaMA-2-7B since its prompts contains instructions. We set the maximum generation length to 200 for both models to control the conciseness of answers.

⁶In addition to LLaMA-2-7B, we also apply our method to T5-large. Its results are in Appendix D.

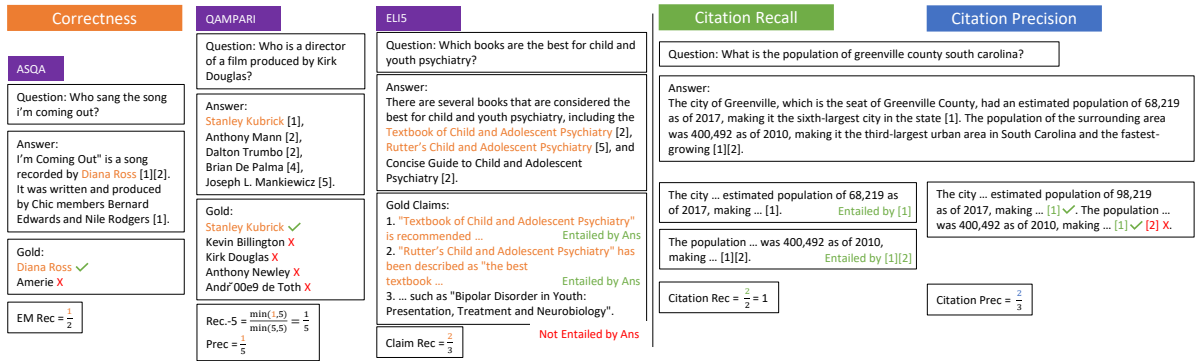


Figure 3: Left: Examples of how the Correctness metrics are computed for ASQA (EM Rec), QAMPARI (Rec-.5, Prec), and ELI5 (Claim Rec) respectively; Right: An example of how the Citation Recall and Citation Precision are computed.

Distillation We instruction-tune our model for 10 epochs with Adam Optimizer. We use a learning rate of $5e-5$ and a linear learning rate scheduler for T5-large. We use a learning rate of $2e-4$ and a cosine learning rate scheduler with a warm-up ratio of 0.03 for LLaMA-2-7B. The best checkpoints are selected using the dev set for RS fine-tuning or RL.

RS fine-tuning When fine-tuning the model, we use the same hyperparameters as in the distillation step, except that we change the number of training epochs to 3.

LoRA We set the rank to 64 and α to 16. The adapters are applied to the Q and V projection layers in the transformer unit. We use the same set of hyperparameters for every adapter.

RL Similar to Wu et al. (2023), we set $\gamma = 1$ and $\lambda = 0.95$. We use Adam Optimizer. The learning rate is set to $1e-5$ for T5-large and $3e-5$ for LLaMA-2-7B, and a linear scheduler is used for both models. Additionally, we sample 4 response sequences from the model for each training sample, with the temperature set to 0.7 and top k set to 20. On T5-large, we train the model for 80,000 episodes in total. On LLaMA-2-7B, for efficiency reasons, we set the number of episodes to 48,000. We use 100 warm-up steps across every setting. The KL coefficient β is set to 0.3. We repeat each RL experiment 3 times, each with a different seed. The results are the average of the inference results of these 3 independent runs.

NLI model For reward computation, we use TRUE⁷ (Honovich et al., 2022), a T5-11B model for NLI, to measure entailment for correctness

⁷See https://huggingface.co/google/t5_xxl_true_nli_mixture.

recall on ELI5, and citation recall and precision across our datasets. For efficiency, we replace TRUE with a smaller NLI model, namely attrscore-flan-t5-xl (Yue et al., 2023), during the RL training stage.

D Complete Main Experiments

To demonstrate the general applicability of our framework, we also apply our method to T5-large (Raffel et al., 2020). We provide the complete main experiment results in Table 4, including results on T5-large and on LLaMA-2-7B checkpoints. These new models are trained in the *separate* setting, which means they are trained on each individual dataset instead of the mixture. We also include additional baselines ICL_{LLaMA-2-13B} and Self-RAG (Asai et al., 2023). Note that RL experiments for the *separate* training settings are run only once, since we observe little variation when repeating them with different seeds.

E Retrieval Analysis

As can be seen from Table 5, around 70% of answers present in the retrieved documents can be captured by our LLMs on ASQA and QAMPARI. On ELI5, this coverage increases to more than 90%.

F Training Curves

We show the f.g.RL training curves of LLaMA-2-7B across the 3 independent runs in figure 4. While correctness recall fluctuates, there is a steady increase in citation recall and citation precision.

