

# Can LLMs Evaluate Complex Attribution in QA? Automatic Benchmarking using Knowledge Graphs

Nan Hu<sup>1,2\*</sup>, Jiaoyan Chen<sup>3†</sup>, Yike Wu<sup>1,2</sup>, Guilin Qi<sup>1,6†</sup>,  
Hongru Wang<sup>7</sup>, Sheng Bi<sup>5</sup>, Yongrui Chen<sup>1,2</sup>, Tongtong Wu<sup>8</sup>, Jeff Z. Pan<sup>4†</sup>

<sup>1</sup> School of Computer Science and Engineering, Southeast University

<sup>2</sup> Key Laboratory of New Generation Artificial Intelligence Technology and its Interdisciplinary Applications (Southeast University), Ministry of Education

<sup>3</sup> The University of Manchester <sup>4</sup> The University of Edinburgh

<sup>5</sup> Law and Innovation Lab, School of Law, Southeast University

<sup>6</sup> State Key Lab. for Novel Software Technology, Nanjing University

<sup>7</sup> The Chinese University of Hong Kong <sup>8</sup> Monash University

## Abstract

Attributed Question Answering (AQA) has attracted wide attention, but there are still several limitations in evaluating the attributions, including lacking fine-grained attribution categories, relying on manual annotations, and failing to compare attributions with only subtle differences. To bridge these gaps, we introduce Complex Attributed Question Answering (CAQA), a large-scale benchmark containing comprehensive attribution categories, automatically generated using Knowledge Graphs (KGs), and complex attribution scenarios. We have conducted extensive experiments to verify the effectiveness of CAQA, including the benchmarking of 25 automatic evaluators, their comparison with human evaluators, the testing of LLM evaluators fine-tuned by CAQA and so on. These experiments also lead to a series of important findings that can benefit the future research of AQA. All the codes and data are publicly accessible at <https://github.com/HuuNan/CAQA-Benchmark>.

## 1 Introduction

Generative AI is increasingly adept together with other techniques like search engines to produce answers to natural language questions. However, their tendency to generate confident yet inaccurate or “hallucinated” contents (Ji et al., 2023; Pan et al., 2023) poses significant risks in high-stakes domains such as medicine (Lee et al., 2023). Therefore, Question Answering (QA) with attribution has been proposed, where not only answers but also citations (or evidence snippets) for supporting the answers are output (Menick et al., 2022; Rashkin et al., 2023; Bohnet et al., 2022; Li et al., 2023).

Despite their potential, state-of-the-art implementations of Attributed QA (AQA), exemplified

by Large Language Models (LLMs) with search engines like Bing Chat, perplexity.ai and YouChat<sup>1</sup>, still often produce erroneous attributions (Liu et al., 2023). To develop systems for higher-quality attributions, it is crucial to explore effective automatic attribution evaluation methods, which can not only compare and analyze different AQA systems, but also provide feedback to improve their attributions (Asai et al., 2024; Yue et al., 2023; Gao et al., 2023a; Bohnet et al., 2022), alleviating the issues of factuality, faithfulness and hallucination (Amouyal et al., 2023; Asai et al., 2024). However, the existing AQA benchmarks (see Table 1) are inadequate due to their limited sizes, incomplete attribution categories (e.g., they do not fully consider partially supportive where only some sub-facts in the answer are supported), and ignore of complex attribution scenarios which require reasoning with multiple evidence under various logic operations, and are common in Bing Chat and retrieve-and-read systems (Malaviya et al., 2024).

In this work, we adopt a comprehensive set of attribution categories including supportive, partially supportive, contradictory and irrelevant (see Table 2 for examples) and define different levels of attribution complexity based on the reasoning required to infer the answer: single, union, intersection, and concatenation (see Table 17 for examples). Based on these foundations, we construct a larger-scale Complex Attributed Question Answering (CAQA) benchmark by an automatic generation method based on a Knowledge Graph (KG) (Pan et al., 2017b,a) composed of relational facts (Hogan et al., 2021; Bollacker et al., 2008) and two existing KGQA datasets containing question-answer pairs and corresponding KG queries. Briefly, the construction method first extends a query by introducing additional logi-

\*Email: nanhu@seu.edu.cn

†Corresponding authors. Email: gqi@seu.edu.cn, jiaoyan.chen@manchester.ac.uk, j.z.pan@ed.ac.uk

<sup>1</sup>bing.com/new, perplexity.ai, https://you.com/

Table 1: CAQA and existing benchmarks. Attribution categories include Supportive (S), Non-supportive (N), Partially Supportive (P), Contradictory (C), Irrelevant (I) and Extrapolatory (E), with E and I treated as equivalent. Comp. denotes whether the benchmark considers attribution complexity. Auto. denotes whether the benchmark is automatically constructed without manual annotation.

Benchmarks	#Sample	Category	Comp.	Auto.
Bohnet et al. (Bohnet et al., 2022)	23,000	S/N	✗	✗
HAGRID (Kamalloo et al., 2023)	2,638	S/N	✗	✗
ExpertQA (Malaviya et al., 2024)	2,177	S/N	✗	✗
AttributionBench (Li et al., 2024b)	17,816	S/N	✗	✗
Liu et al. (Liu et al., 2023)	11,037	S/P/N	✗	✗
ALCE (Gao et al., 2023b)	800	S/P/N	✗	✗
AttrEval-Gen (Yue et al., 2023)	242	S/C/E	✗	✗
AttrEval-Sim (Yue et al., 2023)	64.2K	S/C/E	✗	✓
CAQA (Ours)	161.1K	S/P/C/I	✓	✓

cal operators for increasing reasoning complexity, then employs the extended queries to extract KG sub-graphs (Huang et al., 2024), and finally edits these sub-graphs using different strategies and re-writes them into natural language citations as attributions of different categories. This method is flexible, allowing the generation of benchmarks with varied features, and adaptable to different KGs and KGQA datasets.

We use two existing attribution evaluators that are fine-tuned on specific data and 23 LLM evaluators under the zero-shot, few-shot and fine-tuning settings to demonstrate the effectiveness of CAQA for assessing and developing AQA evaluators. Through extensive evaluation, we get a series of important findings. (1) All the evaluators struggle to identify the nuanced negative attribution categories in both zero-shot and few-shot settings. With fine-tuning, the F1 scores of all the categories exceed 90% for most LLM evaluators. (2) All the evaluators perform poorly in recognizing attributions of partially supportive and perform worse on more complex attribution scenarios, e.g. the GPT models perform worse on those requiring concatenation and intersection, while other open source LLMs perform worse on those requiring union. (3) The automatically generated CAQA is reliable with high consistency as human annotators. (4) When tested on an out-of-distribution dataset, LLM evaluators fine-tuned by CAQA outperform the existing particularly developed evaluators.

## 2 Related Work

**Attributed Question Answering.** LLMs now lead the performance in QA, but often produce hallu-

cinations (Ji et al., 2023; Xiao and Wang, 2021; Wang and Sennrich, 2020; Shuster et al., 2021). To alleviate this issue, some studies (Menick et al., 2022; Nakano et al., 2021; Gao et al., 2023b) train attributed models to answer questions that generate answers along with supporting evidence, typically in the form of citations or references. Some other studies augment LLMs with external tools (Mialon et al., 2023; Shen et al., 2023; Schick et al., 2023; Wang et al., 2024, 2025; He et al., 2025) such as retrievers (Han et al., 2023; Shi et al., 2024; Asai et al., 2024; Izacard et al., 2022; Min et al., 2024, 2025; Shen et al., 2025) and search engines (Nakano et al., 2021; Komeili et al., 2021), or incorporate external references for attribution. While these approaches improve factual grounding, ensuring attribution accuracy and consistency remains an open challenge, underscoring the need for more systematic evaluation.

Attributed QA is conceptually related to the task of Attributing Unanswerable Questions (AUQ) (Moradisani et al., 2024), although the two tasks differ in scope. AUQ focuses on identifying questions that cannot be answered based on the given context, typically due to factors such as negation, entity swaps, or missing information, and explaining why a valid answer does not exist. In contrast, AQA aims to generate answers along with explicit attribution from the context. A well-designed AQA system should naturally refrain from answering when no supporting evidence is available, thereby addressing unanswerability through its attribution mechanism.

**Attribution Evaluation.** Current methods for evaluating attribution predominantly depend on human annotation (Nakano et al., 2021; Bohnet et al., 2022; Liu et al., 2023; Rashkin et al., 2023; Muller et al., 2023), which is costly and inefficient. Recent studies propose automatic attribution evaluators based on LLMs, such as AUTOIS (Gao et al., 2023a; Bohnet et al., 2022) and ATTRSCORE (Yue et al., 2023). However, existing benchmarks are inadequate for evaluating and advancing attribution evaluators due to their limited sizes and restricted evaluation settings, including incomplete attribution categories and the ignorance of attribution complexity. They mostly classify attribution into only two categories: the cited evidence *supports* or *does not support* the answer (Gao et al., 2023b; Li et al., 2024a,b; Malaviya et al., 2024; Bohnet et al., 2022). Some benchmarks (Gao et al., 2023b; Liu et al., 2023; Zhang et al., 2024) add

a third category, *partially supportive*, but their sizes are small and reliance on manual annotation. (Yue et al., 2023) presents a method for automatically generating attribution annotations to construct large-scale samples with categories of *supportive*, *contradictory*, and *extrapolatory* (equivalent to *irrelevant*). However, their method cannot support *partially supportive*, as it relies solely on answer word replacement. Our work addresses these limitations by proposing a novel method based on KGs and KGQA datasets to automatically create a large-scale AQA benchmark with comprehensive attribution categories. Notably, our benchmark is the first to offer fine-grained evaluation for partially supportive evidence and consider varying levels of logical reasoning complexity in attribution.

