

# Enabling Large Language Models to Generate Text with Citations

Tianyu Gao   Howard Yen   Jiatong Yu   Danqi Chen

Department of Computer Science, Princeton University

{tianyug,danqic}@cs.princeton.edu   {hyen,jiatongy}@princeton.edu

## Abstract

Large language models (LLMs) have emerged as a widely-used tool for information seeking, but their generated outputs are prone to hallucination. In this work, we aim to enable LLMs to generate text with *citations*, improving their factual correctness and verifiability. Existing work mainly relies on commercial search engines and human evaluation, making it challenging to reproduce and compare with different modeling approaches. We propose **ALCE**, the first benchmark for **A**utomatic **L**LMs’ **C**itation **E**valuation. ALCE collects a diverse set of questions and retrieval corpora and requires building end-to-end systems to retrieve supporting evidence and generate answers with citations. We build automatic metrics along three dimensions—fluency, correctness, and citation quality—and demonstrate their strong correlation with human judgments. Our experiments with state-of-the-art LLMs and novel prompting strategies show that current systems have considerable room for improvements—for example, on the ELI5 dataset, even the best model has 49% of its generations lacking complete citation support. Our extensive analyses further highlight promising future directions, including developing better retrievers, advancing long-context LLMs, and improving the ability to synthesize information from multiple sources.<sup>1</sup>

## 1 Introduction

Large language models (LLMs; Brown et al., 2020) have gained increasing popularity as a tool for information seeking. While they generate engaging and coherent responses, their output is prone to hallucination and often contains factually incorrect information (Ji et al., 2023). This makes it harder for users to trust and verify LLMs’ generated outputs without any supporting evidence.

<sup>1</sup>Our code and data are available at <https://github.com/princeton-nlp/ALCE>.

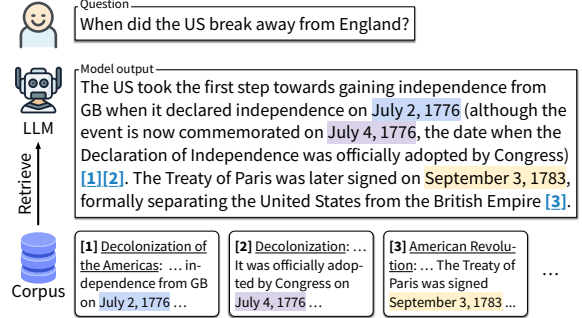


Figure 1: The task setup of ALCE. Given a question, the system generates text while providing *citing passages* from a large retrieval corpus. Each statement may contain multiple citations (e.g., [1][2]).

In this work, we study a new generation paradigm for LLMs, where we require LLMs to provide *citations* to one or a few text passages for any statement they generate (Figure 1). Incorporating citations brings several benefits: (1) Users can easily verify LLMs’ claims with the provided citations; (2) LLMs can generate text that faithfully follows cited passages, which has the promise to improve correctness and alleviate hallucination.

Multiple commercial systems have adopted this paradigm: Bing Chat<sup>2</sup> and perplexity.ai<sup>3</sup> respond to user questions in natural language with references to Web pages. Nakano et al. (2021); Menick et al. (2022) share a similar motivation, but they mainly experiment with search engines and closed-source models, making their results difficult to evaluate. Retrieval-augmented LLMs (Borgeaud et al., 2022; Izacard et al., 2022) incorporate retrieved passages during both training and inference, but do not guarantee faithfulness to retrieved passages or explicitly provide citations. Additionally, previous studies mostly rely on human evaluation (Nakano et al., 2021; Menick et al., 2022; Liu et al., 2023), which is expensive and difficult to reproduce. We argue

<sup>2</sup><https://www.bing.com/new>

<sup>3</sup><https://www.perplexity.ai>

Dataset	Corpus (#passages)	Question type	Example
ASQA	Wikipedia (21M)	Factoid (ambiguous)	Q: When did the US break away from England? A: The US declared independence on July 2, 1776 [1][2] ... The Treaty of Paris was later signed on September 3, 1783 [3].
QAMPARI	Wikipedia (21M)	Factoid (list)	Q: Which films have Gong Li as a member of their cast? A: The Story of Qiu Ju [1], Farewell My Concubine [2], The Monkey King 2 [3], Mulan [3], Saturday Fiction [3] ...
ELI5	Sphere (899M)	Why/How/What	Q: How do student loans affect getting a mortgage? A: Student loans can affect the debt to income ratio [1], which is a key factor in determining the amount that ... [2][3]

Table 1: The three datasets used in our ALCE benchmark. These datasets cover a wide range of question types and the corresponding corpora span from limited domains (Wikipedia) to Web scale (Sphere).

that the absence of automated evaluation hinders the advances of such systems.

We present **ALCE**, the first reproducible benchmark for automatically evaluating LLMs’ generations with citations. ALCE assumes a natural-language question and a retrieval corpus, and requires building end-to-end systems to retrieve relevant passages from the corpus, generate a response to the question, and cite corresponding supporting passages. We collect three datasets that cover different types of questions and corpora—ASQA (Stelmakh et al., 2022), QAMPARI (Rubin et al., 2022), and ELI5 (Fan et al., 2019)—as shown in Table 1. Different from previous benchmarks (Lee et al., 2019; Bohnet et al., 2022), ALCE evaluates long-text generation, focusing on automatically evaluating citation quality, and allows citing multiple passages for individual statements.

We design automatic evaluation methods on the three dimensions: **fluency**, **correctness**, and **citation quality**. Specifically, we use MAUVE (Pillutla et al., 2021) to measure fluency, adopt a natural language inference (NLI) model (Honovich et al., 2022) to measure citation quality, and propose specialized correctness evaluation for each dataset. We showcase how the three dimensions together contribute to a robust evaluation, preventing systems from exploiting shortcuts. Additionally, we conduct human evaluation and demonstrate a strong correlation with our automatic metrics.

We experiment on multiple systems with state-of-the-art LLMs and retrievers, and also propose novel prompting strategies to synthesize and incorporate retrieved text into text generation. Although all systems are capable of providing fluent and coherent responses, there remains substantial room for improvement in terms of correctness and citation quality: For example, on the ELI5 dataset, 49% generations of our ChatGPT baseline are not fully

supported by the cited passages. Additionally, we find that (1) a closed-book model (generating answers without accessing any retrieved documents) with post-hoc citing achieves good correctness but much worse citation quality; (2) though interactive retrieval approaches (Yao et al., 2023; Schick et al., 2023) offer more flexibility in when/what to retrieve, we find that they do not improve the performance; (3) summarizing the retrieved passages in shorter text improves correctness but not citation quality; (4) reranking multiple generations boosts citation quality measured by human evaluation.

Our extensive analyses highlight three major challenges of building LLMs to generate text with citations: (1) The retrieval quality is crucial to the final performance and has a substantial room for improvement; (2) LLMs’ limited context window restricts the number of passages/tokens they can incorporate; (3) Current LLMs struggle to synthesize multiple documents in context without being distracted by irrelevant ones. These challenges pose promising research directions for developing better systems integrating retrieval and LLMs.

## 2 Task Setup and Datasets

Our task is formalized as follows: Given a query  $q$  and a corpus of text passages  $\mathcal{D}$ , the system is required to return an output  $\mathcal{S}$ , which consists of  $n$  statements  $s_1, \dots, s_n$ , and each statement  $s_i$  cites a list of passages  $\mathcal{C}_i = \{c_{i,1}, c_{i,2}, \dots\}$ <sup>4</sup>, where  $c_{i,j} \in \mathcal{D}$ . In this work, we segment LLMs’ output into statements by sentence boundaries, as one sentence usually conveys a coherent statement while being short enough to be easily verified.<sup>5</sup> While LLMs may include sentences that do not require

<sup>4</sup>In practice, we allow at most 3 citations for each statement as more citations usually do not help.

<sup>5</sup>QAMPARI requires a list as the answer, and we choose each entity in the generated list as a statement.

a citation, such as “*I’m happy to help*”, we observe that almost all sentences that LLMs output in response to our input questions provide valuable information that should be verified and require citations, similar to the findings of Liu et al. (2023). In this work, citations are enclosed by box brackets such as [1][2].

We divide the corpus  $\mathcal{D}$  into 100-word passages following previous works on open-domain question answering (Karpukhin et al., 2020; Petroni et al., 2021; Piktus et al., 2021). This is in contrast to commercial systems like Bing Chat, which often cite entire Web pages. Taking short 100-word passages has several advantages: (1) It is easier for humans to verify; (2) It allows for more retrieved passages to fit in LLMs’ limited context.

We choose QA datasets so that (1) they contain factual questions, in which references are important; (2) questions require long-text answers that cover multiple aspects; (3) answering the questions requires synthesizing multiple sources. We select three datasets (Table 1) and introduce them below.

**ASQA** (Stelmakh et al., 2022) is a long-form factoid dataset. Each question is an ambiguous question from AmbigQA (Min et al., 2020) and requires multiple short answers to cover different interpretations of the question. For example, the question “*When did the US break away from England?*” should be answered with both *July 2, 1776* (declaration of independence) and *September 3, 1783* (Treaty of Paris). The ambiguity nature of the questions necessitates synthesizing information from multiple documents. Stelmakh et al. (2022) crowdsourced long-form answers for each question to cover all short, disambiguated answers while providing essential background information, as shown in the example in Figure 1. Human-written answers for ASQA have an average length of 65 words. Since most of the questions can be answered by Wikipedia, we use the 2018-12-20 Wikipedia snapshot as  $\mathcal{D}$ . Using a small corpus allows researchers to deploy and study different retrieval methods at low cost, especially dense retrieval (Karpukhin et al., 2020; Ni et al., 2022).

**QAMPARI** (Rubin et al., 2022) is a factoid QA dataset, where the answer is a list of entities that are usually drawn from different passages. The dataset is automatically constructed by leveraging Wikipedia knowledge graphs and tables, with questions manually paraphrased and answers verified. This dataset represents a common “listing” type of

questions people ask, for example, “*Which films have Gong Li as a member of their cast?*” Models need to synthesize over multiple passages to achieve high precision and recall on the dataset, posing a unique challenge. On average, each question in QAMPARI has 13 answers. Same as ASQA, we use the 2018-12-20 Wikipedia as the corpus.

**ELI5** (Fan et al., 2019) is a long-form QA dataset built on the Reddit forum “Explain Like I’m Five”,<sup>6</sup> with an average answer length of 131 words. Most ELI5 questions are how/why/what questions that require in-depth long answers and multiple passages as evidence. Due to the diverse range of topics discussed in the questions, we use Sphere (Piktus et al., 2021)—a filtered version of Common Crawl<sup>7</sup>—as the corpus. The ELI5 dataset is widely used in related work due to its challenging nature (Nakano et al., 2021; Menick et al., 2022; Liu et al., 2023).

We randomly select 1,000 examples from the development set of each dataset for ALCE. Our benchmark primarily assesses the citation capabilities of existing LLMs and does not provide training data, as there are no available examples that provide supervision for citations in these datasets.

### 3 Automatic Evaluation

Our benchmark measures the following three dimensions of system responses:

- **Fluency**: whether the model’s generated text is fluent and coherent.
- **Correctness**: whether the answer is accurate and covers all aspects of interest.
- **Citation quality**: whether the answer is well supported by the cited passages and no irrelevant passages are cited.

In the following, we present automatic metrics for each dimension, and why combining the three metrics provides a robust evaluation (§3.4).

#### 3.1 Fluency

We use MAUVE (Pillutla et al., 2021) to evaluate the fluency of the output. MAUVE measures the similarity between two text distributions. Specifically, we concatenate the question and the model output and compare it to the distribution of question-gold-answer concatenation. We deploy

<sup>6</sup><https://www.reddit.com/r/explainlikeimfive/>

<sup>7</sup><https://commoncrawl.org>. We also filter out any Web pages on Reddit.