	ASQA				QAMPARI				ELI5			
	Correctness	Citation		Fluency	Correctness	Citation		Fluency	Correctness	Citation		Fluency
	EM Rec	Rec	Prec	MAUVE	Rec.-5	Prec	Rec	Prec	Claim Rec	Rec	Prec	MAUVE
ICL ChatGPT	39.96	74.72	70.97	76.49	18.34	18.57	17.65	20.60	13.47	50.94	47.58	43.13
ICL LLaMA-2-7B	34.15	14.12	15.26	85.31	8.24	9.23	7.51	7.65	7.83	14.44	8.92	43.53
ICL LLaMA-2-13B	35.97	21.59	25.27	88.01	12.66	12.88	8.34	8.42	13.20	6.86	4.51	41.15
Self-RAG _{7B} *	30.0	67.8	66.9	74.3	-	-	-	-	-	-	-	-
Self-RAG _{13B} *	31.7	71.3	70.3	71.6	-	-	-	-	-	-	-	-
<i>T5 (Separate Setting)</i>												
\mathcal{M}_{dist}	33.60	71.98	65.15	66.81	15.98	14.30	15.48	17.96	9.03	48.32	43.29	51.04
+f.g.RL	33.75	74.24	68.30	61.11	18.10	14.56	15.75	17.97	9.63	52.04	47.25	40.31
<i>LLaMA-2-7B (Separate Setting)</i>												
\mathcal{M}_{dist}	34.84	69.30	62.20	59.93	17.66	16.42	16.14	17.93	12.77	49.37	43.73	48.53
+h.RL	36.80	71.57	68.87	77.79	17.30	17.08	16.80	18.56	12.40	49.95	45.83	54.68
+f.g.RL	36.94	71.86	69.75	77.21	18.86	17.66	16.95	18.36	12.23	49.76	47.10	47.23
+h.RS	37.41	77.84	71.28	68.71	16.06	18.01	17.29	18.13	11.30	55.93	51.35	42.75
+f.g.RS	40.24	77.65	74.96	83.10	17.48	20.67	20.62	21.65	11.87	61.27	56.45	53.37
+h.(RS+RL)	37.72	78.47	73.19	68.65	17.58	20.74	21.25	21.61	11.30	56.67	53.69	59.79
+f.g.(RS+RL)	41.33	79.71	77.90	91.14	15.76	21.60	21.36	21.92	11.47	63.45	59.96	51.32
<i>LLaMA-2-7B (Combined Setting)</i>												
\mathcal{M}_{dist}	35.56	74.80	67.99	62.08	17.26	16.18	18.69	18.94	12.03	49.69	45.71	36.03
+h.RL	34.33	75.77	70.12	61.79	17.30	16.44	16.39	18.92	11.52	51.77	49.32	45.46
+f.g.RL	35.99	76.30	72.38	69.99	18.39	18.81	17.82	19.07	11.60	51.29	51.09	40.64
+h.RS	37.96	74.86	68.48	79.04	14.62	15.21	16.71	17.98	11.60	54.10	48.95	47.91
+f.g.RS	40.07	76.71	74.35	87.04	16.14	18.95	18.56	19.50	11.67	58.75	55.03	45.03
+h.(RS+RL)	37.33	74.86	69.37	81.91	15.02	15.67	16.82	18.09	11.21	55.62	50.58	47.76
+f.g.(RS+RL)	40.05	77.83	76.33	89.08	16.65	19.54	19.50	20.03	11.54	60.86	60.23	40.52

Table 4: Complete main result on the three QA datasets. In-Context: In-Context Learning; * indicates results that are reported by the original paper; - indicates numbers not reported by the original paper; Entries that are highlighted in yellow are those that surpass their counterparts from ChatGPT. Our results on ChatGPT slightly deviate from those reported by Gao et al. (2023b), since they use sampling when decoding the responses during inference, and their evaluation on QAMPARI and ELI5 is run on the development split, while ours is run on the test split.




Figure 4: Training curves of LLaMA-2-7B with f.g.RL in the *combined* setting, measured on the development set across 3 independent runs. The shaded region indicates the standard error across these runs.

G Citation Error Analysis

We show examples for each type of citation error in Table 6, 7, 8, and 9.

H Additional Experiment on EXPERTQA

We also evaluate EXPERTQA with our own metrics. The results are shown in Table 10. Similar to ELI5, our method outperforms ChatGPT on citation recall and precision, but still lags behind on claim recall. The reasons can be complicated, while the most probable ones are the same as those for ELI5, including (1) insufficient emphasis on correctness reward (w_1) during RS and RL and (2) incorrect sub-claims derivation by text-davinci-003 that can cause inconsistency between the re-

wards and true correctness recall of the generated responses.