### 3 Definitions in QA Attribution

#### 3.1 Task Formulation

This work aims to evaluate the attribution in AQA. It is to verify whether an evidence, which has one or multiple citations (references) with facts stated, can sufficiently support a generated answer towards a natural language question. Formally, given a question  $q$ , an answer statement  $a$  and an evidence  $e$ , the objective of attribution evaluation is to map them to an attribution category  $t$  (a.k.a. class label). It can be represented by the function  $\mathcal{F} : \mathcal{Q} \times \mathcal{A} \times \mathcal{E} \mapsto \mathcal{T}$ , where  $\mathcal{Q}$ ,  $\mathcal{A}$  and  $\mathcal{E}$  denote the sets of questions, answers and evidence, respectively, and  $\mathcal{T}$  denotes the set of potential categories, such as  $\{\textit{supportive}, \textit{partially supportive}, \textit{contradictory}, \textit{irrelevant}\}$  which mean “ $e$  is supportive, partially supportive, contradictory or irrelevant to the fact that  $a$  is the answer of  $q$ .”

#### 3.2 Attribution Categorization

We analyse the results of practical AQA systems (Gao et al., 2023b) and find that apart from correct attributions that are *supportive*, there are three main categories of incorrect attributions: *partially supportive*, *contradictory* and *irrelevant*. More details are shown in Appendix G. The four attribution categories are defined below:

- **Supportive (Sup.):** The evidence includes facts that can fully support the answer statement.
- **Partially Supportive (Par.):** The evidence lacks a part of the facts that are required to infer the answer statement.
- **Contradictory (Con.):** The evidence includes facts that can infer a different answer statement.

- **Irrelevant (Irr.):** The evidence has no facts that can be used to infer the answer statement.

Table 2 provides examples of the four attribution categories. In the **supportive** example, the answer is backed by citation [1], which confirms that “*Ruth Madoc plays Fruma Sarah in Fiddler on the Roof.*” In the **partially supportive** example, the answer cites [1] but does not fully align with the complete context provided, mentioning only “*the actor Heath Ledger stars in the film 10 Things I Hate About You*” and missing the information “*Heath Ledger plays the character Patrick*”. In the **contradictory** example, the citation [1] states “*The Puppetoon Movie is directed by Arnold Leibovit,*” which contradicts the generated answer. The **irrelevant** example involves citing [1], which discusses an unrelated actor, Chris Rankin, and his career offers no relevant facts to verify the answer. Note that the contradictory category differs from the partially supportive or irrelevant categories in reasoning. In the former, the evidence leads to another answer that conflicts with the generated answer.

#### 3.3 Attribution Complexity

Previous research has not explored different levels of complexity in inferring the answer. Malaviya et al. (2024) has shown that AutoIS (Bohnet et al., 2022), the most commonly used automatic attribution evaluator, often mistakes in scenarios that require multiple pieces of evidence to validate the answer. To advance automatic evaluators, our benchmark incorporates reasoning complexity by categorizing attribution into four levels of complexity, based on the reasoning logic of supporting facts in the citations (see Table 17 for examples):

- **Single:** The answer is supported by one single citation.
- **Union:** The answer is supported by independent facts from multiple citations.
- **Intersection:** The answer is supported by facts with common entities from multiple citations.
- **Concatenation:** The answer is supported by chains of facts from multiple citations.

### 4 Benchmark Construction

In this section, we introduce our benchmark construction method. Figure 1 presents its overview including four key steps: (1) Query Collection: given a KGQA dataset, we collect data corresponding to basic KG logical queries; (2) Query Extension: two logical operators are applied to increase the

Table 2: Examples of the four attribution categories. **Green**, **yellow**, and **red** text indicate the content in the answer that is **supported**, **not supported**, or **contradicted** by the content in the citation, respectively.

Attribution Category	Examples
Supportive	<p><b>Question:</b> Who plays Fruma Sarah in Fiddler on the Roof?</p> <p><b>Answer:</b> Fruma Sarah is a character in the musical “Fiddler on the Roof”, and Ruth Madoc played the role [1].</p> <p><b>Citations:</b> [1] ... In 1971 Ruth Madoc played Fruma Sarah in the film version of the musical “Fiddler on the Roof”, and in 1972 she appeared as ...</p>
Partially Supportive	<p><b>Question:</b> Who plays Patrick in 10 Things I Hate About You?</p> <p><b>Answer:</b> Patrick is played by actor Heath Ledger in the film 10 Things I Hate About You [1].</p> <p><b>Citations:</b> [1] 10 Things I Hate About You is a 1999 American teen romantic comedy-drama film directed by Gil Junger and starring Heath Ledger, Julia Stiles, Joseph Gordon-Levitt, and Larisa Oleynik. The screenplay, written by ...</p>
Contradictory	<p><b>Question:</b> Who directed a George Pal’s production?</p> <p><b>Answer:</b> George Pal directed a production called Puppets [1].</p> <p><b>Citations:</b> [1] ... The Puppetoon Movie is a 1987 animated film written, produced, and directed by Arnold Leibovit ...</p>
Irrelevant	<p><b>Question:</b> Who played the weasley brothers in Harry Potter?</p> <p><b>Answer:</b> James and Oliver Phelps, identical twin actors, played the roles of Fred and George Weasley in the Harry Potter film series [1].</p> <p><b>Citations:</b> [1] Chris Rankin plays of “Bugsy Malone”, “The Lion, The Witch and The Wardrobe” and Harry Potter series ... he plays a brother of Harry Potter’s best friend, ...</p>

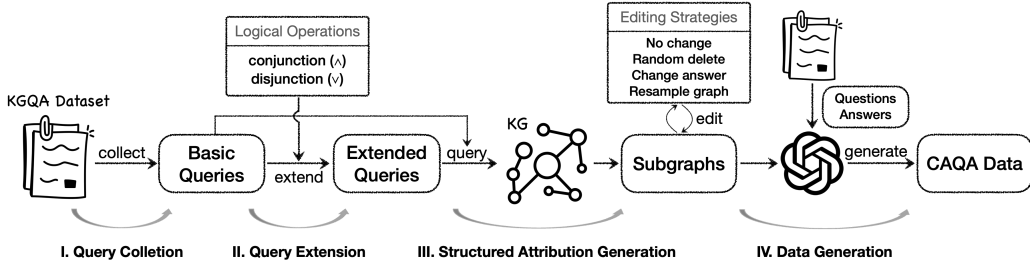


Figure 1: The entire process of constructing the CAQA benchmark.

complexity of the original queries; (3) Structured Attribution Generation: the extended queries are grounded in the KG to obtain relevant subgraphs, which are then probabilistically edited using four strategies to generate new subgraphs with four attribution categories; (4) Data Generation: we produce AQA data, where each instance consists of an extended question, rephrased answer entities, citations derived from subgraphs, as well as the attribution category and complexity labels.

#### 4.1 Query Collection

The selection of KGs and KGQA datasets is primarily motivated by two observations: (1) KGQA is a well-established task with a wealth of open resources, as evidenced by 25 KGQA datasets for 5 KGs reported in (Jiang and Usbeck, 2022); (2) existing KGQA datasets contain high-quality question-answer pairs and corresponding KG logical queries, often expressed in SPARQL, which are capable of deriving the correct answers and can be leveraged to generate evidence.

A KG is composed of relational facts in the form

of triple, i.e.,  $(h, r, t)$ , where  $h$  and  $t$  denote a subject entity and a object entity, respectively, and  $r$  denotes a relation between them. A KGQA dataset  $D = \{S_1, S_2, \dots, S_N\}$  consists of samples in the form of  $S_i = (q_i, a_i, l_i)$ , where  $q_i$  denotes a natural language question,  $a_i$  denotes its answer entity, and  $l_i$  denotes the KG logical query of  $q_i$ . Our data collection focuses on samples where the KG logical query falls into one of three types: **single-triple**, **path-like**, or **tree-like queries**. As shown in the first three columns in Table 3, a single triple query denoted as  $(e_0, r_0, ?a)$  indicates that the answer entity  $?a$  can be obtained via the subject  $e_0$  and the relation  $r_0$ . A path-like query denoted as  $[e_0, r_0, ?v_1, \dots, ?v_{n-1}, r_{n-1}, ?a]$  represents that the answer  $?a$  is reachable through an  $n$ -hop path starting from  $e_0$ , traversing  $n$  relations and  $n - 1$  intermediate entities represented by query variables  $?v_1, \dots, ?v_{n-1}$ . Specifically,  $?v_1$  denotes the first intermediate variable in the query path. Notably, a path-like query reduces to a single-triple query when  $n = 1$ . A tree-like query, formulated as  $\bigwedge_{i=0}^{n-1} (e_i, r_i, ?a)$ , includes  $n$  distinct triples, each



originating from different subjects and converging on the answer  $?a$ .

## 4.2 Query Extension

For each KGQA example  $S_i = (q_i, a_i, l_i)$ , we extend the logical query  $l_i$  to  $l'_i$  using a set of predefined query extension rules designed based on the logical operations *intersection* (i.e., conjunction,  $\wedge$ ) and *union* (i.e., disjunction,  $\vee$ ) (Ren et al., 2023)<sup>2</sup>.

Table 3 outlines the extension rules and we show concrete examples of each extension rule in Appendix B. For a single-triple query  $l$ , the *union* operation is used. Initially, we retrieve entities from the KG that share the same name as  $e_0$  in  $l$ , producing a set of  $m$  entities  $\{e_1, \dots, e_m\}$ , where  $m$  may be zero. Subsequently, we generate logical queries  $(e_1, r_0, ?a), \dots, (e_m, r_0, ?a)$  by combining the retrieved entities and the relation  $r_0$  from  $l$ . These new queries are then merged with  $l$  via *union*, resulting in a union-tree-like query structure. This structure implies that the final answer is derived as the union of the answers obtained from each subquery.