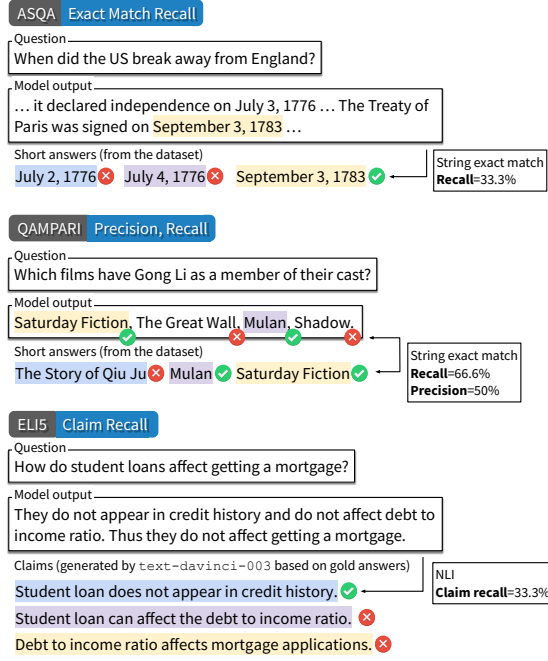


Figure 2: Evaluation of correctness (details in §3.2).

MAUVE for ASQA and ELI5 as QAMPARI only requires a list of short answers as the response. In our experiments, LLMs consistently adhere to the list format and thus we omit reporting fluency scores for the QAMPARI dataset.

### 3.2 Correctness

Our objective is to measure the informativeness and utility of the generation to the question. Liu et al. (2023) propose to directly evaluate *perceived utility* by humans, a process difficult to automate. Therefore, we use correctness—whether the response is accurate compared to a ground truth answer—as a proxy. Evaluating the correctness of long-form generation is a challenging task, and we describe our strategy for each task below.

For **ASQA**, we follow Stelmakh et al. (2022) and calculate the recall of correct short answers by checking whether the short answers (provided by the dataset) are exact substrings of the generation (*exact match recall*; EM recall).<sup>8</sup>

For **QAMPARI**, we follow Rubin et al. (2022) and calculate the *precision* and *recall* of the model prediction, by checking the exact match to the gold answer list,<sup>9</sup> with one additional adjustment: con-

<sup>8</sup>We use aliases of short answers provided by the dataset. We normalize the response and the short answers similarly to Rajpurkar et al. (2016). Stelmakh et al. (2022) also propose a QA-based evaluation which we found to be not as stable.

<sup>9</sup>QAMPARI also provides aliases, and we use them with answer normalization similar to ASQA.

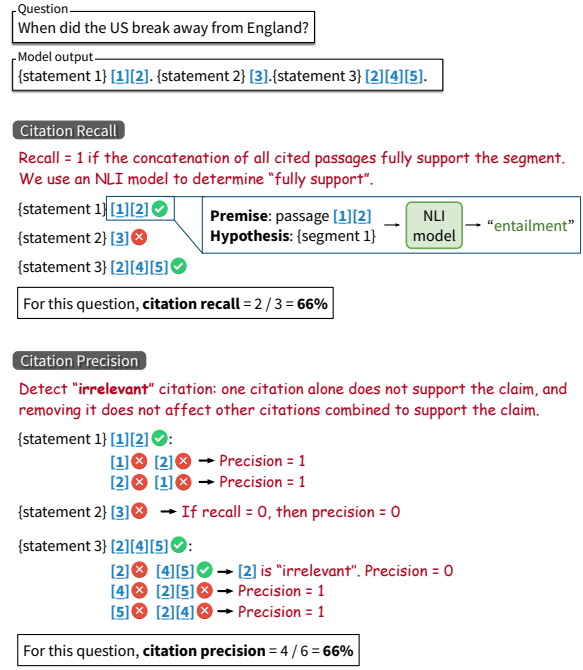


Figure 3: Evaluation of citation quality (details in §3.3). We use an NLI model to verify whether a statement is supported by its citations.

sidering that users often want to know only a few examples of “films that Gong Li starred in”, our evaluation considers recall to be 100% if the prediction includes at least 5 correct answers (*recall-5*).

The **ELI5** dataset only provides web-crawled long answers, and there are no short entity answers similar to ASQA and QAMPARI. Fan et al. (2019) use ROUGE for evaluation, which does not reflect the correctness of the answer well (shown in §A). Inspired by works in summarization evaluation and claim verification (Zhang and Bansal, 2021; Kamoi et al., 2023; Wang et al., 2020), we use Instruct-GPT (text-davinci-003; Ouyang et al., 2022) to generate three “sub-claims” and use a state-of-the-art natural language inference (NLI) model TRUE (Honovich et al., 2022) to check whether the model output entails the sub-claims (*claim recall*). We demonstrate that claim recall provides a more accurate measure of correctness than existing metrics. Please refer to §A for more details.

### 3.3 Citation Quality

We evaluate citation qualities using two metrics: (1) *citation recall*, which determines if the output is entirely supported by cited passages, and (2) *citation precision*, which identifies any irrelevant citations. Although we prioritize citation recall as it entails a well-supported and truthful answer, enhancing



precision is crucial for better user satisfaction, reducing the need for human review of extraneous passages. Figure 3 provides an illustrated example.

We use TRUE<sup>10</sup> (Honovich et al., 2022), a T5-11B (Raffel et al., 2020) model fine-tuned on a collection of NLI datasets to automatically examine whether the cited passages entail the model generation. TRUE targets factual correctness and has been used by previous works in similar context (Bohnet et al., 2022; Gao et al., 2022). We conduct human evaluation (§6) to demonstrate strong human correlation of our automatic evaluation procedure.

**Citation recall.** We first calculate the citation recall of *each statement* (0 or 1) and average over all statements in the model response. For each statement  $s_i$ , its citation recall is 1 if and only if there is at least one citation ( $\mathcal{C}_i \neq \emptyset$ ) and  $\phi(\text{concat}(\mathcal{C}_i), s_i) = 1$ , where  $\phi(\text{premise}, \text{hypothesis})$  is the NLI model that outputs 1 if the premise entails the hypothesis, and 0 otherwise;  $\text{concat}(\mathcal{C}_i)$  concatenates all passages in  $\mathcal{C}_i$  together (details in §B). The NLI evaluation is in accordance with the *attributable to identified sources* (AIS) framework (Rashkin et al., 2021):  $\phi(\text{concat}(\mathcal{C}_i), s_i) = 1$  implies that  $s_i$  is true based solely on  $\text{concat}(\mathcal{C}_i)$ .

**Citation precision.** Our citation precision evaluation detects citations that are irrelevant to the claim, but it does not require citing a minimal set and it permits citing redundant passages entailing similar claims. This design takes the following considerations: Human writing often cites multiple similar sources to enhance the credibility; human readers may appreciate multiple citations, especially when it pertains to critical claims such as medical advice.

We calculate the citation precision for *each citation* (0 or 1) and average over all citations in the response. To calculate the citation precision, we must first determine if a citation is “irrelevant”. Intuitively, a citation  $c_{i,j}$  is “irrelevant” if (a)  $c_{i,j}$  itself cannot support  $s_i$  and (b) removing  $c_{i,j}$  does not affect the rest of the citations to support  $s_i$ . Formally,  $c_{i,j}$  is “irrelevant” if and only if

- (a)  $\phi(c_{i,j}, s_i) = 0$ , AND
- (b)  $\phi(\text{concat}(\mathcal{C}_i \setminus \{c_{i,j}\}), s_i) = 1$ .

$c_{i,j}$  has a precision score of 1 if  $s_i$  has recall=1 and  $c_{i,j}$  is not irrelevant. For example (Figure 3),

when  $s_3$  cites three references [2][4][5] and recall=1, [2] is “irrelevant” if  $\phi([2], s_3) = 0$  and  $\phi([4][5], s_3) = 1$ . For condition (b) to work, we set recall=1 as a prerequisite.

Note that this algorithm overlooks the scenario when one citation partially supports the statement. For example, if [2] entails partial information of  $s_3$  that [4][5] also entails, [2] will be counted as “irrelevant” while it should not be penalized. Liu et al. (2023) conduct human evaluation on citation precision that addresses this problem and we discuss the differences in §D.

### 3.4 ALCE is Robust to Shortcut Cases

We showcase how ALCE evaluations are robust to two possible shortcuts in §C: (1) using the top-1 retrieved passage as the response and citing itself, and (2) using the first two sentences of the top-1 passage. Both cases have almost-perfect citation scores, but (1) has low fluency due to its unnaturally long length compared to human answers, and (2) has low correctness due to low coverage.

## 4 Modeling

In this section, we discuss three major modeling components for an ALCE system—retrieval, synthesis, post-editing—and introduce different variants that we evaluate in experiments.

### 4.1 Retrieval

In our experiments, we explore simple, off-the-shelf retrievers: We use dense retrievers like GTR (Ni et al., 2022) and DPR (Karpukhin et al., 2020) for Wikipedia, and BM25 for Sphere (as dense retrievers are expensive and slow for Web-scale corpus). For each question, we retrieve top-100 passages. Retrieving high-quality and high-coverage passages is crucial for both correctness and citation quality. We leave the exploration of more sophisticated retrievers to future work.

### 4.2 Synthesis

In this part, we focus on how to prompt a fixed LM to interact with the retriever, and synthesize and cite the evidences. One noteworthy challenge is that existing off-the-shelf LLMs all have limited context window and thus can only fit a handful of passages. We describe various strategies that we explore in this work below.

**VANILLA.** In this setup, we simply provide the

<sup>10</sup>[https://huggingface.co/google/t5-xxl-true\\_nli\\_mixture](https://huggingface.co/google/t5-xxl-true_nli_mixture). Details in §B.

---

Instruction: Write an accurate, engaging, and concise answer for ...

<Retrieve for the question>

Document [1](Title: American Decolonization) ...

Document [2](Title: Decolonization) ...

Document [3](Title: American Revolution) ...

...

Question: When did US break away from England?

Answer: The United States took the first step towards gaining independence ... [1][2]. The Treaty of Paris was later signed ... [3].

---

Table 2: An example of our VANILLA method. Different colors represent prompt, model generation, and <actions>. We also provide two in-context demonstrations before the test example.

---

Instruction: ...

<Retrieve for question "...">

Question: When did US break away from England?

Search: Declaration of Independence

<Search the query among the top-100 passages>

Document [1](Title: ...) ...

Output: The United States ... [1].

<Remove Document [1] from context>

Search: Treaty of Paris

<Search the query among the top-100 passages>

Document [3](Title: ...) ...

Output: The Treaty of Paris ... [3].

<Remove Document [3] from context>

End.

---

Table 3: An example of INLINESEARCH.

model the top- $k$ <sup>11</sup> retrieved passages and instruct the model to cite accordingly. Table 2 shows an example of our prompt. We also use in-context learning (Brown et al., 2020) and prepend 2 demonstrations with the same prompts so that the model can learn the format even if it cannot pick up the instruction. The complete instruction is in Table 22.

**SUMM/SNIPPET.** With ChatGPT’s 4,096 context window, we can at most safely fit 5 passages. As shown in Figure 4, top-5 retrieved passages can only cover 56.8% percent of the answers in ASQA.

To tackle this limitation, we propose to provide summaries or snippets of passages instead of the full text (differences are summaries are free-form but snippets are spans from passages). We acquire summaries and snippets by prompting ChatGPT with instruction (prompts in Table 24 and

<sup>11</sup> $k = 3$  for models with 2,048 context window and  $k = 5$  for models with 4,096 context window

25).<sup>12</sup> Then we replace all passages with summaries/snippets. Summaries or snippets significantly reduces the passage length: for ASQA, they reduce passage length by  $6\times$  on average.