I Prompts

The prompts for ASQA, QAMPARI, and ELI5 are shown in Table 11, Table 12, and Table 13 respectively. The prompt for EXPERTQA is the same as that for ELI5, except that we append an additional sentence 'If you don't know the answer, just say "I don't know" and explain why the context is insufficient to answer the question.' to the instruction, in order to explicitly allow the model to abstain from answering the question. The reason is that EXPERTQA questions usually require more intensive domain knowledge to answer. Our prompts are consistent with those used by Gao et al. (2023b). Note that the in-context demonstrations are included in

	ASQA	QAMPARI			ELI5
	P.R.: 57.65	P.R.: 20.53			P.R.: 9.13
	C.R.	C.R.-5	C.R.	C.P.	C.R.
ICL ChatGPT	73.05	68.81	67.57	16.86	93.28
ICL LLaMA-2-7B	65.24	55.16	54.59	8.99	91.15
\mathcal{M}_{dist}	67.68	69.95	68.78	15.59	92.58
<i>RL/RS only</i>					
+h.RL	66.10	69.89	68.73	15.85	92.47
+f.g.RL	68.62	71.19	69.84	18.02	92.34
+h.RS	70.87	63.90	63.03	14.74	92.68
+f.g.RS	74.28	67.27	66.16	18.41	92.73
<i>RS+RL</i>					
+h.(RS+RL)	70.19	64.67	63.79	15.25	92.41
+f.g.(RS+RL)	74.08	67.87	66.75	18.87	92.57

Table 5: Retrieval Analysis. P.R.: Passage Recall (i.e., Correctness Recall@5 of the retrieved passages); C.R.: Correctness Recall calibrated with the answers in the retrieved passages; C.R.-5: Correctness Recall that is 100 if there are at least 5 hits, and calibrated with the retrieved passages; C.P.: Correctness Precision calibrated with the retrieved passages.

Instruction: Provide a ...

Question: Who is a director of a film produced by Kirk Douglas?

Document [1](Title: Kirk Douglas): began producing films as varied as "Paths of Glory" (1957) and ...
Document [2](Title: Kirk Douglas): executive producer as well, raising the \$12 million production cost, making ... Douglas initially selected Anthony Mann to direct, but replaced him early on with Stanley Kubrick, with whom ...
When the film was released, Douglas gave full credit to its screenwriter, Dalton Trumbo, who was on the ...
Document [3] ...

Answers: Stanley Kubrick [1], Anthony Mann [2], Dalton Trumbo [2], Brian De Palma [4], Joseph L. Mankiewicz [5].

Table 6: An example from QAMPARI test set demonstrating the mixing up of passage ids error. Stanley Kubrick is only mentioned in [2], but is erroneously attributed to [1].

Instruction: Write an ...

Question: What exactly is 'spin' in quantum mechanics?

Document [1](Title: quantum mechanics - Does spin have anything to do with a rate of change? - Physics Stack Exchange): quantum mechanics - Does spin have anything to do ... what exactly does it mean, 1/2
Document [2](Title: quantum mechanics - Does spin have anything to do with a rate of change? - Physics Stack Exchange): quantum mechanics - Does spin have anything to do ... why do we relate spin with angular momentum operator in quantum mechanics?
Document [3] ...

Answers: In quantum mechanics, spin refers to the intrinsic angular momentum of particles, such as electrons, which can have either a spin up or a spin down state [1][2]. Spin is ...

Table 7: An example from ELI5 test set demonstrating the redundant citations error. None of the documents from [1] to [5] mention anything about spin states. The LLM retrieves its parametric knowledge when answering the question, but it still inserts the citations.

Instruction: Provide a ...

Instruction: Provide a ...

Question: Michael Kiske is part of what music group?

Document [1] ...
Document [2] ...
Document [3](Michael Kiske): ... In 2006, Michael Kiske released his third full-length solo album under the name of "Kiske". The album contained songs in a mellow, acoustic singer-songwriter style and was a step away from the sound of his previous solo albums. ...
Document [4] ...

Answers: Helloween [1], Unisonic [1], SupaRed [3], Kiske [3], Ill Prophecy [5].

Table 8: An example from QAMPARI test set demonstrating the misinterpretation of a single passage error. "Kiske" is actually the name of Michael Kiske's album, but is incorrectly interpreted by the LLM as the name of a music group where Michael Kiske is a member.

Instruction: Write an ...

Question: Purdue football players in nfl hall of fame?

Document [1] ...

Document [2](Purdue Boilermakers football): ... Among the numerous Boilermakers that have participated in the NFL, CFL, and AFL, a total of 34 have received all-star recognition by their respective

Document [3](Purdue Boilermakers football): ... A total of seventeen Boilermakers have been inducted into the College Football Hall of Fame. Four Boilermakers hold the distinguished title of Pro Football Hall of Fame inductees. ...

Document [4] ...

Answers: As of September 2018, there are a total of 34 Purdue Boilermakers who have received all-star recognition by their respective NFL teams [2]. Four of these players have been inducted into the Pro Football Hall of Fame [3]. ...

Table 9: An example from ASQA test set demonstrating the incorrect synthesis of multiple passages error. The coreference *these* in LLM’s response asserts that the 4 boilermakers who have been inducted into the Pro Football Hall of Fame are among the 34 boilermakers who have received all-star recognition, while such connection is not established by the original passages [2] and [3].

the prompt only in the distillation step (§ 2.1) and the ICL baselines.

J Examples

We provide an example for ASQA, QAMPARI, ELI5, and EXPERTQA each, presented in Table 14, 15, 16, and 17 respectively. The responses are generated by the f.g.(RS + RL) checkpoints trained in the *combined* setting.