For a path-like query or a tree-like query, we apply the *intersection* operation in two distinct ways. In the first way, we identify a unique subject entity  $e_0$  for path-like queries or randomly select a subject  $e_k$  for tree-like queries. We then retrieve corresponding triples  $(e_1, r_n, e_0)$  or  $(e_n, r_n, e_k)$  from the KG, where  $r_n$  represents a relation not present in  $l$ . These new triples are appended to the respective queries, ensuring that  $e_0$  and  $e_k$  are connected entities. This process maintains the overall structure of the path-like or tree-like query. In the second way, we append a new query  $(e_1, r_n, ?a)$  or  $(e_n, r_n, ?a)$  to the respective logical forms, ensuring that the intersection of the answers obtained from the new queries with those from  $l$  is non-empty. Through this extension, both the path-like query and tree-like query are converted into the tree-like structures.

For both a path-like query ( $n \geq 2$ ) and a tree-like query, the two intersection extensions are applied with equal probability. In contrast, for single-triple queries (a special case of path-like queries), four operations are equally likely: union extension, two types of intersection extension, and no extension (to preserve some single-triple queries). The extension process results in four query types: *single-tree*,

<sup>2</sup>Our method can be easily extended with more logical operations like Negation and Kleene Plus (Ren et al., 2023).

*union-tree-like*, *tree-like*, and *path-like*, corresponding to the attribution complexity types (denoted by  $r$ )—*single*, *union*, *intersection*, and *concatenation*.

## 4.3 Structured Attribution Generation

We first obtain a KG subgraph  $\mathcal{G}$  by grounding each extended query  $l'$  in the KG, which returns the entities that are assigned to all the variables in the query for inferring the answer. The subgraph  $\mathcal{G}$  is regarded as the structured attribution to support the answer to the question and falls under the *supportive* attribution category. To get structured attributions of *partially supportive*, *contradictory*, and *irrelevant*, we edit  $\mathcal{G}$  as follows.

- **Partially Supportive.** The *partially supportive* subgraph  $\mathcal{G}_{In}$  is generated by partial deletion, resulting in a subgraph that cannot fully support the answer. For path-like queries, we randomly delete one triple in  $\mathcal{G}$ . For tree-like or union-tree queries, we delete a path connecting one of the subject entities to the answer. In the case of single-triple queries, no deletion is performed.

- **Contradictory** The *contradictory* subgraph  $\mathcal{G}_C$  is constructed by altering  $\mathcal{G}$  such that its reasoning conflicts with the answer. This is done by replacing the answer entity in  $\mathcal{G}$  with a non-answer entity of the same type. Especially for queries involving a union operation, we replace one of the answer entities within  $\mathcal{G}$ .

- **Irrelevant** The *irrelevant* subgraph  $\mathcal{G}_{Ir}$  is obtained by selecting an entirely different subgraph from the KG that is structurally similar to  $\mathcal{G}$  but contains unrelated entities and relations, except for the subject entity in  $\mathcal{G}$ .

## 4.4 Data Generation

We employ ChatGPT with tailored prompts to transform the subgraphs of  $\mathcal{G}$ ,  $\mathcal{G}_{In}$ ,  $\mathcal{G}_C$  and  $\mathcal{G}_{Ir}$  into citations (attributions) of *supportive*, *partially supportive*, *contradictory* and *irrelevant*, respectively. When the original logical query  $l$  is expanded to  $l'$ , the original question  $q$  is similarly extended to a new question  $\tilde{q}$  using ChatGPT. In addition, the answer entity  $a$  is paraphrased into a more detailed answer statement  $\tilde{a}$ . Ultimately, it yields an AQA sample consisting of the question  $q$  or  $\tilde{q}$ , the answer statement  $\tilde{a}$ , the citation  $c$ , the attribution category  $t$ , and the complexity label  $r$ . More details on this step can be found in Appendix A.

Table 3: The rules for extending the logical queries utilizing two query operations: intersection ( $\wedge$ ) and union ( $\vee$ ). All queries are classified according to their structure as single-triple (S.) queries, path-like (P.) queries, tree-like (T.) queries and union-tree-like (U.) queries. The ‘Examples’ column presents corresponding graph representations for the case where  $n = 2$ ,  $m = 2$ , and  $k = 0$ , where grey nodes represent variables for answer entities.



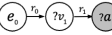
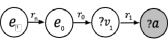
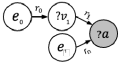

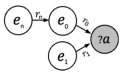
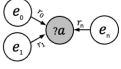
Original Query $l$			Extended Query $l'$		
Definitions	Structures	Examples	Definitions	Structures	Examples
$(e_0, r_0, ?a)$	S.		$(e_0, r_0, ?a)$ $\vee (e_1, r_0, ?a) \vee \dots \vee (e_m, r_0, ?a)$	U.	
$[e_0, r_0, ?v_1, \dots, ?v_{n-1}, r_{n-1}, ?a]$	P.		$[e_0, r_0, ?v_1, \dots, ?v_{n-1}, r_{n-1}, ?a]$ $\wedge (e_1, r_n, e_0)$	P.	
			$[e_0, r_0, ?v_1, \dots, ?v_{n-1}, r_{n-1}, ?a]$ $\wedge (e_1, r_n, ?a)$	T.	
$\wedge_{i=0}^{n-1} (e_i, r_i, ?a)$	T.		$\wedge_{i=0}^{n-1} (e_i, r_i, ?a), i \neq k$ $\wedge (e_n, r_n, e_k) \wedge (e_k, r_k, ?a)$	T.	
			$\wedge_{i=0}^{n-1} (e_i, r_i, ?a) \wedge (e_n, r_n, ?a)$	T.	

Table 4: Statistics of the CAQA dataset

Classes		Train	Test	Total
		137,211	23,963	161,174
Category	Sup.	39,489	6,668	46,157
	Par.	28,868	5,065	33,933
	Con.	36,620	6,423	43,043
	Irr.	32,234	5,807	38,041
Complexity	Single	73,795	10,443	84,238
	Concatenation	46,783	8,455	55,238
	Union	5,347	886	6,233
	Intersection	11,286	4,179	15,465

## 5 Experiment Setup

**The CAQA Benchmark.** It is constructed using two KGQA datasets: GrailQA (Gu et al., 2021) and WebQuestionsSP (Yih et al., 2016), and the Freebase KG (Bollacker et al., 2008). See Table 4 for its statistics. Additionally, we manually annotated the attribution categories of 300 test samples to assess their consistency with the original categories (see results in Section 6.2). More details on CAQA are in Appendix C, and human annotation processes are described in Appendix H.

**ALCE-FineGrained.** We manually annotated 215 samples of the ALCE benchmark with the four attribution categories, and get a new benchmark, ALCE-FineGrained, as an out-of-distribution testing set (see results in Section 6.4). Details of human annotation are given in Appendix H.

**Attribution Evaluators.** We evaluate state-of-the-art LLMs as attribution evaluators across different architectures and sizes: GPT-3.5, GPT-4o, GPT-4o

mini, LLaMA-2/3/3.1 (Grattafiori et al., 2024), Vicuna (Chiang et al., 2023), Gemma-2 (Team et al., 2024), Mistral (Jiang et al., 2023), Phi-3 (Abdin et al., 2024), and Qwen2.5 (Hui et al., 2024). We conduct evaluations in three settings: the zero-shot setting, the few-shot setting where several attribution examples are given, and the fine-tuning setting where the LLM is trained with the samples in the training set. Additionally, we test two specially developed automatic attribution evaluators AUTOIS (Honovich et al., 2022) and ATTRSCORE (Yue et al., 2023). More details on the evaluators’ implementation are given in Appendix D.

**Metrics.** We report F1 score for the performance on each attribution category and micro-F1 score for the performance on each complexity level and overall performance. We also use the FACTSCORES metric (Min et al., 2023) for a fine-grained evaluation of the partially supportive category.

## 6 Evaluation

### 6.1 Overall Results

Table 5 shows the results of the attribution evaluators on CAQA. Appendix E shows more experiments. Our analysis is as follows:

**All the evaluators perform poorly in identifying nuanced negative attribution categories, especially partially supportive, under the zero-shot and few-shot settings.** Smaller LLMs ( $\leq 14$ B), except for Gemma-2 and Qwen-2.5, perform extremely poorly on all three negative categories,