Though SUMM/SNIPPET allows for more retrieved passages, they are lossy compressions. To alleviate this problem, we propose INTERACT, an interactive prompting scheme to allow the model to check certain passages’ full text. At each step, the model can execute one of the three actions: (1) “Check: Document [1][2]” to check the full text of the corresponding documents; (2) “Output:” to output a statement of the answer; (3) “End.” to end the generation. The full passage will disappear after one action to save context space. Table 26 shows an example of this process (details in §B).

**INLINESEARCH.** The above methods all demonstrate retrieval results at the beginning. In INLINESEARCH, we allow LLMs to call “search” during the generation process (Yao et al., 2023; Press et al., 2022; Jiang et al., 2023). At each step, the model can execute one of the three actions: “Search: {query}” to search among the top-100 passages<sup>13</sup> by using GTR; the “Output” and “End” actions are the same as INTERACT. For each “Search” action, we display the best retrieved passage in the context. Similar to INTERACT, the passage is removed after one action to save context space. Table 3 shows an example of INLINESEARCH.

**CLOSEDBOOK.** We also add a simple closed-book baseline, where the model is only prompted with the instruction and the question, without any retrieved passages provided. Consequently, this variant does not cite any evidences.

### 4.3 Post-editing

In this section we discuss several strategies for refining the output to further improve its quality.

**RERANK.** We randomly sample  $n$  ( $n = 4$  in our experiments) responses for each question, and select the best response using the automatic *citation recall* score. RERANK directly optimizes citation quality but at a cost of  $n$  times more compute.

**POSTCITE.** We also provide a post-hoc cite option to models (especially for CLOSEDBOOK), where

<sup>12</sup>We also query ChatGPT whether the passage is relevant to the question, and filter out passages that are “irrelevant”.

<sup>13</sup>We do not search over the entire corpus because {query} may leave out certain context in the question and search among the already-retrieved passages gives better results.

	Fluency	Correct.	Citation	
	(MAUVE)	(EM Rec.)	Rec.	Prec.
<b>ChatGPT</b>				
VANILLA (5-psg)	66.6	40.4	73.6	72.5
w/ RERANK	77.0	40.2	<b>84.8</b>	<b>81.6</b>
SUMM (10-psg)	70.0	<b>43.3</b>	68.9	61.8
w/ INTERACT	69.0	39.1	73.4	66.5
SNIPPET (10-psg)	69.8	41.4	65.3	57.4
INLINESEARCH	58.7	32.4	58.3	58.2
CLOSEDBOOK	52.7	38.3	26.7	26.7
<b>LLaMA (VANILLA, 3-psg)</b>				
LLaMA-13B	68.4	26.9	10.6	15.4
Vicuna-13B	82.6	31.9	51.1	50.1

Table 4: Experiments on ASQA. For CLOSED BOOK, we use POSTCITE to get citations.  $k$ -psg: putting top- $k$  passages from the retrieval results into the context.

	Correctness		Citation	
	Rec.-5	Prec.	Rec.	Prec.
<b>ChatGPT</b>				
VANILLA (5-psg)	20.8	20.8	20.5	20.9
w/ RERANK	22.8	21.4	21.2	21.4
SUMM (10-psg)	23.6	21.2	<b>23.6</b>	<b>25.7</b>
SNIPPET (10-psg)	24.5	21.5	22.9	24.9
w/ INTERACT	21.9	<b>23.0</b>	21.9	23.4
INLINESEARCH	17.2	20.4	14.9	14.9
CLOSEDBOOK	32.9	19.8	10.0	10.0
<b>LLaMA (VANILLA, 3-psg)</b>				
LLaMA-13B	9.7	9.1	6.7	7.1
Vicuna-13B	14.0	15.9	12.5	13.4

Table 5: Experiments on QAMPARI. “Rec.-5”: we set the recall to be 100% if the prediction includes at least 5 correct answers.

for each statement, we find the best matching passage among the top-100 retrieved passages using GTR and cite it. One significant drawback of POSTCITE is at most one citation will be added per statement, while some statements may require multiple citations to fully support the claim.

## 5 Experiments

We demonstrate results of different LLMs with strategies introduced in §4. More experiment details in §B. We report results of two model families.

**OpenAI models.** We use ChatGPT (gpt-3.5-turbo-0301) for most of our development as well as the main results. ChatGPT has a context window of 4,096, larger than any open-source LLMs.

**LLaMA and instruction-tuned LLaMA.** LLaMA (Touvron et al., 2023) is so far the strongest open-source LLM family and we conduct experiments with LLaMA 6.7B, 13B, and 33B.

	Fluency	Correct.	Citation	
	(MAUVE)	(Claim)	Rec.	Prec.
<b>ChatGPT</b>				
VANILLA (5-psg)	57.2	12.0	51.1	50.0
w/ RERANK	56.1	11.4	<b>69.3</b>	<b>67.8</b>
SUMM (10-psg)	40.3	12.5	51.5	48.2
SNIPPET (10-psg)	62.9	14.3	50.4	45.0
w/ INTERACT	68.0	13.3	47.8	45.0
INLINESEARCH	49.7	13.4	45.6	43.7
CLOSEDBOOK	32.6	<b>18.6</b>	15.5	15.5
<b>LLaMA (VANILLA, 3-psg)</b>				
LLaMA-13B	50.0	3.9	3.1	5.3
Vicuna-13B	58.2	10.0	15.6	19.6

Table 6: Experiments on ELI5. We use *claim recall* for the correctness evaluation.

We also explore several human-instruction-tuned LLaMA models, including Alpaca-7B (Taori et al., 2023), Vicuna-13B (Chiang et al., 2023), and Oasst-33B (Köpf et al., 2023). One limitation of LLaMA compared to ChatGPT is its smaller context window (2,048 tokens). We use a light instructions for LLaMA (Table 23) to save space.

### 5.1 Main Results

We present the main results on three datasets in Table 4, 5, and 6 respectively (full results in §F.4). We first note that all models achieve good fluency scores (except CLOSED BOOK on ELI5 due to its longer generations). We summarize main takeaways from experiments below.

**VANILLA achieves strong performance.** Despite its simplicity, ChatGPT with VANILLA achieves close-to-best performance in both correctness and citation quality. This suggests that simply putting retrieved passages in context induces a strong baseline compared to other advanced approaches.

**Using summaries or snippets improves correctness.** We see a universal trend that SUMM or SNIPPET improves the correctness (SUMM works better for ASQA and SNIPPET works better for ELI5 and QAMPARI). Though on ASQA and ELI5, such improvement is at the cost of citation quality due to the lossy compression. Combining INTERACT with SUMM/SNIPPET also does not bring improvement, and we hypothesize that checking the full passages offers limited benefit and current LLMs are not proficient in an interactive usage.

**Retrieving text on-the-fly does not improve performance.** All datasets show that VANILLA outperforms INLINESEARCH on citation quality (and on

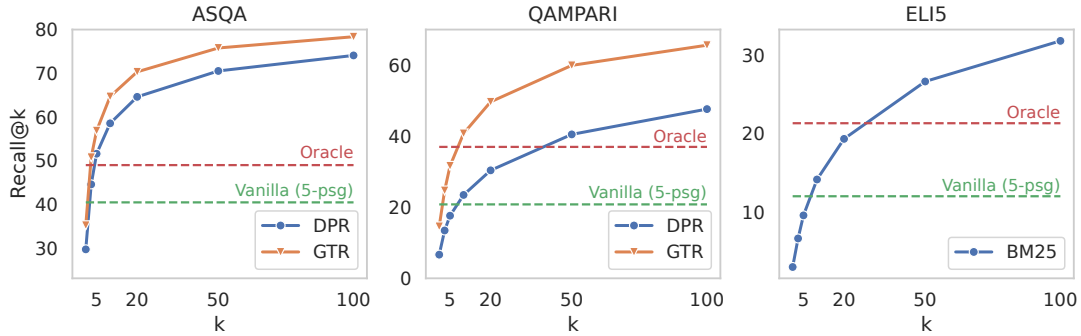


Figure 4: Retrieval recall@ $k$  on ASQA (*EM recall*), QAMPARI (*recall-5*), and ELI5 (*claim recall*). Retrieval recall serves as an upper bound for model performance, and we compare them with two models’ correctness results in the figure (dashed lines): “Vanilla (5-psg)” is ChatGPT VANILLA with top-5 passages in context; “Oracle” is the same model except that it uses 5 gold passages (§F.1), whose recall matches Recall@100 on all three datasets.

correctness for ASQA and ELI5). By manually examining the examples, we find that it is challenging to ask questions without seeing any passages: for ASQA, INLINESEARCH mostly copies or rephrases the original question as the query and it calls the retriever 1.04 times per question on average; for ELI5, given the questions are longer and provide more information, INLINESEARCH calls retrievers more (2.58 times on average) and the corresponding results are better. To improve INLINESEARCH, one may need to provide more context about the questions in advance or encourage the model to call retrievers with more detailed and diverse queries.

**RERANK boosts citation quality.** We observe that RERANK leads to constantly improvement in citation quality (on ASQA and ELI5). As the automatic scores may be biased in RERANK, we also conduct human evaluation (§6) and verify its efficacy. The limitation is that RERANK does not improve the correctness and costs  $n$  times more compute.

**CLOSEDBOOK+POSTCITE delivers strong correctness but poor citation quality.** State-of-the-art LLMs like ChatGPT can achieve strong correctness performance without any retrieval results. CLOSEDBOOK outperforms VANILLA in correctness on ELI5 and QAMPARI, and has only 2% gap on ASQA. However, CLOSEDBOOK cannot provide any citation; when combined with POSTCITE, the citation quality remains inadequate. For instance, on ASQA, VANILLA has a citation recall of 73.6%, whereas CLOSEDBOOK+POSTCITE has a significant lower citation recall of 26.7%.

To understand why CLOSEDBOOK achieves better correctness and why POSTCITE cannot deliver satisfying citation quality, we manually examine

	Fluency	Correct.	Citation	
	(MAUVE)	(EM Rec.)	Rec.	Prec.
<b>Open-source</b> (max #tokens=2,048, 3-psg, light inst.)				
LLaMA-7B	69.8	22.6	6.2	9.2
LLaMA-13B	68.4	26.9	10.6	15.4
LLaMA-33B	83.7	31.0	19.5	23.0
Alpaca-7B	84.2	32.1	12.3	14.1
Vicuna-13B	82.6	31.9	51.1	50.1
Oasst-33B	82.9	34.8	36.2	35.5
<b>ChatGPT</b> (max #tokens=4,096, full inst.)				
ChatGPT (3-psg)	66.6	39.6	72.8	73.9
ChatGPT (5-psg)	66.6	40.4	73.6	72.5

Table 7: Comparison of different LLMs on ASQA (GTR+VANILLA). All open-source models have a context limit of 2,048 tokens, and thus can only use a light version of instructions and at most top-3 passages.

model outputs and find that: (1) open-book models are easily distracted by irrelevant passages and generate responses with lower correctness, a phenomenon also observed by Shi et al. (2023). (2) CLOSEDBOOK often generates texts that are correct but not similar to any existing passages, making it difficult to post-hoc retrieve a reference.

## 5.2 Comparison of Different LLMs

Table 7 compares different LLMs on ASQA using VANILLA (more results in §F.4). Notably, human-instruction-tuned models (Alpaca-7B, Vicuna-13B, and Oasst-33B) outperform the original LLaMA models in correctness and drastically enhance the citation quality. We observe that while the original LLaMA models are able to copy facts from the context, they struggle with accurately citing the sources or simply do not cite.