	EXPERTQA			
	Correctness	Citation	Fluency	
	Claim Rec	Rec	Prec	MAUVE
ICL ChatGPT	19.27	47.79	47.30	48.68
ICL LLaMA-2-7B	10.32	10.09	7.79	34.27
\mathcal{M}_{dist}	15.28	49.03	46.22	40.63
+f.g.RL	15.53	49.73	51.11	45.92
+f.g.RS	15.48	59.46	57.58	44.67
+f.g.(RS+RL)	15.23	58.94	59.80	42.13

Table 10: Result on EXPERTQA.

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: Who played galen in planet of the apes?

Document [1](Title: Planet of the Apes): {text 1}
Document [2](Title: Planet of the Apes (1968 film)): {text 2}
Document [3](Title: Planet of the Apes (1968 film)): {text 3}
Document [4](Title: Planet of the Apes): {text 4}
Document [5](Title: Planet of the Apes): {text 5}

Answer:In the 1968 film Planet of the Apes, Galen was played by Wright King [2]. And in the tv series Planet of the Apes, Galen was played by Roddy McDowall [1].

Question: Which is the most rainy place on earth?

Document [1](Title: Cherrapunji): {text 1}
Document [2](Title: Cherrapunji): {text 2}
Document [3](Title: Mawsynram): {text 3}
Document [4](Title: Earth rainfall climatology): {text 4}
Document [5](Title: Going to Extremes): {text 5}

Answer:Several places on Earth claim to be the most rainy, such as Lloróof3, Colombia, which reported an average annual rainfall of 12,717 mm between 1952 and 1989, and Lóof3pez de Micay, Colombia, which reported an annual 12,892 mm between 1960 and 2012 [3]. However, the official record is held by Mawsynram, India with an average annual rainfall of 11,872 mm [3], although nearby town Sohra, India, also known as Cherrapunji, holds the record for most rain in a calendar month for July 1861 and most rain in a year from August 1860 to July 1861 [1].

Question: {question}

Document [1](Title: title 1): {text 1}
Document [2](Title: title 2): {text 2}
Document [3](Title: title 3): {text 3}
Document [4](Title: title 4): {text 4}
Document [5](Title: title 5): {text 5}

Answer:

Table 11: Prompt and demonstrations for ASQA.

Instruction: Provide a list of accurate answers for the given question using only the provided search results (some of which might be irrelevant) and cite them properly. Always cite one and only one document for each answer. Separate answers by commas. For questions that have more than 5 answers, write at least 5 answers.

Question: Which film has Gong Li as a member of its cast?

Document [1](Title: Gong Li): {text 1}
Document [2](Title: Gong Li): {text 2}
Document [3](Title: Gong Li): {text 3}
Document [4](Title: Zhang Yimou): {text 4}
Document [5](Title: Gong Li): {text 5}

Answer: The Story of Qiu Ju [1], Farewell My Concubine [2], Flirting Scholar [2], The Monkey King 2 [3], Mulan [3], Saturday Fiction [3], Coming Home [3].

Question: Glenn Ford was a member of cast in which film?

Document [1](Title: Glenn Ford): {text 1}
Document [2](Title: Glenn Ford): {text 2}
Document [3](Title: CBS Thursday Night Movie): {text 3}
Document [4](Title: The Trouble with Girls (film)): {text 4}
Document [5](Title: Trouble in the Glen): {text 5}

Answer: So Ends Our Night [1], Heaven with a Barbed Wire Fence [1], Happy Birthday to Me [2], The Greatest Gift [2], The Gift [2], The Brotherhood of the Bell [3].

Question: {question}

Document [1](Title: title 1): {text 1}
Document [2](Title: title 2): {text 2}
Document [3](Title: title 3): {text 3}
Document [4](Title: title 4): {text 4}
Document [5](Title: title 5): {text 5}

Answer:

Table 12: Prompt and demonstrations for QAMPARI.

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: What's the difference between Shia vs. Sunni Islam?

Document [1](Title: The Sunni vs Shia Divide - Explained - Globaloi): {text 1}

Document [2](Title: What's the difference between Sunni and Shia Islam? 2019 MacroSnap): {text 2}

Document [3](Title: Difference between Sunni and Shia Muslims | Sunni vs Shia Muslims): {text 3}

Document [4](Title: What is the difference between Shia and Sunni Islam? - Islam Stack Exchange): {text 4}

Document [5](Title: What is the difference between Sunni and Shia Islam? | Patrick Syder Travel): {text 5}

Answer: The main difference between Shia and Sunni Muslim is related to ideological heritage and issues of leadership [1]. This difference is first formed after the death of the Prophet Muhammad in 632 A.D. [1][2]. The ideological practice of the Sunni branch strictly follows Prophet Muhammad and his teachings, while the Shia branch follows Prophet Muhammad's son-in-law Ali [2]. Nowadays, Sunni and Shia are the major branches of Islam [3].