Settings	Evaluators (LLM Size)	Category					Complexity			
		Sup.	Par.	Con.	Irr.	Overall	Single	Concatenation	Intersection	Union
Zero-Shot	LLaMA-3 (8B)	0.467	0.120	0.072	0.007	0.296	0.304	0.271	0.283	0.259
	LLaMA-3.1 (8B)	0.544	0.049	0.130	0.017	0.318	0.319	0.326	0.319	0.285
	LLaMA-3.1 (70B)	0.688	0.168	0.547	0.609	0.544	0.545	0.549	0.545	0.499
	Mistral-v0.3 (7B)	0.661	0.160	0.051	0.334	0.362	0.363	0.374	0.356	0.337
	Mixtral-v1.0 (8x7B)	0.677	0.094	0.170	0.635	0.494	0.495	0.516	0.484	0.487
	Vicuna (7B)	0.513	0.100	0.064	0.199	0.327	0.343	0.273	0.312	0.256
	Vicuna (13B)	0.634	0.211	0.393	0.275	0.405	0.432	0.314	0.361	0.374
	Gemma-2 (9B)	0.667	0.280	0.498	0.624	0.556	0.557	0.572	0.552	0.508
	Gemma-2 (27B)	0.653	0.184	0.569	0.646	0.566	0.566	0.579	<u>0.566</u>	0.537
	Qwen-2.5 (14B)	0.680	0.132	<u>0.708</u>	<u>0.660</u>	<u>0.617</u>	0.640	<b>0.622</b>	<b>0.611</b>	0.547
	Qwen-2.5 (72B)	0.629	0.266	0.701	0.471	0.571	0.593	0.583	0.565	0.530
	GPT-4	<b>0.771</b>	<b>0.456</b>	<b>0.745</b>	0.473	<b>0.630</b>	<b>0.685</b>	0.451	0.514	<b>0.616</b>
	GPT-4o	<u>0.769</u>	<u>0.445</u>	0.598	0.626	<b>0.630</b>	<u>0.676</u>	<u>0.591</u>	0.470	<u>0.588</u>
GPT-4o-mini	0.718	0.297	0.632	<b>0.703</b>	0.616	0.672	0.473	0.444	0.559	
Few-Shot	LLaMA-3 (8B)	0.573	0.202	0.234	0.156	0.336	0.356	0.279	0.310	0.294
	LLaMA-3.1 (8B)	0.631	0.101	0.307	0.157	0.353	0.354	0.375	0.343	0.362
	LLaMA-3.1 (70B)	0.698	0.259	0.713	0.665	0.627	0.626	<u>0.638</u>	<u>0.632</u>	<u>0.589</u>
	Mistral-v0.3 (7B)	0.571	0.072	0.252	0.134	0.342	0.343	0.351	0.336	0.326
	Mixtral-v1.0 (8x7B)	0.589	0.084	0.356	0.558	0.455	0.456	0.473	0.451	0.416
	Vicuna (7B)	0.578	0.183	0.081	0.324	0.325	0.337	0.272	0.354	0.311
	Vicuna (13B)	0.633	0.208	0.383	0.288	0.403	0.427	0.315	0.397	0.374
	Gemma-2 (9B)	0.705	0.390	0.568	0.593	0.572	0.571	0.586	0.578	0.526
	Gemma-2 (27B)	0.646	0.231	0.670	0.572	0.570	0.570	0.585	0.569	0.511
	Qwen-2.5 (14B)	0.699	0.257	<b>0.741</b>	<u>0.676</u>	0.646	0.680	<b>0.656</b>	<b>0.638</b>	<b>0.608</b>
	Qwen-2.5 (72B)	0.721	0.400	<u>0.736</u>	0.503	0.617	0.635	0.608	0.626	0.592
	GPT-4	<b>0.794</b>	<b>0.520</b>	0.728	0.653	<b>0.680</b>	<b>0.745</b>	0.492	0.473	0.559
	GPT-4o	<u>0.783</u>	<u>0.507</u>	0.683	0.641	<u>0.664</u>	0.730	0.559	0.449	0.529
GPT-4o-mini	0.763	0.435	0.705	<b>0.700</b>	0.657	<u>0.741</u>	0.430	0.404	0.588	
Fine-Tuning	LLaMA-3 (8B)	0.935	0.901	0.935	0.928	0.926	0.935	0.820	0.930	0.924
	LLaMA-3.1 (8B)	<b>0.946</b>	0.919	<u>0.944</u>	<u>0.934</u>	<u>0.941</u>	<u>0.953</u>	<u>0.850</u>	<b>0.939</b>	<b>0.945</b>
	Mistral-v0.3 (7B)	<u>0.944</u>	<u>0.921</u>	<b>0.947</b>	<b>0.935</b>	<b>0.942</b>	<b>0.956</b>	<b>0.852</b>	<u>0.937</u>	<u>0.941</u>
	Vicuna (7B)	0.937	0.907	0.940	0.906	0.932	<b>0.956</b>	0.823	<b>0.936</b>	0.939
	Vicuna (13B)	0.942	<b>0.923</b>	0.939	0.923	<b>0.933</b>	0.950	0.847	0.935	0.940
		Existing Attribution Evaluators								
	AUTOIS (11B)	0.609	-	-	-	-	-	-	-	-
	ATTRSCORE (13B)	0.687	-	0.523	0.541	0.521	0.559	0.410	0.432	0.353

Table 5: Overall results on CAQA. Bold and underlined values are the best and second-best performance, respectively. “-” indicates that the method is not applicable. More complete results can be found in Table 16.

suggesting that none of them are capable of distinguishing subtle differences between negative attributions. Larger LLMs ( $\geq 27B$ ) demonstrate strong performance on certain negative categories but consistently underperform on *partially supportive* category. For example, in both zero-shot and few-shot settings, the highest scores—0.456 and 0.52, respectively—are achieved by GPT-4. This highlights the challenges faced by the attribution evaluator in accurately identifying *partially supportive* category. We find that evaluators often classify *partially supportive* as *supportive*, even though it is apparent that part of the information is missing.

**Fine-tuning effectively enhances evaluator performance, while few-shot prompting is particularly beneficial for large-scale evaluators.** All fine-tuned evaluators trained on our dataset achieve similar performance, exceeding 90%, except for Mistral-v0.2. AutoIS and AttrScore, which were fine-tuned on other benchmarks, demonstrate limited ability to identify fine-grained attribution categories. Large-scale LLMs ( $\geq 70B$ ) and GPT-

series models exhibit significantly stronger in-context learning capabilities. Under few-shot settings, these models improve by an average of 4.84% in distinguishing subtle attribution differences, demonstrating deeper contextual understanding. However, few-shot prompting provides only marginal performance gains for smaller LLMs and can even be detrimental to certain models. See more evaluation of the few-shot setting in Appendix E.

**The evaluators are often impacted by the keyword co-occurrence and fail to capture the reasoning in complex attribution scenarios.** Considering examples of *irrelevant*: the question is “What is the soundtrack of the video game X?” The answer is, “The video game X’s soundtrack is Y,” and the evidence is, “Z is a video game designer who has designed games such as X”. The evaluators often incorrectly treats the attribution as *supportive* due to the co-occurring keywords “video game” and “X”, neglecting the semantics of the relation “Soundtrack\_Of”. This issue is more serious in the *partially supportive* category, where many

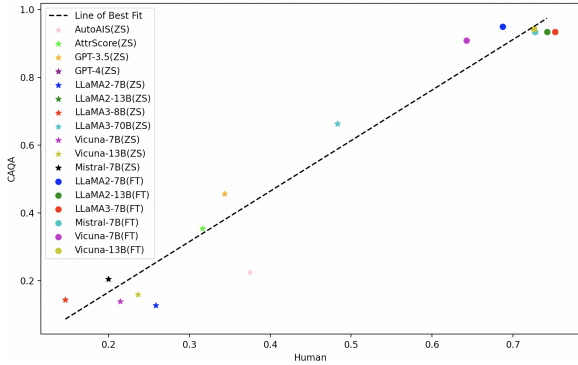


Figure 2: Micro-F1 scores of the evaluators on CAQA based on (1) CAQA’s own categories (y-axis) and (2) human-annotated categories (x-axis).

instances are incorrectly identified as *supportive*. This misclassification reduces the true positives for *partially supportive* cases while increasing false positives for *supportive* ones. Moreover, apart from GPT models, most other evaluators struggle to understand reasoning logic that requires integrating multiple pieces of evidence to prove the answer. The observation can be verified by the fact that all the evaluators usually perform worse on complex attribution with *concatenation*, *intersection* and *union* and on less complex attribution.

## 6.2 Consistency with Human Annotations

We assess the consistency between the automatically generated categories in CAQA and new categories annotated by humans. With these two sets of categories as gold standard, we compute and compare the overall micro-F1 scores of each evaluator. Some of the results as shown in Figure 2. The results demonstrate that the evaluators mostly have consistent results between the two kinds of categories, with a Pearson correlation coefficient of 0.97. This confirms that evaluation using our automatically generated categories closely align with manual evaluation.

## 6.3 Fine-grained Evaluation on Partially Supportive

Our CAQA benchmark enables a more fine-grained evaluation of attribution, particularly for partially supportive cases, compared to existing benchmarks. To achieve this, we employ the metric FACTSCORES (Min et al., 2023), which is the proportion of sub-facts in the answer that are supported by the citations. For each partially supportive case, human annotators identify the supported sub-facts, allowing us to

Table 6: FACTSCORES results on 200 partially supportive samples in CAQA. Hum-Gap measures the discrepancy of FACTSCORES w.r.t. human annotations (i.e., 0.58). Fine-tuning uses the training set of CAQA.

Evaluators		FACTSCORES	Hum-Gap
<b>Zero-Shot</b>	LLaMA-3 (70B)	0.85	0.27
	GPT-3.5-turbo	0.93	0.35
	GPT-4	0.84	0.26
<b>Fine-Tuning</b>	LLaMA-3 (8B)	0.19	0.39
	Vicuna (7B)	0.19	0.39
	Vicuna (13B)	0.18	0.40
AUTOIS (11B)		0.44	0.14
ATTRSCORE (13B)		0.25	0.33
CAQA Annotations		0.62	<b>0.04</b>
Human Annotations		0.58	-

calculate FACTSCORES manually. Additionally, CAQA can automatically compute FACTSCORES by comparing the triples in the initial subgraph with those in the subgraph after deletion. Various automatic evaluators, including LLM evaluators, as well as existing methods such as AutoIS and AttrScore, generate their FACTSCORES by first converting the triples in the initial subgraph  $\mathcal{G}$  into natural language sub-facts using ChatGPT. Then, we apply the Retrieve→LM method (Min et al., 2023) to obtain FACTSCORES (see details in Appendix D.2).