The three instruction-tuned LLaMA models exhibit comparable correctness while demonstrating



	Fluency	Correct.	Citation	
	(MAUVE)	(EM)	Rec.	Prec.
<b>ChatGPT (VANILLA)</b>				
DPR (5-psg)	61.8	36.1	65.0	65.6
GTR (1-psg)	69.5	38.4	56.0	64.0
GTR (3-psg)	66.6	39.6	72.8	<b>73.9</b>
GTR (5-psg)	66.6	<b>40.4</b>	<b>73.6</b>	72.5

Table 8: Retriever ablation on ASQA.

notable variations in citation quality. By analyzing their training data, we find that data of Vicuna and Oasst contain examples similar to our citing format, while Alpaca’s data do not have such instances. We hypothesize that ChatGPT’s training data even utilize more of such examples.

### 5.3 Retrieval Analysis

The retrieval results play a crucial role to the correctness and the citation quality. Figure 4 presents the retrieval recall@ $k$  with different datasets and retrievers. As the number of passages increases, retrieval recall steadily improves. Additionally, Figure 4 shows the correctness performance of two models: (1) ChatGPT VANILLA with top-5 passages (our primary baseline); (2) an oracle version of the same model employing 5 gold passages (§F.1; the 5 gold passages match the retrieval recall@100). Notably, both models’ correctness lags behind the corresponding retrieval recall (except for ELI5 top-5). The discrepancy suggests that despite the presence of accurate answers in context, LLMs struggle to generate them in the answers.

We compare the impact of different retrievers and different numbers of passages to LLMs. Table 8 shows that GTR outperforms DPR in both correctness and citation quality, emphasizing the importance of deploying better retrievers. Contrary to the retrieval recall trend in Figure 4, more passages in context do not yield substantial improvement. Specifically, correctness plateaus at top-1 passage and citation quality plateaus at top-3 passages. This further indicates the limited ability of LLMs in utilizing multiple passages within context.

### 5.4 Other Ablations

We provide additional ablations in §F. In summary, we find that (1) using comprehensive instructions significantly enhances the citation quality of instruction-tuned models (§F.2); (2) including at least one demonstration improves the performance

	Human		ALCE	
	Rec.	Prec.	Rec.	Prec.
ChatGPT VANILLA	74.7	76.6	75.3	74.4
w/ RERANK	<b>79.3</b>	<b>81.9</b>	<b>83.9</b>	<b>80.8</b>
Vicuna-13B VANILLA	51.6	51.5	50.3	50.1

Table 9: Comparing human vs. ALCE’s citation quality evaluation on ASQA.

	Human scores		ALCE scores	
	Rec.	Prec.	Rec.	Prec.
ChatGPT VANILLA	50.8	52.4	52.8	50.4
w/ RERANK	<b>59.7</b>	<b>60.6</b>	<b>63.0</b>	<b>60.6</b>
Vicuna-13B VANILLA	13.4	19.2	13.6	18.1

Table 10: Human citation quality evaluation vs ALCE citation quality evaluation on ELI5.

of instruction-tuned models, while non-instruction-tuned models benefit from more demonstrations (§F.3). These results further underscore the importance of developing longer-context LLMs.

## 6 Human Evaluation

To verify that our automatic evaluation correlates with human judgement, we conduct human evaluation via Surge AI<sup>14</sup>. We randomly sample 100 examples from ASQA and ELI5 and annotate outputs of selected models: ChatGPT VANILLA, ChatGPT RERANK, and Vicuna-13B VANILLA.

We request workers to judge model generations on three dimensions similar to Liu et al. (2023)—(1) utility: a 1 to 5 score indicating whether the answers are helpful to the questions<sup>15</sup>; (2) citation recall: the annotator is given a sentence and all passages that the sentence cited, and is asked to judge whether the passages fully support the sentence; (3) citation precision: given a sentence and one of its citations, the annotator is asked to judge whether the citation “fully support”, “partially support”, or “does not support” the sentence. Each citation gets a precision score 1 if the output sentence has a citation recall of 1 and this citation is at least “partially support”. See Appendix E for more details. The results, along with our automated metrics, are presented in Table 9 and Table 10.

**Model outputs score high utility.** The utility scores do not differ significantly between models,

<sup>14</sup><https://www.surgehq.ai/>

<sup>15</sup>note that this is a measure of fluency and relevance to the question as opposed to correctness

ranging 3.7-3.9 for ASQA and 3.5-3.6 for ELI5. Upon inspection, all tested models are mostly able to output fluent answers that are related to the question, despite the differences in factual correctness.

**Our automatic evaluation of citation quality strongly correlates with human judgements.** As shown in Table 9 and Table 10, the relative rankings induced by human and our automatic metrics are consistent. The absolute citation scores from human and ALCE are very close except for RERANK (which uses the automated citation recall for reranking). This suggests that an improvement on ALCE citation metrics translates to improvements in human preferences. Furthermore, the Cohen’s kappa coefficient between human and ALCE suggests substantial agreement for citation recall (0.698) and moderate agreement for citation precision (0.525).

We further evaluate the accuracy of our automatic metrics by treating the human annotations as gold labels. For citation recall, ALCE achieves an accuracy of 85.1%; for citation precision, ALCE has an accuracy of 77.6%. Regarding detecting insufficient citations, ALCE has a recall of 82.3% and a precision of 84.2%; regarding detecting “irrelevant” citations, ALCE has a recall of 75.6% and a precision of 66.1%—ALCE is effective in detecting “irrelevant” citations, but due to the limitation of the NLI model (cannot detect “partial support”), it has a relatively high false positive rate.

## 7 Related Work

**Evaluating citations.** Generating text with citations is closely related to attribution. Rashkin et al. (2021) define the “attributable to identified sources” (AIS) score to measure how faithful a generated text is to its sources. Bohnet et al. (2022) apply AIS scores on a single-document short-answer question answering dataset. Honovich et al. (2022); Yue et al. (2023) study building automatic evaluations for the AIS score. A concurrent work (Liu et al., 2023) conduct human evaluation on commercial generative search engines like Bing Chat to examine their citation qualities.

Scientific citation text generation (Funkquist et al., 2022) is a related task to ALCE where the model is provided the papers-to-cite and context and is required to recover the citing text. It is different from ALCE as all citations are provided and the model only needs to perform the summarization.

**Retrieval-augmented LLMs.** Many studies have

explored augmenting LMs with externally retrieved information. Guu et al. (2020); Borgeaud et al. (2022); Izacard et al. (2022) pre-train language models with retrieved passages, while Khandelwal et al. (2020); Zhong et al. (2022) augment LLMs’ output by interpolating it with a  $k$ NN module; though none of them explicitly provide citations to the retrieved sources. Other works prompt or fine-tune LLMs to “retrieve on-the-fly” (Parisi et al., 2022; Schick et al., 2023; Shuster et al., 2022; Jiang et al., 2023; Yao et al., 2023; Press et al., 2022), which offers flexibility of when and what to search. Gao et al. (2022); He et al. (2022) propose to first generate text without accessing external documents and then retrieve relevant documents and revise the generation to be consistent.

Among previous explorations, Nakano et al. (2021); Menick et al. (2022) are the closest to our setting, where LLMs are trained to answer questions while providing citations. However, they do not explore retrieval strategies and simply use commercial search engines, which are not reproducible, and their models and training data are closed-source. To the best of our knowledge, we are the first to implement end-to-end systems that retrieve, synthesize, and cite documents with LLMs.

## 8 Conclusion

We propose ALCE, the first automatic benchmark for evaluating LLM generations with citations. We deploy automatic metrics to measure fluency, correctness, and citation quality, and verify their efficacy via human evaluation. We explore a variety of strategies for incorporating citations in LLMs and demonstrate that current systems have considerable room for improvement on ALCE.

Our experiments highlight a number of promising research directions, including (1) enhancing retrieval and refining retrieval integrations in LLMs, (2) developing long-context LLMs, and (3) advancing LLMs’ ability to synthesize multiple sources. What’s even more intriguing is that these research proposals extend beyond the ALCE setup (for example, long-context LLMs have numerous exciting applications), and ALCE can serve as a valuable testbed for their development.

## Limitation

Our evaluation is not without flaw: (1) MAUVE is found to be sensitive to output length and may provide unstable results. (2) For the ELI5’s cor-

rectness evaluation, the automatically generated claims may not cover all possible answers due to the open-ended nature of the questions. (3) Our citation quality evaluation is limited by the accuracy of the NLI model; For citation precision, the NLI model cannot detect the case of “partially support” and thus leads to a lower citation precision score than the human evaluation.

Although we believe our curated datasets closely resemble the distribution of real-world user questions, we acknowledge that they do not cover more challenging scenarios, such as multi-hop reasoning, math reasoning, and code completion.

In our experiments, we focus on prompting LLMs without updating their model weights. Training a model directly to incorporate citations remains challenging due to the lack of supervised data. However, we observe that certain human-instruction datasets contain examples similar to our task setup. We leave the exploration of training LLMs to generate citations for future work.

## Acknowledgments

We appreciate the helpful feedback from the members of the Princeton NLP group. We thank Nelson Liu, Tianyi Zhang, and Sadhika Malladi for the valuable discussion. We thank Surge AI (especially Anna Folinsky and Edwin Chen) for their support with the human evaluation. This research is supported by an NSF CAREER award (IIS-2239290), a Sloan Research Fellowship, and Microsoft Azure credits through the “Accelerate Foundation Models Academic Research” Initiative.

## References

- Bernd Bohnet, Vinh Q Tran, Pat Verga, Roei Aharoni, Daniel Andor, Livio Baldini Soares, Jacob Eisenstein, Kuzman Ganchev, Jonathan Herzig, Kai Hui, et al. 2022. [Attributed question answering: Evaluation and modeling for attributed large language models](#). *arXiv preprint arXiv:2212.08037*.
- Sebastian Borgeaud, Arthur Mensch, Jordan Hoffmann, Trevor Cai, Eliza Rutherford, Katie Millican, George Bm Van Den Driessche, Jean-Baptiste Lespiau, Bogdan Damoc, Aidan Clark, Diego De Las Casas, Aurelia Guy, Jacob Menick, Roman Ring, Tom Hennigan, Saffron Huang, Loren Maggiore, Chris Jones, Albin Cassirer, Andy Brock, Michela Paganini, Geoffrey Irving, Oriol Vinyals, Simon Osindero, Karen Simonyan, Jack Rae, Erich Elsen, and Laurent Sifre. 2022. [Improving language models by retrieving from trillions of tokens](#). In *International Conference on Machine Learning (ICML)*, volume 162, pages 2206–2240.
- Samuel Bowman, Gabor Angeli, Christopher Potts, and Christopher D Manning. 2015. A large annotated corpus for learning natural language inference. In *Empirical Methods in Natural Language Processing (EMNLP)*.
- Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica, and Eric P. Xing. 2023. [Vicuna: An open-source chatbot impressing gpt-4 with 90%\\* chatgpt quality](#).
- 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.
- Martin Funkquist, Ilia Kuznetsov, Yufang Hou, and Iryna Gurevych. 2022. [Citebench: A benchmark for scientific citation text generation](#). *arXiv preprint arXiv:2212.09577*.
- Luyu Gao, Zhuyun Dai, Panupong Pasupat, Anthony Chen, Arun Tejasvi Chaganty, Yicheng Fan, Vincent Y Zhao, Ni Lao, Hongrae Lee, Da-Cheng Juan, et al. 2022. [Rarr: Researching and revising what language models say, using language models](#). *arXiv preprint arXiv:2210.08726*.
- Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Ming-Wei Chang. 2020. Realm: Retrieval-augmented language model pre-training. In *International Conference on Machine Learning (ICML)*, ICML’20.
- Hangfeng He, Hongming Zhang, and Dan Roth. 2022. [Rethinking with retrieval: Faithful large language model inference](#). *arXiv preprint arXiv:2301.00303*.
- Or Honovich, Roei Aharoni, Jonathan Herzig, Hagai Taitelbaum, Doron Kukliansky, 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.
- 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*.

- Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. 2023. [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, Yiming Yang, Jamie Callan, and Graham Neubig. 2023. [Active retrieval augmented generation](#). *arXiv preprint arXiv:2305.06983*.
- Ryo Kamoi, Tanya Goyal, Juan Diego Rodriguez, and Greg Durrett. 2023. [Wice: Real-world entailment for claims in wikipedia](#). *arXiv preprint arXiv:2303.01432*.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. [Dense passage retrieval for open-domain question answering](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6769–6781, Online. Association for Computational Linguistics.
- 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 (ICLR)*.
- Tushar Khot, Ashish Sabharwal, and Peter Clark. 2018. [Scitail: A textual entailment dataset from science question answering](#). In *Conference on Artificial Intelligence (AAAI)*, volume 32.
- Kalpesh Krishna, Aurko Roy, and Mohit Iyyer. 2021. [Hurdles to progress in long-form question answering](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4940–4957, Online. Association for Computational Linguistics.
- Andreas Köpf, Yannic Kilcher, Dimitri von Rütte, Sotiris Anagnostidis, Zhi-Rui Tam, Keith Stevens, Abdullah Barhoum, Nguyen Minh Duc, Oliver Stanley, Richárd Nagyfi, Shahul ES, Sameer Suri, David Glushkov, Arnav Dantuluri, Andrew Maguire, Christoph Schuhmann, Huu Nguyen, and Alexander Mattick. 2023. [Openassistant conversations – democratizing large language model alignment](#). *arXiv preprint arXiv:2304.07327*.
- Kenton Lee, Ming-Wei Chang, and Kristina Toutanova. 2019. [Latent retrieval for weakly supervised open domain question answering](#). In *Association for Computational Linguistics (ACL)*, pages 6086–6096.
- Nelson F Liu, Tianyi Zhang, and Percy Liang. 2023. [Evaluating verifiability in generative search engines](#). *arXiv preprint arXiv:2304.09848*.
- Jacob Menick, Maja Trebacz, Vladimir Mikulik, John Aslanides, Francis Song, Martin Chadwick, Mia Glaese, Susannah Young, Lucy Campbell-Gillingham, Geoffrey Irving, et al. 2022. [Teaching language models to support answers with verified quotes](#). *arXiv preprint arXiv:2203.11147*.
- Sewon Min, Julian Michael, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2020. [AmbigQA: Answering ambiguous open-domain questions](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5783–5797, Online. Association for Computational Linguistics.
- Reiichiro Nakano, Jacob Hilton, Suchir Balaji, Jeff Wu, Long Ouyang, Christina Kim, Christopher Hesse, Shantanu Jain, Vineet Kosaraju, William Saunders, et al. 2021. [Webgpt: Browser-assisted question-answering with human feedback](#). *arXiv preprint arXiv:2112.09332*.
- 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 *Empirical Methods in Natural Language Processing (EMNLP)*, pages 9844–9855.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. [Training language models to follow instructions with human feedback](#). *Advances in Neural Information Processing Systems (NeurIPS)*, 35:27730–27744.
- Aaron Parisi, Yao Zhao, and Noah Fiedel. 2022. [Talm: Tool augmented language models](#). *arXiv preprint arXiv:2205.12255*.
- Fabio Petroni, Aleksandra Piktus, Angela Fan, Patrick Lewis, Majid Yazdani, Nicola De Cao, James Thorne, Yacine Jernite, Vladimir Karpukhin, Jean Maillard, Vassilis Plachouras, Tim Rocktäschel, and Sebastian Riedel. 2021. [KILT: a benchmark for knowledge intensive language tasks](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2523–2544, Online. Association for Computational Linguistics.
- Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Dmytro Okhonko, Samuel Broscheit, Gautier Izacard, Patrick Lewis, Barlas Oğuz, Edouard Grave, Wen-tau Yih, et al. 2021. [The web is your oyster—knowledge-intensive nlp against a very large web corpus](#). *arXiv preprint arXiv:2112.09924*.
- 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*.



- Ofir Press, Muru Zhang, Sewon Min, Ludwig Schmidt, Noah A Smith, and Mike Lewis. 2022. [Measuring and narrowing the compositionality gap in language models](#). *arXiv preprint arXiv:2210.03350*.
- 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).
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [SQuAD: 100,000+ questions for machine comprehension of text](#). In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 2383–2392.
- Hannah Rashkin, Vitaly Nikolaev, Matthew Lamm, Lora Aroyo, Michael Collins, Dipanjan Das, Slav Petrov, Gaurav Singh Tomar, Iulia Turc, and David Reitter. 2021. [Measuring attribution in natural language generation models](#). *arXiv preprint arXiv:2112.12870*.
- 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*.
- 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*.
- Tal Schuster, Adam Fisch, and Regina Barzilay. 2021. [Get your vitamin C! robust fact verification with contrastive evidence](#). In *North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, pages 624–643.
- Freda Shi, Xinyun Chen, Kanishka Misra, Nathan Scales, David Dohan, Ed Chi, Nathanael Schärli, and Denny Zhou. 2023. [Large language models can be easily distracted by irrelevant context](#). In *International Conference on Machine Learning (ICML)*.
- Kurt Shuster, Mojtaba Komeili, Leonard Adolphs, Stephen Roller, Arthur Szlam, and Jason Weston. 2022. [Language models that seek for knowledge: Modular search & generation for dialogue and prompt completion](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 373–393.
- Ivan Stelmakh, Yi Luan, Bhuwan Dhingra, and Ming-Wei Chang. 2022. [Asqa: Factoid questions meet long-form answers](#). *arXiv preprint arXiv:2204.06092*.
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. 2023. [Stanford alpaca: An instruction-following llama model](#).
- James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2018. [FEVER: a large-scale dataset for fact extraction and VERification](#). In *North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, pages 809–819.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. [Llama: Open and efficient foundation language models](#). *arXiv preprint arXiv:2302.13971*.
- Alex Wang, Kyunghyun Cho, and Mike Lewis. 2020. [Asking and answering questions to evaluate the factual consistency of summaries](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5008–5020, Online. Association for Computational Linguistics.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. [A broad-coverage challenge corpus for sentence understanding through inference](#). In *North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*.
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik R Narasimhan, and Yuan Cao. 2023. [React: Synergizing reasoning and acting in language models](#). In *International Conference on Learning Representations (ICLR)*.
- Xiang Yue, Boshi Wang, Kai Zhang, Zirui Chen, Yu Su, and Huan Sun. 2023. [Automatic evaluation of attribution by large language models](#). *arXiv preprint arXiv:2305.06311*.
- Shiyue Zhang and Mohit Bansal. 2021. [Finding a balanced degree of automation for summary evaluation](#). In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 6617–6632.
- Yuan Zhang, Jason Baldridge, and Luheng He. 2019. [PAWS: Paraphrase adversaries from word scrambling](#). In *North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, pages 1298–1308.
- Zexuan Zhong, Tao Lei, and Danqi Chen. 2022. [Training language models with memory augmentation](#). In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 5657–5673.

## A Generating Claims for ELI5

We elect not to use ROUGE-L as our main correctness metrics since it does not account for the different ways of expressing the same answer and it can be easily gamed (Krishna et al., 2021). We further illustrate this issue in Table 11. A system can easily achieve high ROUGE-L score by retrieving and returning the top passage from a BM25 index. However, the claims evaluation metric does not reward this approach since the output often lacks different aspects of the answers.

	ROUGE-L	Claim recall
ChatGPT VANILLA	20.6	12.0
ChatGPT ORACLE	21.2	21.3
LLaMa-13B VANILLA	16.2	3.9
Top-1 passage	19.1	3.0

Table 11: Comparison between ROUGE-L and claim recall scores on ELI5.

Instead, we leverage the original answers to generate sub-claims and use them to serve as an estimate of the different aspects of the answers that we expect the model to cover. This approach is inspired by works in summarization evaluation and claim verification (Zhang and Bansal, 2021; Kamoi et al., 2023; Wang et al., 2020).

Specifically, we use text-davinci-003 to generate the sub-claims. We first manually annotate three question and answer pairs from the original ELI5 training set with 3 sub-claims each. Then, we prompt text-davinci-003 with these pairs as demonstrations. The full prompt with an example is shown in Table 21.

**InstructGPT generates coherent and faithful sub-claims.** To ensure that the generated sub-claims are of good quality, we manually inspect a random sample of 40 answers and their generated sub-claims (totally 120 sub-claims). For each sub-claim, we assign a score of 1 if it is relevant to the question and faithful to the facts presented in the ground truth, and 0 otherwise. We found that 112 out of the 120 (93.33%) sub-claims received a score of 1, meaning that our generated sub-claims are of high quality and faithful to the ground truth. Furthermore, the average number of words in the generated sub-claims is 14 words, and they are typically just one sentence long. This is aligned with the intent behind the metric—to capture short factual claims made by the original answer.

**NLI model accurately predicts the entailment of sub-claims.** We further analyze our sub-claim evaluation metrics by checking the error rate of the final prediction of the NLI model. To this end, we first manually annotate the entailment scores between 40 outputs and their sub-claims (totally 120 pairs; these are the same questions from the previous analysis). We then use the NLI model to obtain the entailment scores for the output and sub-claims. Using the human annotations as the ground truth label, we found that the NLI model achieved an accuracy of 80.0%.

## B Implementation Details

**NLI model.** We use the version of TRUE model from [https://huggingface.co/google/t5-xxl\\_true\\_nli\\_mixture](https://huggingface.co/google/t5-xxl_true_nli_mixture), which is trained on SNLI (Bowman et al., 2015), MNLI (Williams et al., 2018), Fever (Thorne et al., 2018), SciTail (Khot et al., 2018), PAWS (Zhang et al., 2019), and VitaminC (Schuster et al., 2021). This model uses the following prompt: “premise: {PREMISE} hypothesis: { }” and outputs “1” if the premise entails the hypothesis. We format each passage (when used as premise) by the format of “Title: {TITLE}\n{TEXT}” and concatenate all passages with “\n” as a separator.

**MAUVE.** When running MAUVE, we concatenate the question and the model output (or human answer) by space. We truncate both the references and the model generations to 100 words, as we found MAUVE results are unstable beyond this length for ELI5 (this is due to that ELI5 has a lot of extremely long human answers).

**Output truncation.** Before evaluation, we truncate model output by new lines, as non-instruction-tuned models may generate more content after new lines that are irrelevant.

**INTERACT.** Empirically we found that models tend to execute too many consecutive “check” actions, so we force the model to always “output” after each “check”. We also limit the maximum number of passages to check as 3 to avoid exceeding the length limit.

**Main experiments.** For all experiments except ChatGPT RERANK, we run each model three times with different seeds and each time we sample two demonstrations from a pool of four. We report averaged scores for all experiments in the main paper

and we report the standard deviations in Appendix F.4.

## C ALCE Catches Shortcut Cases

	Fluency	Correct.	Citation	
	(MAUVE)	(EM Rec.)	Rec.	Prec.
ChatGPT	66.6	40.4	73.6	63.0
Top-1 passage	20.8	35.1	99.4	99.4
First 2 sents	67.2	18.9	98.7	98.7

Table 12: ASQA cheating cases. “ChatGPT”: the ChatGPT VANILLA model with GTR-retrieved top-5 passages. “Top-1 passage”: use the top-1 retrieved passage as the response. “First 2 sents”: use the first 2 sentences of the top-1 retrieved passage.

Table 12 demonstrates the experiments to show that ALCE is robust to shortcut cases. Using the top-1 passages or first two sentences of the top-1 passages induces almost perfect citation quality, but fluency and correctness are dramatically lower.