Question: How do student loans affect getting a mortgage?

Document [1](Title: Student Loans 2019 How do they work? | The Financial Review): {text 1}

Document [2](Title: How Does Student Loan Debt Affect Buying a Home? | Experian): {text 2}

Document [3](Title: Studentloanify - How your student loans affect your home mortgage prospects): {text 3}

Document [4](Title: How do student loans affect your credit score? | Student Loan Planner): {text 4}

Document [5](Title: Does Student Loan Debt Affect Getting A Mortgage?): {text 5}

Answer: When applying for a mortgage, student loans can affect the debt to income ratio, which is a key factor in determining the amount that an individual can afford to pay for the mortgage [1]. While student loan repayments do not appear in an individual's credit history and do not affect credit scores, lenders do consider the amount of an individual's student loan repayments when assessing their mortgage application [1][2][3]. Some 83% of non-homeowners say student loan debt is preventing them from buying a home, according to the National Association of Realtors [2]. It is important to note that student loans do not prevent an individual from getting a mortgage [1].

Question: {question}

Document [1](Title: title 1): {text 1}

Document [2](Title: title 2): {text 2}

Document [3](Title: title 3): {text 3}

Document [4](Title: title 4): {text 4}

Document [5](Title: title 5): {text 5}

Answer:

Instruction: Write an . . .

Question: Who sang the song i'm coming out?

Document [1](Title: I'm Coming Out): I'm Coming Out "I'm Coming Out" is a song recorded by American singer Diana Ross. It was written and produced by Chic members Bernard Edwards and Nile Rodgers, and released in August 22, 1980 as the second single from Ross' self-titled tenth album "Diana" (1980). In 1979, Ross commissioned Chic founders Nile Rodgers and Bernard Edwards to create material for a new album after taking her daughters to see the band in concert, following the success of Ross's final Ashford & Simpson-composed and produced LP, "The Boss." "Rodgers got the idea for "I'm Coming Out" after noticing three different drag

Document [2](Title: I'm Coming Out): the Loren Dawson remix. The original single version (which is similar to the original) can be found on both CD singles and the original motion picture soundtrack to "Maid in Manhattan". Ross' version appears on the film's soundtrack album as well. A music video was produced to promote the single. I'm Coming Out "I'm Coming Out" is a song recorded by American singer Diana Ross. It was written and produced by Chic members Bernard Edwards and Nile Rodgers, and released in August 22, 1980 as the second single from Ross' self-titled tenth album "Diana" (1980). In 1979, Ross commissioned Chic

Document [3](Title: I'm Coming Out): "I'm Coming Out" and its resulting popularity has led to Ross's song often being sampled, most notably by Stevie J, who sampled the song for rapper The Notorious B.I.G.'s 1997 song "Mo Money Mo Problems". There's also been a cover by Marcia Hines. It was featured as a cover for the video game "Karaoke Revolution" as well. In more recent times, the song was sampled in Keyshia Cole's remake of "I'm Coming Out" (featuring Iggy Azalea) for the original soundtrack of the 2014 movie The Other Woman, starring Cameron Diaz, Kate Upton and Leslie Mann. The song is sampled in

Document [4](Title: I'm Coming Out): "Break Your Heart Right Back" from Ariana Grande's sophomore album "My Everything" (2014). This song was also covered for the soundtrack to the 2016 movie "Trolls". In 1997, German comedian Hella von Sinnen sampled "I'm Coming Out" for her single "Mein Coming Out", using the term "coming out" in the context of coming out of the closet, or revealing one's homosexuality. In 2003, American singer Amerie covered "I'm Coming Out" for the soundtrack to the film "Maid in Manhattan". Her version was released as a single in select European countries and Australia. The final single version to be released was

Document [5](Title: Coming Out (album)): "Poinciana (The Song of the Tree)" featured a solo by Michael Brecker. "Coming Out" debuted on Billboard's Top Pop Album chart on September 18, 1976 and peaked at No. 48. The album was re-released on CD in March 1993. Manhattan Transfer Musicians Production Coming Out (album) Coming Out is the third album by The Manhattan Transfer, released August 19, 1976 on Atlantic Records. On the advice of their manager Aaron Russo, the group downplayed the "nostalgic" tenor of their previous breakthrough album in favor of more contemporary songs. However, the throwback single "Chanson d'Amour" became a breakout single in Europe

Answer: "I'm Coming Out" is a song recorded by Diana Ross [1][2]. It was written and produced by Chic members Bernard Edwards and Nile Rodgers [1].

Gold: "I'm Coming Out" is a song recorded by American singer Diana Ross. It was written and produced by Chic members Bernard Edwards and Nile Rodgers, and released on August 22, 1980, as the second single from Ross' self-titled tenth album Diana (1980). In 2003, American singer Amerie covered "I'm Coming Out" for the soundtrack to the film "Maid in Manhattan". Her version was released as a single in select European countries and Australia.

Gold Answers: [Diana Ross, Amerie]

Table 14: ASQA example

Instruction: Provide a . . .