The results in Table 6 reveal a significant performance gap between the selected evaluators and human annotations. The three evaluators fine-tuned on CAQA, which support four categories, and AttrScore, which identifies three, exhibit worse performance (much higher Hum-Gap) compared to AutoIS which identifies only two categories. Meanwhile, the evaluators in the zero-shot setting have much higher FACTSCORES than human annotations, which reflects the finding in Section 6.1, i.e., their attribution is impacted by keyword co-occurrence in sub-facts and citations. Additionally, the FACTSCORES of the automated annotations in CAQA differ from human annotations by only 4%, demonstrating that CAQA is a reliable benchmark for automated fine-grained evaluation.

## 6.4 Exploration of Out-of-Domain Data

We compare the existing evaluators AutoIS (based on T5-11B) and AttrScore (based on Vicuna-13B) that are trained by some other benchmarks, with T5-11B and Vicuna-13B fine-tuned by CAQA (denoted as T5-11B\* and Vicuna-13B\*), by testing them on OOD benchmark ALCE-FineGrained. We also further fine-tune Vicuna-13B\* with a small number of samples in ALCE-FineGrained. The



Table 7: Performance on ALCE-FineGrained by (1) T5-11B\* and Vicuna-13B\* and (2) AutoIS and AttrScore.

Evaluators	ALCE-FineGrained				Overall
	Sup.	Non-Sup.			
AutoIS (T5-11B)	0.31	0.65		0.54	
T5-11B*	<b>0.44</b>	<b>0.72</b>		<b>0.63</b>	
	Sup.	Par.	Con.	Irr.	Overall
AttrScore (Vicuna-13B)	0.52	-	0.21	0.42	0.36
Vicuna-13B*	0.54	0.24	0.30	0.34	0.38
Vicuna-13B* Fine-Tuning	<b>0.69</b>	<b>0.36</b>	<b>0.40</b>	<b>0.46</b>	<b>0.52</b>

results are in Table 7. Note for comparison with AutoIS, we merge the three negative categories into Non-Supportive. Compared to AutoIS and AttrScore, T5-11B\* and Vicuna-13B\*, have competitive performance in individual classes and the overall score, demonstrating the effectiveness of CAQA for developing attribution evaluators. Table 7 also verifies that fine-tuning with a few samples of the domain of the testing dataset further improve the evaluators. See more details in Appendix F.

## 7 Conclusion and Future Work

This work has advanced the evaluation of AQA by proposing a new benchmark CAQA containing comprehensive attribution categories and different attribution complexity levels, and an automatic and general benchmark construction method. Extensive evaluation has been conducted over 25 automatic evaluators, including two specifically developed baselines, which has demonstrated the effectiveness of CAQA for evaluator assessment and development, and leads to a series of important findings in AQA, such as the shortage of logical reasoning capabilities in dealing with complex attribution scenario. In the future, we will study more robust attribution evaluators with CAQA and its variants from other KGs, supporting logical reasoning.

## Limitations

Although we consider comprehensive attribution categories and complex reasoning scenarios, some more situations including lengthy answers and citations, mathematical, temporal and spatial reasoning, and reasoning requiring domain specific knowledge could be considered in the future. For illustration, consider the question: “When did England last reach the quarterfinals of the World Cup?” The provided answer is “England last made the quarterfinals in 1990,” with a citation noting that “The England national football team finished in fourth place in 2018 and reached the semifinals in 1990.” To accurately attribute the answer, it is essential

to understand that finishing in fourth place implies participation in the quarterfinals and that 2018 is more recent than 1990. All these features could be reflected in new versions of our benchmark via considering corresponding KGs (or ontologies) and queries in our automatic construction method.

## Acknowledgments

This work is partially supported by National Nature Science Foundation of China under No. 62476058, and by the EPSRC project OntoEm (EP/Y017706/1). We gratefully acknowledge the Big Data Computing Center of Southeast University, the Edinburgh International Data Facility (EIDF), and the Data-Driven Innovation Programme at the University of Edinburgh for providing computational resources and support for the numerical experiments conducted in this paper.

## References

- Marah Abdin, Jyoti Aneja, Hany Awadalla, Ahmed Awadallah, Ammar Ahmad Awan, Nguyen Bach, Amit Bahree, Arash Bakhtiari, Jianmin Bao, Harkirat Behl, Alon Benhaim, Misha Bilenko, Johan Bjorck, Sébastien Bubeck, and Martin Cai etc. 2024. [Phi-3 technical report: A highly capable language model locally on your phone](#). *Preprint*, arXiv:2404.14219.
- Samuel Amouyal, Tomer Wolfson, Ohad Rubin, Ori Yoran, Jonathan Herzig, and Jonathan Berant. 2023. [QAMPARI: A benchmark for open-domain questions with many answers](#). In *Proceedings of the Third Workshop on Natural Language Generation, Evaluation, and Metrics (GEM)*, pages 97–110, Singapore. Association for Computational Linguistics.
- Akari Asai, Zeqiu Wu, Yizhong Wang, Avirup Sil, and Hannaneh Hajishirzi. 2024. [Self-rag: Learning to retrieve, generate, and critique through self-reflection](#). In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.
- Agnes Axelsson and Gabriel Skantze. 2023. [Using large language models for zero-shot natural language generation from knowledge graphs](#). In *Proceedings of the Workshop on Multimodal, Multilingual Natural Language Generation and Multilingual WebNLG Challenge (MM-NLG 2023)*, pages 39–54, Prague, Czech Republic. Association for Computational Linguistics.
- 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*.

- Kurt Bollacker, Colin Evans, Praveen Paritosh, Tim Sturge, and Jamie Taylor. 2008. Freebase: a collaboratively created graph database for structuring human knowledge. In *Proceedings of the 2008 ACM SIGMOD international conference on Management of data*, pages 1247–1250.
- 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](#).
- Luyu Gao, Zhuyun Dai, Panupong Pasupat, Anthony Chen, Arun Tejasvi Chaganty, Yicheng Fan, Vincent Zhao, Ni Lao, Hongrae Lee, Da-Cheng Juan, et al. 2023a. Rarr: Researching and revising what language models say, using language models. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 16477–16508.
- Tianyu Gao, Howard Yen, Jiatong Yu, and Danqi Chen. 2023b. [Enabling large language models to generate text with citations](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023*, pages 6465–6488. Association for Computational Linguistics.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, Aurelien Rodriguez, Austen Gregerson, Ava Spataru, Baptiste Roziere, Bethany Biron, Binh Tang, Bobbie Chern, Charlotte Caucheteux, Chaya Nayak, Chloe Bi, Chris Marra, Chris McConnell, and . Christian Keller etc. 2024. [The llama 3 herd of models](#). *Preprint*, arXiv:2407.21783.
- Yu Gu, Sue Kase, Michelle Vanni, Brian Sadler, Percy Liang, Xifeng Yan, and Yu Su. 2021. Beyond iid: three levels of generalization for question answering on knowledge bases. In *Proceedings of the Web Conference 2021*, pages 3477–3488.
- Xiaoqi Han, Ru Li, Hongye Tan, Wang Yuanlong, Qinghua Chai, and Jeff Pan. 2023. [Improving sequential model editing with fact retrieval](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 11209–11224, Singapore. Association for Computational Linguistics.
- Jie He, Jennifer Neville, Mengting Wan, Longqi Yang, Hui Liu, Xiaofeng Xu, Xia Song, Jeff Z. Pan, and Pei Zhou. 2025. GenTool: Enhancing Tool Generalization in Language Models through Zero-to-One and Weak-to-Strong Simulation. In *In the Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (ACL 2025)*.
- Aidan Hogan, Eva Blomqvist, Michael Cochez, Claudia d’Amato, Gerard De Melo, Claudio Gutierrez, Sabrina Kirrane, José Emilio Labra Gayo, Roberto Navigli, Sebastian Neumaier, et al. 2021. Knowledge graphs. *ACM Computing Surveys (Csur)*, 54(4):1–37.
- Or Honovich, Roei Aharoni, Jonathan Herzig, Hagai Taitelbaum, Doron Kukliansy, Vered Cohen, Thomas Scialom, Idan Szpektor, Avinatan Hassidim, and Yossi Matias. 2022. [TRUE: Re-evaluating factual consistency evaluation](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3905–3920, Seattle, United States. Association for Computational Linguistics.
- Wenyu Huang, Guancheng Zhou, Hongru Wang, Pavlos Vougiouklis, Mirella Lapata, and Jeff Z. Pan. 2024. Less is More: Making Smaller Language Models Competent Subgraph Retrievers for Multi-hop KGQA. In *In the Findings of the Association for Computational Linguistics: EMNLP 2024 (EMNLP 2024 Findings)*, pages 15787–15803.
- Binyuan Hui, Jian Yang, Zeyu Cui, Jiayi Yang, Dayiheng Liu, Lei Zhang, Tianyu Liu, Jiajun Zhang, Bowen Yu, Keming Lu, Kai Dang, Yang Fan, Yichang Zhang, An Yang, Rui Men, Fei Huang, Bo Zheng, Yibo Miao, Shanghaoran Quan, Yunlong Feng, Xingzhang Ren, Xuancheng Ren, Jingren Zhou, and Junyang Lin. 2024. [Qwen2.5-coder technical report](#). *Preprint*, arXiv:2409.12186.
- Gautier Izacard, Patrick Lewis, Maria Lomeli, Lucas Hosseini, Fabio Petroni, Timo Schick, Jane Dwivedi-Yu, Armand Joulin, Sebastian Riedel, and Edouard Grave. 2022. 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.
- Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. Mistral 7b. *arXiv preprint arXiv:2310.06825*.
- Longquan Jiang and Ricardo Usbeck. 2022. Knowledge graph question answering datasets and their generalizability: Are they enough for future research? In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 3209–3218.
- Ehsan Kamaloo, Aref Jafari, Xinyu Zhang, Nandan Thakur, and Jimmy Lin. 2023. Hagrid: A human-llm collaborative dataset for generative information-seeking with attribution. *arXiv preprint arXiv:2307.16883*.