## D Citation Recall Discussion

Liu et al. (2023) conduct human evaluation on citation precision in a different way: For each citation, they ask annotators to judge whether the citation (1) fully support, (2) partially support, or (3) does not support  $s_i$ . One citation  $c_{i,j}$  is precise if (a)  $c_{i,j}$  fully supports  $s_i$  or (b)  $\mathcal{C}_i$  fully supports  $s_i$ ,  $c_{i,j}$  partially supports  $s_i$ , and no  $c \in \mathcal{C}_i$  alone fully supports  $s_i$ . This evaluation solved the corner case we mentioned in the main paper (one citation partially supports the claim but is identified as “irrelevant”). However, it is challenging to conduct such evaluation automatically, as there is no existing model that can judge whether a citation “partially” supports a claim. We also explore prompting ChatGPT to conduct such a task, which yields poor results. We defer it to future work to collect such supervised data to train a better  $\phi$  that is equipped with the ability to detect “partial support”.

## E Human Evaluation

We employ Surge AI (<https://www.surgehq.ai/>) for our human evaluation. The average pay to workers is 20 USD per hour.

### E.1 Utility

To check if the model output is useful to downstream users, we measure the utility of the response  $S$ . We first show the query  $q$  and model response

	R@1	R@3	R@5	R@20	R@100
DPR	29.6	44.5	51.5	64.6	74.1
GTR	35.1	50.7	56.8	70.3	78.4
Oracle	63.8	72.8	78.4	-	-

Table 13: Retrieval results for ASQA (EM recall).

	R@1	R@3	R@5	R@20	R@100
DPR	6.7	13.5	17.6	30.4	47.6
GTR	14.6	24.7	31.6	49.7	65.6
Oracle	44.3	58.7	65.6	-	-

Table 14: Retrieval results for QAMPARI (recall-5).

$S$  to the worker and ask them to rate their agreement with the statement "The response is a helpful and informative answer to the query" on a Likert scale of 1-5, corresponding to *Strongly Disagree*, *Disagree*, *Neutral*, *Agree*, and *Strongly Agree*.

## E.2 Citation Recall

The annotators are shown the question  $q$ , the statement  $s_i$ , and all of its citations  $\mathcal{C}_i$ , and they rate if the joint set of citations fully support the statement (recall=1) or if they do not support all the claims (recall=0). We calculate the overall recall score for the generation by taking an average of all the statements’ recall scores.

## E.3 Citation Precision

We show the question  $q$  and a pair of a statement  $s_i$  and one of its citation  $c_{i,j} \in \mathcal{C}_i$  to the annotator. We ask the annotator if the citation *fully supports*, *partially supports*, or *does not support* the factual claims in  $s_i$ . Citation  $c_{i,j}$  has a citation precision of 1 if  $s_i$  has a recall of 1, and  $c_{i,j}$  fully or partially supports  $s_i$ . Finally, we take an average of precision scores of all citations in the statement  $S$  to obtain the citation precision score.

## F More Experiments

### F.1 Retrieval Analysis

**Oracle.** Since the original datasets do not contain gold passages at the same granularity level as our setting (100-word passages), we approximate gold passages by running the following algorithm on the top-100 retrieved passages: We calculate the recall score for each passage. Then, we sort the passages using their recall score and take the top 5 passages as our initial oracle set. Finally, we iterate through all passages that were not initially in the oracle set

	R@1	R@3	R@5	R@20	R@100
BM25	3.0	6.6	9.6	19.3	31.8
Oracle	25.3	29.7	31.8	-	-

Table 15: Retrieval results for ELI5 (claim recall).

and try to replace the passages in the oracle set in a greedy fashion: we calculate the change in the recall score of the oracle set for every possible replacement and proceed with the replacement that results in the largest recall improvement.

**Detailed retrieval results.** We show detailed retrieval results in Table 13, 14, and 15.

## F.2 Effect of Instructions

	<u>Fluency</u>	<u>Correct.</u>	<u>Citation</u>	
	(MAUVE)	(EM Rec.)	Rec.	Prec.
<b>ChatGPT (VANILLA, 5-doc)</b>				
Light instruction	64.1	39.5	69.6	73.2
Full instruction	66.6	40.4	73.6	72.5

Table 16: Instruction

Table 16 shows results of using a full instruction (Table 22) and a light version of the instruction (Table 23). We see that the full version induces stronger correctness and citation recall, while the two instructions lead to similar citation precision.

## F.3 Effect of Demonstrations

	<u>Fluency</u>	<u>Correct.</u>	<u>Citation</u>	
	(MAUVE)	(EM)	Rec.	Prec.
<b>ChatGPT (VANILLA)</b>				
#demo = 0	74.5	41.9	69.3	73.4
#demo = 1	68.9	39.8	74.6	73.2
#demo = 2	66.6	40.4	73.6	72.5

Table 17: Different demonstrations on ASQA.

Table 17 shows results on effect of different numbers of demonstrations. We see that numbers of demonstrations do not affect ChatGPT’s correctness but using at least one demonstration ensures high citation recall. For the original LLaMA model, Table 17 shows the trend that more demonstrations lead to better performance.

## F.4 Main Results

We show full results of our experiments along with the standard deviation in Table 18, 19, and 20. We

repeat all experiments with three different random seeds. However, for ChatGPT RERANK, we use only one seeded run since each run repeats the generation step four times, and more experiments would incur significant costs.

## G Prompts

We show detailed prompts used in our paper in Table 22, 23, 24, 25, 26, 27, and 28.

## H Examples

In Table 29 and 30 we show some examples of questions and model generated outputs.



	Fluency	Correct.	Citation			
	(MAUVE)	(EM Rec.)	Rec.	Prec.	ROUGE-L	Length
<b>ChatGPT</b>						
VANILLA (5-psg)	66.8 (2.0)	40.4 (0.6)	73.6 (1.1)	72.5 (1.8)	37.0 (0.4)	40.0 (3.1)
w/ RERANK	77.0 (0.0)	40.2 (0.0)	84.8 (0.0)	81.6 (0.0)	36.9 (0.0)	40.8 (0.0)
SUMM (10-psg)	70.0 (1.2)	43.3 (0.8)	68.8 (0.6)	61.8 (1.1)	36.9 (0.2)	49.8 (4.3)
w/ INTERACT	69.0 (2.7)	39.1 (0.5)	73.4 (0.2)	66.5 (4.9)	35.7 (0.2)	34.0 (0.9)
SNIPPET (10-psg)	69.8 (2.5)	41.4 (0.6)	65.3 (0.6)	57.4 (0.9)	36.4 (0.4)	43.0 (3.5)
INLINESEARCH	58.7 (1.3)	32.4 (0.6)	58.3 (1.3)	58.3 (1.3)	58.2 (1.1)	23.7 (1.1)
CLOSEDBOOK	52.7 (4.9)	38.2 (0.1)	26.7 (1.1)	26.7 (1.1)	37.1 (0.3)	61.1 (4.5)
ORACLE(5-psg)	64.4 (0.6)	48.9 (1.2)	74.5 (0.6)	72.7 (1.0)	38.2 (1.0)	37.4 (3.0)
<b>Open-source</b>						
LLaMA-7B VANILLA (3-psg)	69.8 (2.0)	22.6 (0.9)	6.2 (2.7)	9.2 (2.9)	29.1 (0.2)	61.3 (14.3)
Alpaca-7B VANILLA (3-psg)	84.2 (2.7)	32.1 (1.7)	12.3 (7.2)	14.1 (7.0)	33.1 (0.8)	51.7 (12.8)
Vicuna-7B VANILLA (3-psg)	82.9 (5.0)	34.6 (0.7)	40.3 (0.5)	42.6 (1.0)	35.9 (0.7)	48.9 (6.6)
LLaMA-13B VANILLA (3-psg)	68.4 (6.4)	26.9 (0.4)	10.6 (4.7)	15.4 (5.2)	29.8 (0.5)	67.1 (19.1)
w/ RERANK	60.9 (14.5)	25.2 (2.5)	28.1 (9.3)	37.0 (7.2)	27.9 (2.4)	50.5 (14.3)
LLaMA-13B SUMM (10-psg)	76.8 (4.7)	33.3 (0.7)	19.6 (3.9)	23.7 (4.7)	32.1 (0.3)	54.4 (1.5)
LLaMA-13B SNIPPET (10-psg)	72.0 (0.8)	31.3 (1.1)	18.2 (3.1)	21.1 (3.6)	30.8 (0.4)	50.5 (4.5)
LLaMA-13B ORACLE (3-psg)	69.5 (11.4)	34.3 (0.9)	10.8 (4.9)	15.8 (5.9)	30.6 (0.1)	67.3 (17.9)
Vicuna-13B VANILLA (3-psg)	82.6 (9.4)	31.9 (3.9)	51.1 (1.4)	50.1 (2.5)	34.9 (1.3)	39.1 (6.6)
w/ RERANK	73.5 (2.1)	32.9 (1.3)	71.9 (1.9)	65.4 (1.5)	34.6 (0.3)	35.7 (4.2)
Vicuna-13B SUMM (10-psg)	67.7 (0.3)	43.2 (0.1)	52.7 (2.6)	50.0 (2.1)	36.7 (0.2)	66.0 (1.2)
Vicuna-13B SNIPPET (10-psg)	81.4 (3.0)	42.1 (1.2)	53.4 (1.9)	48.7 (1.6)	36.9 (0.4)	61.2 (7.4)
Vicuna-13B ORACLE (3-psg)	72.9 (3.5)	42.5 (1.6)	52.2 (0.8)	50.7 (1.6)	36.5 (0.9)	38.7 (3.5)
LLaMA-33B VANILLA (3-psg)	83.7 (5.4)	31.0 (0.8)	19.5 (5.3)	23.0 (5.3)	32.3 (0.6)	44.1 (9.3)
w/ RERANK	82.1 (3.0)	31.3 (1.1)	41.3 (6.4)	44.7 (5.5)	32.5 (0.9)	39.4 (8.0)
LLaMA-33B SUMM (10-psg)	72.0 (3.0)	33.1 (1.9)	34.7 (5.8)	35.2 (6.0)	31.1 (0.8)	43.7 (5.0)
LLaMA-33B SNIPPET (10-psg)	70.8 (3.1)	30.9 (1.4)	31.4 (4.2)	31.5 (5.3)	30.1 (0.7)	42.8 (3.6)
LLaMA-33B ORACLE (3-psg)	82.6 (7.1)	39.3 (2.9)	20.2 (6.2)	23.9 (6.3)	33.1 (0.9)	42.0 (9.3)
Oasst-33B VANILLA (3-psg)	82.9 (2.7)	34.8 (1.5)	36.2 (1.7)	38.3 (2.7)	35.5 (0.7)	45.2 (6.3)
w/ RERANK	83.2 (2.4)	35.1 (1.4)	66.7 (0.2)	64.3 (1.0)	35.0 (0.6)	41.8 (6.0)
Oasst-33B SUMM (10-psg)	74.3 (4.6)	40.9 (1.1)	45.5 (1.9)	44.0 (2.9)	35.8 (0.6)	54.3 (4.8)
Oasst-33B SNIPPET (10-psg)	79.3 (1.0)	40.1 (0.9)	45.0 (1.3)	43.3 (2.2)	35.8 (0.2)	50.9 (4.1)
Oasst-33B ORACLE (3-psg)	85.1 (2.8)	44.3 (2.4)	37.0 (1.0)	39.6 (1.5)	36.5 (1.1)	44.2 (5.8)

Table 18: ASQA full results.