Question: Who is a director of a film produced by Kirk Douglas?

Document [1](Title: Kirk Douglas): began producing films as varied as "Paths of Glory" (1957) and "Spartacus" (1960). In those two films, he collaborated with the then-relatively-unknown director Stanley Kubrick taking lead roles in both films. Douglas has been praised for helping to break the Hollywood blacklist by having Dalton Trumbo write "Spartacus" with an official on-screen credit, although this has been disputed by others. He produced and starred in "Lonely Are the Brave" (1962), considered a classic, and "Seven Days in May" (1964), opposite Burt Lancaster, with whom he made seven films. In 1963, he starred in the Broadway play "One Flew Over the

Document [2](Title: Kirk Douglas): executive producer as well, raising the \$12 million production cost, making it one of the most expensive films made up to that time. Douglas initially selected Anthony Mann to direct, but replaced him early on with Stanley Kubrick, with whom he previously collaborated in "Paths of Glory". When the film was released, Douglas gave full credit to its screenwriter, Dalton Trumbo, who was on the Hollywood blacklist, and thereby effectively ended it. About that event, he said, "I've made over 85 pictures, but the thing I'm most proud of is breaking the blacklist." However the film's producer Edward Lewis and

Document [3](Title: Spartacus (film)): credibility. Trumbo was publicly given credit for two major films. Otto Preminger made public that Trumbo wrote the screenplay for his film "Exodus", and Kirk Douglas publicly announced that Trumbo was the screenwriter of "Spartacus". Further, President John F. Kennedy publicly ignored a demonstration organized by the American Legion and went to see the film. After David Lean turned down an offer to direct, "Spartacus" was to be directed by Anthony Mann, then best known for his Westerns such as "Winchester '73" and "The Naked Spur". Douglas fired Mann at the end of the first week of shooting, in which

Document [4](Title: Home Movies (film)): directing films. Kirk Douglas plays a film instructor loosely modeled on the director himself. Keith Gordon is one of his pupils who films everything that happens at home. Many of the events that happen to Gordon's character were modeled on events from De Palma's own adolescence, particularly rivalry with a more favored brother, a mother prone to dramatic outbursts, and a philandering father. Home Movies (film) Home Movies is a 1979 independent film directed by Brian De Palma and starring Kirk Douglas, Nancy Allen, Vincent Gardenia, Keith Gordon, Theresa Saldana, and Gerrit Graham. De Palma had been teaching film at

Document [5](Title: Kirk Douglas): in the picture," adding that "he's as bright a person as I've met in the acting profession." She says that his "pragmatic approach to acting" would later be a "philosophy that ended up rubbing off on me." Between 1970 and 2008, Douglas made nearly 40 movies and appeared on various television shows. He starred in a western, "There Was a Crooked Man..." (1970), alongside Henry Fonda. The film was produced and directed by Joseph L. Mankiewicz. In 1973, he directed his first film, "Scalawag". Also in 1973, Douglas appeared in a made-for-TV musical version of "Dr. Jekyll and Mr. Hyde".

Answer: [Stanley Kubrick \[1\]](#), [Anthony Mann \[2\]](#), [Dalton Trumbo \[2\]](#), [Brian De Palma \[4\]](#), [Joseph L. Mankiewicz \[5\]](#).

Gold: [Stanley Kubrick](#), [Kevin Billington](#), [Kirk Douglas](#), [Anthony Newley](#), [Andr e de Toth](#)

Gold Answers: [\[\[Stanley Kubrick\]](#), [\[Kevin Billington\]](#), [\[Kirk Douglas\]](#), [\[Anthony Newley\]](#), [\[Andr e de Toth\]\]](#)

Table 15: QAMPARI example

Instruction: Write an . . .

Question: What exactly is 'spin' in quantum mechanics?

Document [1](Title: quantum mechanics - Does spin have anything to do with a rate of change? - Physics Stack Exchange): quantum mechanics - Does spin have anything to do with a rate of change? - Physics Stack Exchange the electron. Not the answer you're looking for? Browse other questions tagged quantum-mechanics operators quantum-spin rotation spinors or ask your own question. How do you rotate spin of an electron? Did the Goudsmit-Uhlenbeck analysis of spin consider relativity? Is there an Ehrenfest-like result for the expectation value of orbital angular momentum? Can spin be related to a shift in angle? What does spin 0 mean exactly? Does String Theory explain spin? How can it be derived that particles described by the Dirac equation must have spin 1/2? When we say electron spin is 1/2, what exactly does it mean, 1/2

Document [2](Title: quantum mechanics - Does spin have anything to do with a rate of change? - Physics Stack Exchange): quantum mechanics - Does spin have anything to do with a rate of change? - Physics Stack Exchange your own question. Did the Goudsmit-Uhlenbeck analysis of spin consider relativity? Can spin be related to a shift in angle? What does spin 0 mean exactly? Does String Theory explain spin? How can it be derived that particles described by the Dirac equation must have spin 1/2? When we say electron spin is 1/2, what exactly does it mean, 1/2 of what? Determine spin-state with $B = B_x$ in the S_z base why do we relate spin with angular momentum operator in quantum mechanics?