- Mojtaba Komeili, Kurt Shuster, and Jason Weston. 2021. Internet-augmented dialogue generation. *arXiv preprint arXiv:2107.07566*.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Haoteng Zhang, and Ion Stoica. 2023. Efficient memory management for large language model serving with pagedattention. *Proceedings of the 29th Symposium on Operating Systems Principles*.
- Peter Lee, Sebastien Bubeck, and Joseph Petro. 2023. Benefits, limits, and risks of gpt-4 as an ai chatbot for medicine. *New England Journal of Medicine*, 388(13):1233–1239.
- Dongfang Li, Zetian Sun, Xinshuo Hu, Zhenyu Liu, Ziyang Chen, Baotian Hu, Aiguo Wu, and Min Zhang. 2023. A survey of large language models attribution. *arXiv preprint arXiv:2311.03731*.
- Xinze Li, Yixin Cao, Liangming Pan, Yubo Ma, and Aixin Sun. 2024a. Towards verifiable generation: A benchmark for knowledge-aware language model attribution. In *Findings of the Association for Computational Linguistics, ACL 2024, Bangkok, Thailand and virtual meeting, August 11-16, 2024*, pages 493–516. Association for Computational Linguistics.
- Yifei Li, Xiang Yue, Zeyi Liao, and Huan Sun. 2024b. Attributionbench: How hard is automatic attribution evaluation? In *Findings of the Association for Computational Linguistics, ACL 2024, Bangkok, Thailand and virtual meeting, August 11-16, 2024*, pages 14919–14935. Association for Computational Linguistics.
- Nelson F. Liu, Tianyi Zhang, and Percy Liang. 2023. Evaluating verifiability in generative search engines. In *Findings of the Association for Computational Linguistics: EMNLP 2023, Singapore, December 6-10, 2023*, pages 7001–7025. Association for Computational Linguistics.
- Chaitanya Malaviya, Subin Lee, Sihao Chen, Elizabeth Sieber, Mark Yatskar, and Dan Roth. 2024. ExpertQA: Expert-curated questions and attributed answers. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 3025–3045, Mexico City, Mexico. Association for Computational Linguistics.
- 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*.
- Grégoire Mialon, Roberto Dessì, Maria Lomeli, Christoforos Nalmpantis, Ram Pasunuru, Roberta Raileanu, Baptiste Rozière, Timo Schick, Jane Dwivedi-Yu, Asli Celikyilmaz, et al. 2023. Augmented language models: a survey. *arXiv preprint arXiv:2302.07842*.
- Dehai Min, Nan Hu, Rihui Jin, Nuo Lin, Jiaoyan Chen, Yongrui Chen, Yu Li, Guilin Qi, Yun Li, Nijun Li, and Qianren Wang. 2024. Exploring the impact of table-to-text methods on augmenting llm-based question answering with domain hybrid data. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies: Industry Track, NAACL 2024, Mexico City, Mexico, June 16-21, 2024*, pages 464–482. Association for Computational Linguistics.
- Dehai Min, Zhiyang Xu, Guilin Qi, Lifu Huang, and Chenyu You. 2025. Unihgrk: Unified instruction-aware heterogeneous knowledge retrievers. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL 2025 - Volume 1: Long Papers, Albuquerque, New Mexico, USA, April 29 - May 4, 2025*, pages 4577–4594. Association for Computational Linguistics.
- Sewon Min, Kalpesh Krishna, Xixi Lyu, Mike Lewis, Wen-tau Yih, Pang Wei Koh, Mohit Iyyer, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2023. Factscore: Fine-grained atomic evaluation of factual precision in long form text generation. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023*, pages 12076–12100. Association for Computational Linguistics.
- Hadiseh Moradisani, Fattane Zarrinkalam, Julien Serbanescu, and Zeinab Noorian. 2024. Unanswerwgen: A systematic approach for generating unanswerable questions in machine reading comprehension. In *Proceedings of the 2024 Annual International ACM SIGIR Conference on Research and Development in Information Retrieval in the Asia Pacific Region, SIGIR-AP 2024, Tokyo, Japan, December 9-12, 2024*, pages 280–286. ACM.
- Benjamin Muller, John Wieting, Jonathan H Clark, Tom Kwiatkowski, Sebastian Ruder, Livio Baldini Soares, Roei Aharoni, Jonathan Herzig, and Xinyi Wang. 2023. Evaluating and modeling attribution for cross-lingual question answering. *arXiv preprint arXiv:2305.14332*.
- 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*.
- Jeff Z. Pan, Simon Razniewski, Jan-Christoph Kalo, Sneha Singhania, Jiaoyan Chen, Stefan Dietze, Hajira Jabeen, Janna Omelivanenko, Wen Zhang, Matteo Lissandrini, Russa Biswas, Gerard de Melo, Angela Bonifati, Edlira Vakaj, Mauro Dragoni, and Damien Graux. 2023. Large Language Models and Knowledge Graphs: Opportunities and Challenges. *Transactions on Graph Data and Knowledge*, pages 1–38.



- J.Z. Pan, D. Calvanese, T. Eiter, I. Horrocks, M. Kifer, F. Lin, and Y. Zhao, editors. 2017a. *Reasoning Web: Logical Foundation of Knowledge Graph Construction and Querying Answering*. Springer.
- J.Z. Pan, G. Vetere, J.M. Gomez-Perez, and H. Wu, editors. 2017b. *Exploiting Linked Data and Knowledge Graphs for Large Organisations*. Springer.
- Hannah Rashkin, Vitaly Nikolaev, Matthew Lamm, Lora Aroyo, Michael Collins, Dipanjan Das, Slav Petrov, Gaurav Singh Tomar, Iulia Turc, and David Reitter. 2023. Measuring attribution in natural language generation models. *Computational Linguistics*, pages 1–66.
- Hongyu Ren, Mikhail Galkin, Michael Cochez, Zhaocheng Zhu, and Jure Leskovec. 2023. Neural graph reasoning: Complex logical query answering meets graph databases. *arXiv preprint arXiv:2303.14617*.
- 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*.
- Yongliang Shen, Kaitao Song, Xu Tan, Dongsheng Li, Weiming Lu, and Yueting Zhuang. 2023. Hugging-gpt: Solving ai tasks with chatgpt and its friends in huggingface. *arXiv preprint arXiv:2303.17580*.
- Zhili Shen, Chenxin Diao, Pavlos Vougiouklis, Pascual Merita, Shriram Piramanayagam, Enting Chen, Damien Graux, Andre Melo, Ruofei Lai, Zeren Jiang, Zhongyang Li, YE QI, Yang Ren, Dandan Tu, and Jeff Z. Pan. 2025. GeAR: Graph-enhanced Agent for Retrieval-augmented Generation. In *In the Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (ACL 2025)*.
- Weijia Shi, Sewon Min, Michihiro Yasunaga, Minjoon Seo, Richard James, Mike Lewis, Luke Zettlemoyer, and Wen-tau Yih. 2024. **REPLUG: retrieval-augmented black-box language models**. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, NAACL 2024, Mexico City, Mexico, June 16-21, 2024, pages 8371–8384. Association for Computational Linguistics.
- Kurt Shuster, Spencer Poff, Moya Chen, Douwe Kiela, and Jason Weston. 2021. **Retrieval augmentation reduces hallucination in conversation**. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 3784–3803, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, Johan Ferret, Peter Liu, Pouya Tafti, Abe Friesen, Michelle Casbon, Sabela Ramos, Ravin Kumar, Charline Le Lan, and Sammy Jerome etc. 2024. **Gemma 2: Improving open language models at a practical size**. *Preprint*, arXiv:2408.00118.
- Chaojun Wang and Rico Sennrich. 2020. **On exposure bias, hallucination and domain shift in neural machine translation**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3544–3552, Online. Association for Computational Linguistics.
- Hongru Wang, Wenyu Huang, Yufei Wang, Yuanhao Xi, Jianqiao Lu, Huan Zhang, Nan Hu, Zeming Liu, Jeff Z Pan, and Kam-Fai Wong. 2025. Rethinking stateful tool use in multi-turn dialogues: Benchmarks and challenges. *arXiv preprint arXiv:2505.13328*.
- Hongru Wang, Yujia Qin, Yankai Lin, Jeff Z. Pan, and Kam-Fai Wong. 2024. Empowering Large Language Models: Tool Learning for Real-world Interaction. In *In Proc. of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '24)*, pages 2983 – 2986.
- Yijun Xiao and William Yang Wang. 2021. **On hallucination and predictive uncertainty in conditional language generation**. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 2734–2744, Online. Association for Computational Linguistics.
- Wen-tau Yih, Matthew Richardson, Christopher Meek, Ming-Wei Chang, and Jina Suh. 2016. The value of semantic parse labeling for knowledge base question answering. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 201–206.
- Xiang Yue, Boshi Wang, Ziruo Chen, Kai Zhang, Yu Su, and Huan Sun. 2023. **Automatic evaluation of attribution by large language models**. In *Findings of the Association for Computational Linguistics: EMNLP 2023, Singapore, December 6-10, 2023*, pages 4615–4635. Association for Computational Linguistics.
- Weijia Zhang, Mohammad Aliannejadi, Yifei Yuan, Jiahuan Pei, Jia-Hong Huang, and Evangelos Kanoulas. 2024. **Towards fine-grained citation evaluation in generated text: A comparative analysis of faithfulness metrics**. In *Proceedings of the 17th International Natural Language Generation Conference, INLG 2024, Tokyo, Japan, September 23 - 27, 2024*, pages 427–439. Association for Computational Linguistics.