	Correctness		Citation		Num Pred.
	Rec.-5	Prec.	Rec.	Prec.	
ChatGPT					
VANILLA (5-psg)	20.8 (2.2)	20.8 (0.2)	20.5 (0.7)	20.9 (0.7)	5.0 (0.5)
w/ RERANK	22.8 (0.0)	21.4 (0.0)	21.2 (0.0)	21.4 (0.0)	5.4 (0.0)
SUMM (10-psg)	23.6 (0.9)	21.2 (0.5)	23.6 (0.7)	25.7 (0.8)	6.7 (0.4)
SNIPPET (10-psg)	24.5 (1.4)	21.5 (1.8)	22.9 (1.6)	24.9 (0.4)	7.2 (0.9)
w/ INTERACT	21.9 (0.9)	23.0 (0.4)	21.9 (1.2)	23.4 (0.9)	6.7 (0.4)
INLINESEARCH	17.2 (1.1)	20.4 (0.8)	14.9 (0.8)	14.9 (0.8)	6.7 (0.2)
CLOSEDBOOK	32.9 (1.1)	19.8 (1.6)	10.0 (0.4)	10.0 (0.4)	17.0 (2.9)
ORACLE	37.0 (3.1)	36.9 (0.6)	24.1 (1.2)	24.6 (1.3)	5.3 (0.6)
Open-source					
LLaMA-7B VANILLA (3-psg)	7.8 (3.4)	7.4 (2.7)	5.1 (0.5)	5.7 (0.8)	5.7 (0.6)
Alpaca-7B VANILLA (3-psg)	9.4 (3.7)	9.5 (3.6)	6.4 (0.5)	6.8 (0.5)	5.1 (0.1)
Vicuna-7B VANILLA (3-psg)	11.3 (1.4)	13.3 (2.3)	10.1 (0.6)	10.9 (0.5)	3.9 (0.3)
LLaMA-13B VANILLA (3-psg)	9.7 (3.6)	9.1 (3.1)	6.7 (0.9)	7.1 (0.9)	5.9 (0.6)
w/ RERANK	10.0 (3.3)	10.7 (3.3)	9.9 (1.2)	10.2 (1.1)	5.4 (0.5)
LLaMA-13B SUMM (10-psg)	14.8 (2.5)	12.6 (1.5)	7.4 (0.5)	8.0 (0.6)	8.1 (0.9)
LLaMA-13B SNIPPET (10-psg)	17.7 (1.4)	15.7 (0.9)	8.8 (0.7)	9.9 (0.6)	8.2 (0.4)
LLaMA-13B ORACLE (3-psg)	16.8 (6.6)	15.4 (5.6)	7.7 (1.0)	8.3 (1.1)	5.7 (0.7)
Vicuna-13B VANILLA (3-psg)	14.0 (0.6)	15.9 (1.7)	12.5 (0.8)	13.4 (0.7)	4.7 (0.3)
w/ RERANK	13.0 (0.7)	17.2 (2.2)	17.3 (0.8)	17.7 (0.6)	4.4 (0.3)
Vicuna-13B SUMM (10-psg)	21.1 (1.4)	17.1 (0.3)	15.7 (0.2)	17.8 (0.1)	6.9 (0.7)
Vicuna-13B SNIPPET (10-psg)	21.9 (0.8)	18.2 (0.3)	16.8 (0.3)	19.7 (0.6)	7.5 (0.4)
Vicuna-13B ORACLE (3-psg)	25.9 (1.6)	28.4 (2.6)	15.8 (1.4)	16.8 (1.4)	4.9 (0.5)
LLaMA-33B VANILLA (3-psg)	14.7 (3.3)	12.0 (2.2)	7.9 (0.7)	8.3 (0.6)	7.2 (0.7)
w/ RERANK	14.0 (3.4)	13.9 (2.6)	10.7 (0.6)	11.1 (0.5)	6.4 (0.7)
LLaMA-33B SUMM (10-psg)	19.0 (1.9)	14.8 (0.8)	12.5 (0.2)	15.0 (0.3)	7.6 (0.6)
LLaMA-33B SNIPPET (10-psg)	19.6 (1.1)	15.7 (0.1)	12.8 (1.1)	15.2 (1.2)	7.8 (0.5)
LLaMA-33B ORACLE (3-psg)	23.9 (6.9)	20.3 (5.2)	9.8 (1.2)	10.4 (1.2)	6.8 (0.9)
Oasst-33B VANILLA (3-psg)	15.5 (1.5)	14.9 (1.4)	9.0 (1.6)	10.1 (1.8)	5.6 (0.3)
w/ RERANK	14.1 (1.1)	15.8 (1.0)	15.0 (1.6)	15.9 (1.6)	4.7 (0.3)
Oasst-33B SUMM (10-psg)	21.0 (0.6)	17.5 (1.0)	12.9 (1.2)	16.6 (1.2)	7.1 (0.4)
Oasst-33B SNIPPET (10-psg)	22.0 (0.4)	17.4 (0.3)	13.6 (1.7)	17.7 (1.6)	7.5 (0.1)
Oasst-33B ORACLE (3-psg)	26.9 (3.7)	26.0 (3.3)	11.7 (1.0)	12.9 (1.2)	5.6 (0.4)

Table 19: QAMPARI full.

	Fluency	Correct.	Citation			
	(MAUVE)	(Claim)	Rec.	Prec.	ROUGE-L	Length
<b>ChatGPT</b>						
VANILLA (5-psg)	57.2 (1.6)	12.0 (0.6)	51.1 (4.2)	50.0 (4.8)	20.6 (0.2)	91.5 (6.5)
w/ RERANK	56.1 (0.0)	11.4 (0.0)	69.3 (0.0)	67.8 (0.0)	20.3 (0.0)	103.4 (0.0)
SUMM (10-psg)	40.2 (1.2)	12.5 (0.2)	51.5 (1.1)	48.2 (2.0)	20.3 (0.1)	90.0 (6.6)
SNIPPET (10-psg)	62.9 (2.2)	14.3 (0.1)	50.4 (1.1)	45.0 (2.7)	21.0 (0.1)	100.0 (6.8)
w/ INTERACT	68.0 (5.8)	13.3 (0.5)	47.8 (3.3)	45.0 (3.1)	20.1 (0.2)	99.8 (6.1)
INLINESEARCH	49.7 (4.6)	13.4 (1.1)	45.6 (2.5)	43.7 (3.9)	20.4 (0.3)	103.0 (18.1)
CLOSEDBOOK	32.6 (1.1)	18.6 (0.5)	15.4 (0.3)	15.4 (0.3)	22.8 (0.1)	108.3 (8.9)
ORACLE (5-psg)	59.4 (4.1)	21.3 (0.2)	57.8 (3.7)	56.0 (3.8)	21.2 (0.3)	93.0 (7.8)
<b>Open-source</b>						
LLaMA-7B VANILLA (3-psg)	28.6 (17.9)	1.6 (0.9)	1.2 (0.0)	2.7 (0.1)	12.2 (1.3)	46.9 (1.2)
Alpaca-7B VANILLA (3-psg)	45.9 (5.3)	9.2 (0.1)	4.5 (1.6)	5.2 (1.9)	18.8 (0.3)	67.1 (1.2)
Vicuna-7B VANILLA (3-psg)	43.2 (3.9)	10.0 (0.5)	12.6 (2.3)	16.3 (2.6)	19.1 (0.4)	68.7 (2.0)
LLaMA-13B VANILLA (3-psg)	50.0 (2.0)	3.9 (0.4)	3.1 (0.9)	5.3 (1.3)	16.1 (0.5)	63.3 (2.0)
w/ RERANK	46.7 (2.9)	4.3 (0.4)	9.7 (2.1)	15.0 (2.2)	16.1 (0.7)	63.0 (2.3)
LLaMA-13B SUMM (10-psg)	28.6 (1.8)	2.9 (0.1)	2.5 (0.8)	3.8 (0.8)	8.5 (0.3)	33.1 (0.6)
LLaMA-13B SNIPPET (10-psg)	48.4 (3.1)	5.7 (0.9)	5.8 (0.6)	7.6 (0.9)	15.1 (1.1)	60.2 (3.2)
LLaMA-13B ORACLE (3-psg)	49.5 (2.4)	6.4 (0.6)	3.7 (0.7)	6.5 (1.0)	16.8 (0.5)	64.5 (1.7)
Vicuna-13B VANILLA (3-psg)	58.2 (25.1)	10.0 (0.3)	15.6 (2.2)	19.6 (2.0)	19.1 (0.3)	69.6 (0.6)
w/ RERANK	45.9 (4.3)	9.2 (0.0)	31.7 (2.9)	38.2 (1.6)	18.6 (0.5)	69.7 (1.0)
Vicuna-13B SUMM (10-psg)	22.4 (3.0)	4.9 (0.1)	9.7 (1.3)	12.2 (1.2)	9.3 (0.4)	33.0 (3.7)
Vicuna-13B SNIPPET (10-psg)	48.1 (5.3)	11.2 (1.4)	27.2 (3.6)	27.9 (1.9)	18.4 (1.9)	76.8 (8.7)
Vicuna-13B ORACLE (3-psg)	41.6 (3.1)	17.1 (0.4)	20.2 (3.0)	26.5 (3.0)	20.0 (0.3)	72.0 (0.3)
LLaMA-33B VANILLA (3-psg)	58.8 (4.3)	6.2 (0.0)	9.3 (3.0)	12.1 (4.2)	16.9 (0.2)	60.0 (1.3)
w/ RERANK	65.9 (2.5)	6.0 (0.7)	22.5 (5.2)	26.1 (6.9)	17.5 (0.4)	61.0 (1.2)
LLaMA-33B SUMM (10-psg)	23.3 (2.0)	3.0 (0.2)	6.2 (0.5)	8.2 (0.7)	7.5 (0.4)	26.2 (2.3)
LLaMA-33B SNIPPET (10-psg)	53.2 (4.0)	7.4 (1.3)	13.7 (0.5)	15.1 (0.4)	14.4 (1.7)	53.3 (8.5)
LLaMA-33B ORACLE (3-psg)	63.7 (2.8)	11.4 (0.5)	11.9 (2.6)	15.4 (3.6)	17.9 (0.2)	61.7 (2.6)
Oasst-33B VANILLA (3-psg)	46.8 (7.6)	9.5 (0.2)	16.0 (2.5)	21.6 (3.5)	18.6 (0.3)	67.8 (1.5)
w/ RERANK	52.1 (6.1)	8.5 (0.5)	34.4 (2.9)	41.5 (2.5)	18.2 (0.3)	67.0 (1.5)
Oasst-33B SUMM (10-psg)	24.8 (2.8)	3.9 (0.3)	12.3 (0.2)	16.3 (0.3)	9.1 (0.3)	31.6 (1.5)
Oasst-33B SNIPPET (10-psg)	50.7 (4.6)	10.7 (1.2)	25.8 (3.3)	26.7 (2.3)	17.8 (1.8)	69.6 (8.6)
Oasst-33B ORACLE (3-psg)	50.7 (12.1)	15.8 (0.1)	20.8 (2.8)	28.0 (3.2)	19.4 (0.1)	70.3 (1.1)

Table 20: ELI5 full results.

---

Read the original question and passage, and generate 3 additional claims that are supported by the passage and answer the question.

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

Passage: The main difference between Shia and Sunni Muslim is related to ideological heritage and issues of leadership. This difference is first formed after the death of the Prophet Muhammad in 632 A.D. 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. Nowadays, Sunni and Shia are the major branches of Islam.

Claim 1: The major branches of Islam are Sunni and Shia.

Claim 2: Prophet Muhammad died in 632 A.D.

Claim 3: The ideological practice of the Sunni branch strictly follows Prophet Muhammad and his teachings.

Original question: What causes Bi-polar disorder?

Passage: Bipolar disorder is an emotional disorder that causes extreme mood swings between excitement and depression. The spectrum of mood swing may span from days to months. We are still not certain of the exact factors that cause such disorder, but genetics is considered a major factor.

Claim 1: One symptom of Bi-polar disorder is extreme mood swings between excitement and depression.

Claim 2: Genetics could be one of the major factors that causes Bi-polar disorder.