Document [3](Title: Search results for 'Biology' - PhilArchive): Search results for 'Biology' - PhilArchive biased to satisfy the collective goals of the ensemble. Using this postulate, we mathematically derive the correct spin probability distribution without using quantum mechanical formalism (operators and Born's rule) and exactly reproduce the quantum mechanical spin correlation in entangled pairs. Using our concept, we also mathematically derive the form of quantum mechanical wave function of free particle which is conventionally a postulate of quantum mechanics. Thus, we prove that the origin of quantum mechanical results lies in the will (or consciousness) of the objects biased by the collective goal of ensemble or universe. This biasing by the group on individuals

Document [4](Title: The Upside to Quantum Mechanics): The Upside to Quantum Mechanics Questions About Quantum Mechanics Which is named Quantum Tunneling. The quantum simulator is installed as part of the QDK. There's no macroscopic property that's equivalent to quantum spin. The traditional manner of operating the Hopfield network is to continue picking neurons at random and updating them by thinking about the connected neurons, together with their weights. Basically, the concept is that, when you have a look at a system (such as an electron that's spinning in two directions at exactly the same time) the electron custom essays is going to be made to quit spinning in two directions simultaneously, pick

Document [5](quantum mechanics - How exactly is fusion possible? - Physics Stack Exchange): quantum mechanics - How exactly is fusion possible? - Physics Stack Exchange Pauli's Exclusion Principle Can two electrons have the same momentum and spin directions? Nature of the quantum degeneracy pressure How does the Pauli-exclusion principle work if space is infinitely divisible? What exactly is meant by a quantum state in QM? Why a half-integer spin implies antisymmetry on the wave function? What is the physical (i.e. non-mathematical) cause off the difference between a system of identical bosons and fermions? How do two exactly the same (in space) wavefunctions of two non-distinguishable free particles contribute to Pauli's exclusion principle? How do I derive Pauli's exclusion principle with path integrals?

Answer: In quantum mechanics, spin refers to the intrinsic angular momentum of particles, such as electrons, which can have either a spin up or a spin down state [1][2]. Spin is not related to a rate of change, but rather is a fundamental property of particles that can be described by quantum mechanical operators and the Pauli exclusion principle [1][5]. While there are various interpretations of quantum mechanics, including the idea that particles described by the Dirac equation must have spin 1/2 [2], the origin of quantum mechanical results is not necessarily related to the will or consciousness of the objects, as suggested in document [3].

Gold: Spin is the quantum mechanical analogue to rotational angular momentum. Rotational angular momentum is a measure of how strongly an object is rotating on its own axis. The Earth's revolution about its own axis (the one that causes days) is an example of rotational angular momentum.

The reason I said that spin is an "analogue" to rotational angular momentum is that it's not an idea that really makes sense in QM. In order for us to see how an object is rotating, we have to be able to clearly see what shape it is. If you can't see a shape, rotation is sort of meaningless, since... well, you can't actually tell if something is rotating if it has no shape.

Quantum mechanical objects don't have well-defined shapes. They're sort of fuzzy and nebulous, so you can't really pin them down like you'd want. We hence can't look at, say, an electron and go "oh, it's rotated by 180 degrees after two seconds, so its angular momentum is X." Electrons can't really rotate, it doesn't make sense.

However, electrons do exhibit certain properties that make it *seem like* they're rotating. For example, when a charged (non-quantum mechanical object) rotates, it makes a magnetic field. Electrons also make a magnetic field in the exact same way. So, they seem to be acting like they're rotating, even though we can't really pin down what that means exactly.

We thus call it "spin." It's... like they're rotating, hence "spin," but it's understood by physicists to not mean the exact same thing.

Spin is "quantized," which means that there are only certain values that it can have. If you measure an electron's spin, you will only either get $\hbar/2$ (clockwise), or $-\hbar/2$ (counterclockwise) (1). \hbar is the "reduced [Planck constant](_URL_0_)," which shows up a lot in physics.

^ (1. note that when I say "measure," what's implied is that we're choosing an axis upon which to measure. For example, if you wanted to measure the Earth's angular momentum, you would measure it along the Earth's axis. The axis you choose to measure on for a quantum mechanical object is irrelevant though.)

Gold Answer Claims:

Spin in quantum mechanics is analogous to rotational angular momentum.

Spin is not a concept that makes sense in quantum mechanics as objects do not have well-defined shapes.

Spin is quantized and can only have certain values when measured.

Table 16: ELI5 example

Instruction: Write an . . .

Question: Which books are the best for child and youth psychiatry?