## A Generation of Natural Language Questions, Answers and Attributions

This section presents examples for generating natural language questions, answers, and citations using the GPT-3.5-turbo model. Our approach involves using distinct prompts for transforming subgraphs



into comprehensible natural language citations, extending original questions, and converting answer entities into detailed answer statements. Table 9 demonstrates the conversion of knowledge graph subgraphs into natural language citations. Table 10 illustrates the example of generating the extended question. Table 11 provides an example of how answer entities are transformed into long-form answer statements.

Additionally, we verify the quality of the KG-to-natural language text generated by ChatGPT. We ensure the quality of the natural language text generated by focusing on grammatical coherence and content accuracy. Building on prior work (Axelsson and Skantze, 2023), which demonstrates that ChatGPT excels in KG-to-Text tasks with high grammatical correctness and coherence, we primarily evaluate content accuracy to ensure consistency with the corresponding triples. Specifically, we adopted the evaluation framework outlined in the work (Axelsson and Skantze, 2023), assessing whether the generated text accurately reflects the triples. We randomly sampled 100 examples (half of the examples are from generated natural language citations, and the other half are from generated extended questions) and employed two independent annotators to label each instance according to one of three exclusive categories: (1) **Full Coverage**: The text fully and correctly states all triples. (2) **Absent**: The text misses some triples. (3) **Hallucinated**: The text introduces content that actively contradicts the triples. The results are shown in Table 8.

Annotator	Full Coverage	Absent	Hallucinated
A	92	5	3
B	95	3	2

Table 8: The human evaluation results of the KG-to-natural language text generated by ChatGPT.

The Cohen’s Kappa score is 0.758, indicating substantial agreement between two annotators. The annotation results demonstrate that ChatGPT reliably generates accurate and coherent text within our benchmark.

## B Concrete Cases of Query Extension

We incorporate concrete examples alongside explanations of each query extension strategy described in Section 4.2. These examples correspond to the rules summarized in Table 3 and illustrate the process of each extension strategy.

**Instruction:** Your task is to convert a specific subgraph from the knowledge graph into one or more coherent sentences that summarize the information encapsulated within the subgraph.

**Subgraph:** [{"Wii/DS NA", "type", "computer\_game\_region"}, {"LostWinds: Winter of the Melodias", "type", "game\_version"}, {"LostWinds: Winter of the Melodias", "game\_version.regions", "Wii/DS NA"}, {"LostWinds: Winter of the Melodias", "game\_version.distributed\_through", "WiiWare"}, {"Frontier Developments", "cvg\_developer.game\_versions\_developed", "LostWinds: Winter of the Melodias"}].

**Sentences:** *LostWinds: Winter of the Melodias is a computer game with a specific version tailored for the Wii/DS NA region. This version is distributed through WiiWare. The game was developed by Frontier Developments, showcasing their involvement in creating various game versions.*

Table 9: An example about converting a subgraph to a natural language citation using GPT-3.5-turbo.

**Instruction:** Given knowledge graph triples, your task is to generate a question using all the triples. The generated questions should contain all the relationships.

# Extended Triples

**Triples:** [(?x, type, cvg.computer\_videogame), (?x, computer\_videogame.influenced\_by, Sengoku Rance), (?x, fictional.setting, Touhou Project)]

**Extended question:** *What computer video game was influenced by Sengoku Rance and is set in the Touhou Project fictional universe?*

Table 10: An example about generating the extended question using GPT-3.5-turbo.

**Case 1: Single-triple Query with the Union Extension Strategy.** The grounded graph of the original query  $l$ : (Rick Scott (Q439729), place of birth, Bloomington). By querying the KG for other entities with the same name (e.g., Rick Scott), we may discover an additional entity Rick Scott (Q7329369) and retrieve its place of birth: (Rick Scott (Q7329369), place of birth, Stone Mountain). The grounded graph of the extended query  $l'$  after Union extension becomes: [(Rick Scott (Q439729), place of birth, Bloomington), (Rick Scott (Q7329369), place of birth, Stone Mountain)].

**Case 2: Path-like Query with Intersection Extension Strategies.** The grounded graph of the original query  $l$ : [(Rick Scott, place of birth, Bloomington), (Bloomington, capital of, McLean

**Instruction:** Your task is to convert a question along with its concise answer into a comprehensive answer statement.

**Question:** What group fought in the Battle of Vicksburg that was based in Montgomery?

**Answer:** Army of Mississippi

**Answer statement:** The group that fought in the Battle of Vicksburg and was based in Montgomery was the Army of Mississippi.

Table 11: An example about converting the answer entity to a long answer statement using GPT-3.5-turbo.

County)].

*First Intersection Strategy (Head Entity Extension):* We extend the query by adding a triple related to the head entity Rick Scott, e.g.: (Anna Scott, spouse, Rick Scott). Then the grounded graph of the extended query  $l'$ : [(Anna Scott, spouse, Rick Scott), (Rick Scott, place of birth, Bloomington), (Bloomington, capital of, McLean County)].

*Second Intersection Strategy (Answer Entity Extension):* We extend based on the answer entity Bloomington, adding: (McLean County, country, United States).

Then the grounded graph of the extended query  $l'$  becomes: [(Rick Scott, place of birth, Bloomington), (Bloomington, capital of, McLean County), (McLean County, country, United States)].

**Case 3: Tree-like Query (Intersection Extension Strategies).** The grounded graph of the original query  $l$ : [(Rick Scott, educated at, University of Missouri–Kansas City), (Sharice Davids, educated at, University of Missouri–Kansas City)].

*First Intersection Strategy (Head Entity Extension):* Add a triple involving the head entity Rick Scott, such as: (Anna Scott, spouse, Rick Scott). Then the grounded graph of the extended query  $l'$ : [(Anna Scott, spouse, Rick Scott), (Rick Scott, educated at, University of Missouri–Kansas City), (Sharice Davids, educated at, University of Missouri–Kansas City)].

*Second Intersection Strategy (Answer Entity Extension):* Extend the query by adding a triple about the answer entity University of Missouri–Kansas City: (University of Missouri–Kansas City, located in, Kansas City). Then the grounded graph of the extended query  $l'$ : [(Rick Scott, educated at, University of Missouri–Kansas City), (Sharice Davids, educated at, University of Missouri–Kansas City),

(University of Missouri–Kansas City, located in, Kansas City)].

## C CAQA Benchmark Construction and Statistics

The CAQA benchmark is built on the top of two KGQA datasets, GrailQA<sup>3</sup> and WebQuestionsSP, with the knowledge graph Freebase, forming a comprehensive attribution evaluation testbed. We selectively include samples from these two datasets whose logical queries align with single-triple, path-like, or tree-like queries, as delineated in Section 4.1. For path queries, we collect the example with a path length of at most two hops. We treat paths incorporating CVT (Compound Value Type) nodes as one-hop. For example, [(Harper Lee, person.education ?cvt), (?cvt education.institution, Monroe County High School)] is a one-hop path, where the node ?cvt holds no actual meaning. Regarding tree-liked queries, we restrict our selection to those with a maximum of two non-answer nodes, meaning up to two subject entities.

The length distribution (i.e., the number of tokens) of citations in the training and test sets of the CAQA benchmark is depicted in Figures 3 and 4. These distributions reveal a concentration of citations around 25 tokens, with a minority exceeding 60 tokens. In future work, we aim to enhance the complexity and length of natural language references by developing more intricate subgraphs. Additionally, Figure 5 presents the domain distribution within the CAQA benchmark. This distribution underscores the benchmark’s broad domain coverage and its encompassment of various sub-domains, highlighting the diversity of our benchmark.

## D Implementation Details

### D.1 Details of Evaluators and Prompts

AutoIS is a natural language inference (NLI) model<sup>4</sup> based on T5-11B that outputs a “1” to indicate that the citation supports the answer statement or a “0” to indicate a lack of support. AttrScore is a uniform name for attribution evaluators developed on various LLMs, and we use the best-performing attribution evaluator (Vicuna-13B) on the original work for comparison. Since AutoIS can only recognise *supportive* and *non-supportive* attribution categories, we only report its F1 score on *supportive* in Table 5. In the experiments on the

<sup>3</sup>GrailQA is licensed under CC BY-SA 4.0.

<sup>4</sup>[https://huggingface.co/google/t5\\_xxl\\_true\\_nli\\_mixture](https://huggingface.co/google/t5_xxl_true_nli_mixture)

### GPT series

**Instruction:** Your task is to evaluate the relationship between a provided citation and the answer to a specific question. There are four possible types of relationships:

1. Supportive: Choose this if the citation directly confirms or is fully in alignment with the answer, providing all necessary information to substantiate it.
2. Insufficient: Choose this when the citation provides only partial backing for the answer, lacking some essential details or evidence needed for full support.
3. Contradictory: Choose this option if the citation is consistent with the intent of the question but directly opposes or contradicts the answer.
4. Irrelevant: Select this option if the citation does not match the intent of the question and contains information that is not useful for answering.

For each example provided: First, you need to look at the question given and the answer provided. Then, compare them with the content of the citation. Finally, select the appropriate relationship category based on whether the citation supports the answer, is missing information, contradicts itself, or is irrelevant to the answer.

Example:

**Question:** {question}

**Answer:** {answer statement}

**Reference:** {citation}

**Relationship Category:**

---

### AttrScore

**premise:** {question|answer statement}

**hypothesis:** {citation}

---

### AutoIS

Below is an instruction that describes a task, paired with an input that provides further context. Write a response that appropriately completes the request.