Claim 3: The mood swing from Bi-polar disorder can last days to months.

Original question: How do we hear differences in sound besides volume and pitch?

Passage: Pitch refers to the frequency of soundwave, and volume refers to the amplitude of the soundwave. Besides volume and pitch, we can also tell the difference between sounds based on the tone of sound. For example, we can differentiate the sound of different instruments based on the tone of the sounds.

Claim 1: Volume of sound is the amplitude of the soundwave.

Claim 2: Pitch is the frequency of soundwave.

Claim 3: We can use the tone of the sounds to differentiate the sound of different instruments.

Original question: How are we able to discern whether a sound is coming from in front of us or behind us?

Passage: There are multiple explanations for why we can localize sounds. One explanation is that sounds travelling to the corresponding side of one's ear will be slightly louder. Another explanation is that there is a slight difference in the hitting time to one's left and right ear based on the sound's direction. However, these explanation means that when a sound is exactly in front of someone or exactly behind someone, he or she can not tell the difference.

Claim 1: We can localize sounds by recognizing that the sound travelling to the corresponding side of one's ear will be slightly louder.

Claim 2: We can also localize sounds by recognizing the difference in hitting time to one's left and right ear based on the sound's direction.

Claim 3: We cannot tell the difference between a sound that is exactly in front of us or exactly behind us.

---

Table 21: Prompt used to generate the sub-claims for ELI5 questions. **Blue text is model generation. Brown text is the ELI5 example that we want to generate sub-claims for.** We construct the prompt by manually writing the sub-claims for three questions from the training set.

---

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.

---

Table 22: Instruction for VANILLA.

---

Instruction: Write a high-quality answer for the given question using only the provided search results and cite them properly using [1][2][3].

---

Table 23: Instruction for VANILLA light.



---

Summarize the following document within 50 words with the question of interest "{QUESTION}"  
Return "irrelevant" if the document is irrelevant to the question. Try to keep all the important  
dates, numbers, and names.

Title: {TITLE}  
Text: {TEXT}  
Summary:

---

Table 24: Prompts for SUMM.

---

Given the follow passage and the question "{QUESTION}", extract a useful span from the passage  
that can answer the question. Resolve all the coreference issues to make the extracted span  
understandable standalone. If the passage is not helpful for answering the question, return  
"irrelevant".

Title: {TITLE}  
Text: {TEXT}  
Extracted span:

---

Table 25: Prompts for SNIPPET.

---

Instruction: Write an accurate, engaging, and concise answer for the given question using only  
the provided search results and cite them properly. Use an unbiased and journalistic tone.  
Always cite for any factual claim.  
You are provided summaries/snippets of the search results. You can use "Check: Document [1][2]"  
to check the corresponding full documents (you should only check relevant documents and you can  
at most check 3 documents at a time) and use "Output:" to output a sentence in the answer. In the  
answer, cite properly by using [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. Use "End" to end the generation.

<Retrieve for question "...">  
<Get summaries/snippets for the passages and delete those that are "irrelevant">  
Document [1](Title: ...) {SUMMARY OR SNIPPET}  
...

Question: When did US break away from England?  
Check: Document [1][3]  
Document [1] {FULL TEXT}  
Document [2] {FULL TEXT}  
Output: The United States ... [1] ... [2].  
<Remove the full text of [1][2] from context>  
Check: Document [3]  
Document [3] {FULL TEXT}  
Output: The Treaty of Paris ... [3].  
<Remove the full text of [3] from context>  
End.

---

Table 26: Instruction for INTERACT

---

Instruction: Write an accurate, engaging, and concise answer for the given question using only  
the provided search results and cite them properly. Use an unbiased and journalistic tone.  
Always cite for any factual claim.  
You can use "Search: key words" to check the most relevant document's full text and use  
"Output:" to output a sentence in the answer. In the answer, cite properly by using [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. Use "End" to end  
the generation.

---

Table 27: Instruction for INLINESEARCH.

---

Instruction: Write an accurate, engaging, and concise answer for the given question. Use an  
unbiased and journalistic tone.

---

Table 28: Prompts for CLOSEDBOOK.

---

Instruction: Write an accurate, engaging, and concise answer for ...

Document [1](Title: How to Treat and Prevent Food Poisoning - MsPrepper) just a typical gastro upset. Salmonella is most commonly caused by eating undercooked or raw foods like eggs or meat. You know how your mom always warned you not to eat raw cookie dough? This is why. Most people do eat cookie dough and they are fine, but salmonella is a risk. If you do contract salmonella, you could start to feel bad within in a couple of hours after eating contaminated food, and sometimes it could take a day or two. Common symptoms are nausea and vomiting, loose stools (sometimes bloody), flu like symptoms, and stomach cramps. To treat

Document [2](Title: FDA Issues Warning About Eating Raw Cookie Dough, But Not For Salmonella Risks) FDA Issues Warning About Eating Raw Cookie Dough, But Not For Salmonella Risks Used to licking the spoon or placating yourself with full-on chunks of raw cookie dough? The Food and Drug Administration issued a warning on Tuesday that strongly advises against continuing the habit. The agency asserted that consuming raw batter of any kind, whether for bread, cookies or pizza, could make a person sick. While you may have been warned in the past against eating raw dough due to the risk of contracting salmonella from raw eggs, the FDA is citing raw flour as the culprit for a

Document [3](Title: It's Probably OK to Eat Raw Cookie Dough - As Long As You're Smart About It - The Crux - Very Top Secret Information) First, when most people think about health risks and cookie dough, they think about raw egg. Eggs can be contaminated with salmonella bacteria, and food safety recommendations encourage people to cook eggs until the white and yolk are firm in order to kill any bacteria. However, anyone making cookies can do things to reduce this risk by using pasteurized egg products. When my kids and I make cookie dough, we never use regular eggs. Instead, we use shell eggs that have been pasteurized to kill any harmful bacteria without actually cooking the egg itself. (A great public health innovation, if

Document [4](Title: How Dangerous Is It to Eat Raw Cookie Dough? | Men's Health) Can Eating Raw Cookie Dough Really Make You Sick? Scientists reveal the truth about this supposedly dangerous delicacy By Katherine Dempsey There are few things more tempting in life than eyeing a bowl of cookie dough and deciding whether or not to stick your finger in for a scoop. It's a bit like playing Russian roulette. You could get lucky and enjoy the delicious dough without consequence, but there's always the risk of getting seriously sick with a food-borne illness. That's because multiple ingredients within the dough could be contaminated by pathogens such as Salmonella and E. coli, says Soohyoun

Document [5](Title: How Dangerous Is It to Eat Raw Cookie Dough? | Men's Health) is usually pasteurized, so it's not likely that the egg would make you sick. However, other ingredients in the dough could potentially harbor pathogens. Experts say that a prime suspect in a 2009 E. coli outbreak linked to prepackaged cookie dough was actually the flour. Nuts and chocolate have also been linked to Salmonella outbreaks. Bottom line: You're better off skipping raw cookie dough. At least you'll have killer cookies as a consolation. The article How Bad Is It To Eat Raw Cookie Dough? originally ran on Prevention.com The 9 Foods Most Likely to Make You Sick Mmm, Just Don't ...

Question: Why is it bad to eat cookie dough for risk of salmonella but things like Cookie Dough Bites are ok?

Answer: Raw cookie dough is not recommended to be eaten due to the risk of salmonella [1][2]. Eating raw flour is also a risk for food poisoning [2]. The egg in cookie dough is usually pasteurized, so it's not likely that the egg would make you sick, but other ingredients in the dough could potentially harbor pathogens [4][5]. However, prepackaged cookie dough like Cookie Dough Bites is safe to eat because the dough is made with pasteurized egg products and heat-treated flour [2][3]..

Ground truth claims:

1. Cookie Dough Bites are safe to eat since they do not contain eggs.
2. Salmonella is a bacteria found in eggs.
3. Eating raw cookie dough carries the risk of salmonella infection.

---

Table 29: ELI5 example 1.

---

Instruction: Write an accurate, engaging, and concise answer for ...

Document [1](Title: Is Snapchat really worth \$19 billion? - CSMonitor.com) reporting that the Los Angeles-based company is aiming to raise \$500 million at a valuation of \$16 billion to \$19 billion, making it the third most highly valued tech start-up backed by venture capitalists. The Chinese handset maker Xiaomi is valued at \$45 billion, while Uber is estimated to be valued at about \$40 billion, according to data from CB Insights. Read MoreVC investment hits \$86B thanks to Uber, Xiaomi Snapchat was valued at \$10 billion in August, according to a Dow Jones report. Some of its investors from previous rounds include Benchmark, Lightspeed Venture Partners and Kleiner Perkins Caufield

Document [2](Title: What Are Venture Capital Investments? - DollarsAndSense.my) Ever wondered how highly valued technology giants like Google and Facebook were able to grow so fast and pay their employees so well in such a short amount of time, or how still growing start-ups like Uber are able to lose 1.2 billion US dollars in just the first half of this year alone and still command a valuation upwards of 50 billion US dollars? The answer lies with a special category of investment activity known as venture capital. Venture capitalists are professional investors who invest in a number of highly scalable high-risk technology ventures hoping to make a multi-fold

Document [3](Title: Opinion | What Dara Khosrowshahi Must Do to Save Uber - The New York Times) at a discount. These are troubling signs. Every start-up must one day fulfill the market's demand that it turn a profit, but Uber has never figured out how to do that. While ride sharing in some form will probably survive, it's more likely that without some drastic changes, Uber won't be around in three to five years. Mr. Khosrowshahi must avoid the mistakes of his predecessor by accepting that "pivots" (Silicon Valley-speak for the desperate changes troubled companies make to reassure their venture capitalist funders) are not the answer. None of the pivots Mr. Kalanick tried - like on-demand delivery

Document [4](Title: Snapchat raising funding round at \$19 billion valuation: Report) Snapchat raising funding round at \$19 billion valuation: Report The mobile messaging app Snapchat is said to be seeking funding at a \$19 billion valuation, according to a report, citing a person with knowledge of the matter. Bloomberg is reporting that the Los Angeles-based company is aiming to raise \$500 million at a valuation of \$16 billion to \$19 billion, making it the third most highly valued tech start-up backed by venture capitalists. The Chinese handset maker Xiaomi is valued at \$45 billion, while Uber is estimated to be valued at about \$40 billion, according to data from CB Insights.

Document [5](Title: Unicorns And Wall Street | MoneyTips) not start showing profits soon. Even Uber, with its massive \$62-billion-plus valuation, will test investor's patience if it does not start showing signs of a profit in the next few years. Cracks are starting to show in terms of pre-IPO devaluations. Dropbox and Snapchat have seen a nearly one-quarter drop in their pre-IPO valuations, while Square's pre-IPO value has dropped by approximately one-third. Larger investors and venture capitalists are sending a message to the unicorns: Outline a solid path to profits and start producing them soon, or expect your valuation to be slashed. Diving ahead with an IPO is not

...

Question: How are firms like snapchat, uber etc valued so highly while still not making a profit? Do venture capitalists not expect some form of repayment within a number of years?

Answer: Venture capitalists invest in a number of highly scalable high-risk technology ventures hoping to make a multi-fold return on their investment in a short amount of time [2]. Highly valued companies like Snapchat, Uber, and Xiaomi are able to grow so fast and command high valuations because they have the potential to generate significant revenue in the future [2][4]. Venture capitalists expect a return on their investment within a number of years, but they also invest in the long-term potential of the company [2]. It is important to note that every start-up must eventually turn a profit, but some start-ups like Uber have not yet figured out how to do that [3][5].

Ground truth claims:

1. Firms like Snapchat and Uber need to establish their brand and amass users before introducing ads.
  2. Introducing ads too early can deter potential users.
  3. Uber is reinvesting a lot of money to make their service better.
- 

Table 30: ELI5 example 2.