Document [1](Title: job details - Consulting Psychiatrist - Department of Psychiatry, Ambulatory mental health program, BC Children's Hospital in Vancouver): job details - Consulting Psychiatrist - Department of Psychiatry, Ambulatory mental health program, BC Children's Hospital in Vancouver themes. Through the affiliation with the BCHRI, the Department has established research strengths in neuroimaging, genetics, metabolism and clinical research. The Department is actively involved in supporting best practice through multiple provincial networks for which it provides infrastructure support (Eating Disorders, Health Literacy, Youth Concurrent Disorders and Child & Youth Mental Health and Substance Use), through its robust and award winning Child & Youth Mental Health Literacy Program and Child Health BC. Interested candidates will have obtained fellowship standing in the Royal College of Physicians & Surgeons of Canada in Psychiatry along with certification in Child and Adolescent Psychiatry subspecialty

Document [2](Title: Textbook of Child and Adolescent Psychiatry - Google Books): Textbook of Child and Adolescent Psychiatry - Google Books the Journal of the American Academy of Child and Adolescent Psychiatry. She has published more than 60 articles, book chapters, and books, including Concise Guide to Child and Adolescent Psychiatry, and Helping Parents, Youth, and Teachers Understand Medications for Behavioral and Emotional Problems, published by American Psychiatric Publishing, Inc. She is the principal author of the American Academy of Child and Adolescent Psychiatry Practice parameters for the assessment and treatment of children, adolescents, and adults with attention-deficit/hyperactivity disorder and a co-author of the American Academy of Child and Adolescent Psychiatry practice parameters for the use of stimulant medications. Title Textbook

Document [3](Title: Journal of the American Academy of Child and Adolescent Psychiatry Page 5): Journal of the American Academy of Child and Adolescent Psychiatry Page 5 in evidence-based treatments for youth depression in recent decades, overall treatment effects are modest at best, with 30-50% of youth being non-responders. Practice parameters consistently recommend systematic assessment and routine monitoring of depressive symptoms, or measurement-based care (MBC), to enhance youth depression treatment. However, the literature offers few guidelines on how to use assessment results to inform care decisions or detect real and clinically meaningful change. (Source: Journal of the American Academy of Child and Adolescent Psychiatry) Source: Journal of the American Academy of Child and Adolescent Psychiatry - February 27, 2019 Category: Psychiatry Authors: Freda F. Liu, Molly C.

Document [4](Title: Browse In Child and Adolescent Psychiatry | Oxford Medicine Online - Oxford Medicine): Browse In Child and Adolescent Psychiatry | Oxford Medicine Online - Oxford Medicine Child and Adolescent Psychiatry x Download complete list of books in this Child and Adolescent Psychiatry Collection (.pdf) (.xls) Bipolar Disorder in Youth: Presentation, Treatment and Neurobiology Stephen M. Strakowski, Melissa P. DelBello, and Caleb M. Adler (eds) Bipolar disorders affect up to 3% of the world's population and are the 6th leading cause of disability worldwide. As the disorder presents differently in youth as it does in adults, ... More Bipolar disorders affect up to 3% of the world's population and are the 6th leading cause of disability worldwide. As the disorder presents differently in youth as it does

Document [5](Title: Rutter's Child and Adolescent Psychiatry -ORCA): Rutter's Child and Adolescent Psychiatry -ORCA new look." -Professor Sir Michael Rutter, extract from Foreword. Reviews of previous editions: "This book is by far the best textbook of Child & Adolescent Psychiatry written to date." -Dr Judith Rapoport, NIH "The editors and the authors are to be congratulated for providing us with such a high standard for a textbook on modern child psychiatry. I strongly recommend this book to every child psychiatrist who wants a reliable, up-to-date, comprehensive, informative and very useful textbook. To my mind this is the best book of its kind available today." -Journal of Child Psychology and Psychiatry

Answer: [There are several books that are considered the best for child and youth psychiatry, including the Textbook of Child and Adolescent Psychiatry \[2\], Rutter's Child and Adolescent Psychiatry \[5\], and Concise Guide to Child and Adolescent Psychiatry \[2\].](#)

Gold: The books recommended for child and youth psychiatry include "Textbook of Child and Adolescent Psychiatry" which is noted in the Journal of the American Academy of Child and Adolescent Psychiatry and "Rutter's Child and Adolescent Psychiatry", which was described as "the best textbook of Child & Adolescent Psychiatry written to date" and a "very useful textbook" by Dr. Judith Rapoport at NIH and the Journal of Child Psychology and Psychiatry, respectively . Other resources could be found in the category of Child and Adolescent Psychiatry in Oxford Medicine Online, which includes a list of books such as "Bipolar Disorder in Youth: Presentation, Treatment and Neurobiology" .

Gold Answer Claims:

"Textbook of Child and Adolescent Psychiatry" is recommended for child and youth psychiatry.,

"Rutter's Child and Adolescent Psychiatry" has been described as "the best textbook of Child & Adolescent Psychiatry written to date".,

Oxford Medicine Online has a list of books in the category of Child and Adolescent Psychiatry, such as "Bipolar Disorder in Youth: Presentation, Treatment and Neurobiology" .

Table 17: EXPERTQA example