**Instruction:** Verify whether a given reference can support the claim. Options: Attributable, Extrapolatory or Contradictory.

**Claim:** {question|answer statement}

**Reference:** {citation}

**Response:**

---

### Other Evaluators

Below is an instruction that describes a task, paired with an input that provides further context. Write a response that appropriately completes the request.

**Instruction:** Verify whether a given reference can support the claim. Options: Supportive, Insufficient, Contradictory or Irrelevant.

**Claim:** {question|answer statement}

**Reference:** {citation}

**Response:**

Table 12: Different prompts designed for different evaluators.

ALCE-FineGrained benchmark, to be able to compare the evaluator trained on our benchmark with AutoIS, we merge the three incorrect categories into the *non-supportive* category, and then compute F1 scores of *supportive* and *non-supportive* as well as overall micro-F1 score.

Table 12 describes the different prompt designs against the various attribution evaluators. During our experiments, we have explored providing more detailed explanations for each category to non-GPT models. However, we observed that overly detailed prompts often led to a decline in their performance. To maximize the performance of all models and to select the best attribution evaluator, we opted for prompts that were calibrated to achieve the best results for all models.

In the few-shot setting, we select one sample per

attribution category as a demonstration, as shown in Table 13. We explore on more few-shot settings in Appendix E. For model fine-tuning, we use the prompt of “Other Evaluators” depicted in Table 12 as input of all models, and the output of the model is one of the four attribution categories proposed. We use two A100 80G GPUs for full parameter fine-tuning and two A100 80G GPU for the inference phase (four A100 80G GPUs for inference with model parameters greater than 70B). During inference, text generation is conducted with a temperature setting of 0 and vLLM (Kwon et al., 2023) is used to acceleration. If LLMs produce an attribution category with an explanation, we extract the predicted label using regular expression techniques. Due to high costs, we evaluate GPT models on only 3000 test samples.

### GPT-3.5 series

**Instruction:** Your task is to evaluate the relationship between a provided citation and the answer to a specific question. There are four possible types of relationships:

1. Supportive: Choose this if the citation directly confirms or is fully in alignment with the answer, providing all necessary information to substantiate it.
2. Insufficient: Choose this when the citation provides only partial backing for the answer, lacking some essential details or evidence needed for full support.
3. Contradictory: Choose this option if the citation is consistent with the intent of the question but directly opposes or contradicts the answer.
4. Irrelevant: Select this option if the citation does not match the intent of the question and contains information that is not useful for answering.

Please read the examples and choose the most appropriate relationship category for the test example.

Example 1: {Support Example}

Example 2: {Missing Example}

Example 3: {Contradictory Example}

Example 4: {Irrelevant Example}

Test Example:

**Question:** {question}

**Answer:** {answer statement}

**Reference:** {citation}

**Relationship Category:**

### Other Evaluators

Below is an instruction that describes a task, paired with an input that provides further context. Write a response that appropriately completes the request.

**Instruction:** Verify whether a given reference can support the claim. Options: Supportive, Insufficient, Contradictory or Irrelevant.

{Support Example}

{Missing Example}

{Contradictory Example}

{Irrelevant Example}

**Claim:** {question|answer statement}

**Reference:** {citation}

**Response:**

Table 13: Different few-shot prompts designed for different evaluators.

## D.2 Implementations of Various Evaluators for Calculating FACTSCORES

We describe how the various attribution evaluators calculate FACTSCORES in the CAQA benchmark. For each *partially supportive* case, we first use GPT-3.5 to convert triples into natural language subfacts with the prompt: “Your task is to convert a triple into natural language statement”. Then, following the Retrieve→LM method (Min et al., 2023), we calculate the FACTSCORES for each attribution evaluator. Specifically, the prompt is fed into the evaluator, which predicts True or False to calculate the FACTSCORES. For the zero-shot evaluator, we use the prompt: “*Judge this fact based on the given context.*\n\n*Fact: {sub-fact}*\n\n*Text: {citation}*\n\n*True or False?*\n\n*Output:*”. For fine-tuned and existing evaluators, the prompt provided in Table 12 is used. Because these evaluators can generate more than two attribution categories, we categorize supportive as True and all other categories as False for calculating the FACTSCORES. Human annotation, as described in Appendix H,

N-shot (GPT-3.5-turbo)	CAQA				
	Sup.	Par.	Con.	Irr.	Overall
1-shot	0.613	0.026	0.318	<b>0.609</b>	0.476
2-shot	<b>0.627</b>	<b>0.034</b>	0.359	0.593	<b>0.486</b>
3-shot	0.599	0.015	<b>0.378</b>	0.581	0.478

Table 14: The performance of GPT-3.5-turbo under various few-shot settings on CAQA.

involves annotators determining whether each subfact is supported by its citation. The FACTSCORES is the proportion of predictions classified as True compared to the total number of subfacts evaluated.

## E Detailed Experimental Results

### E.1 Full Experimental Results

We present the full experimental results in Tables 16. We find that the newest LLMs demonstrate slight improvements in overall performance for attribution recognition compared to their earlier versions in all three settings. Specifically, in zero-shot and few-shot settings, they show more balanced and enhanced results across scenarios with varying levels of complexity, highlighting advancements



in their reasoning capabilities. However, they still struggle to accurately distinguish fine-grained attribute categories, with notable deficiencies persisting in certain attribution recognition.

## E.2 Exploration of Few-shot Settings

We investigate three few-shot settings: 1-shot, 2-shot, and 3-shot in 5,000 test instances employing GPT-3.5-turbo. In these settings, 1, 2, and 3 examples, respectively, are provided for each attribution category. The outcomes, as displayed in Table 14, suggest that increasing the number of examples yields negligible improvement in performance. Consequently, considering the associated costs, we have opted to use the 1-shot setting in all subsequent experiments.

## E.3 Exploration of Advanced Baselines

We have included more advanced and diverse baselines, particularly incorporating LLM-based evaluators with CoT strategies. Specifically, inspired by the approach proposed in (Min et al., 2023), we designed a multi-step CoT evaluation strategy that enhances the model’s attribution evaluation capabilities:

1. Decomposition: The LLM is first prompted to break down the answer into a series of atomic facts. The prompt is: “*Break the following sentence into independent sub-facts. If the sentence is already a minimal fact, return it directly. Each sub-fact should be concise, self-contained, and listed on a new line. \n Sentence: Sentence\n Sub-facts:*”.

2. Atomic Fact Attribution Evaluation: For each atomic fact, the LLM assesses its relationship with the evidence, categorizing it as supportive, contradictory, or irrelevant. The prompt is: “*Evaluate the relationship between the statement and reference. Choose one of the following categories based on how the citation relates to the statement:\n 1. Supportive: Select this if the statement is present and validated by the reference.\n 2. Contradictory: Select this if the reference contains the fact that directly opposes or contradicts the statement.\n 3. Irrelevant: Select this if the reference is completely unrelated to the statement.\n For each example provided: First, you need to understand precisely what the statement claims. Then, understand what facts or information the reference actually provides. Finally, compare the statement with the information of the reference and select the appropriate relationship category.\n\n Example:\n \n Statement: sub\_fac\n Reference: reference\n*”.

*Relationship Category:*”.

3. Final Attribution Classification: Based on the attribution types of the atomic facts, we determine the overall attribution label for the answer as follows: Supportive: All atomic facts are supported by the evidence. Partially Supportive: At least one atomic fact is supported, and none are contradictory. Contradictory: At least one atomic fact is contradicted by the evidence. Irrelevant: All atomic facts are irrelevant to the evidence.

The experiment results are shown in table 15. Compared with the naive LLM evaluators (see Table 5), the experimental results of our proposed CoT-based evaluators reveal two key findings: 1. Effectiveness in Complex Scenarios: The CoT-based method demonstrates strong performance in complex attribution scenarios, such as Concatenation, Intersection, and Union, achieving up to a 16% improvement over the naive LLM evaluator in complex attribution scenarios. However, in the simpler Single scenario, decomposition errors will hinder performance, resulting in an overall score that remains largely unchanged in the zero-shot setting. 2. Improvement with Few-shot Learning: Under the few-shot setting, the CoT-based evaluator benefits significantly from in-context examples, which help guide both fact decomposition and attribution classification. Compared to the equally naive LLM evaluator, the overall performance of the CoT-based evaluator improved by as much as 5.2%, highlighting the value of in-context learning in improving the evaluation of attribution.

## F Details of Experiments on ALCE-FineGrained

ALCE-FineGrained consists of 215 manually labelled samples containing 104 supportive samples, 58 partially supportive samples, 25 contradictory samples, and 28 irrelevant samples. For the few-shot setting, we select one sample for each attribution category as demonstration. For the fine-tuning setting, we employ GPT-4 to annotate 800 samples from the ALCE benchmark as the training set. Since there are fewer contradictory and irrelevant attribution categories in the ALCE benchmark, we use GPT-4 to edit the evidence to construct contradictory and irrelevant samples, thus ensuring a balanced number of the four categories.

Table 18 presents two ALCE-FineGrained examples, illustrating the attribution categories *partially supportive* and *irrelevant*, respectively. It shows

Settings	Evaluators (Size)	Category					Complexity			
		Sup.	Par.	Con.	Irr.	Overall	S.	C.	I.	U.
Zero-Shot	Gemma-2 (27B)	0.663	0.216	0.55	0.597	0.556	0.557	0.667	0.542	0.375
	LLaMA-3.1 (70B)	0.695	0.166	0.54	0.58	0.537	0.527	0.681	0.542	0.5
	Qwen-2.5 (72B)	0.653	0.294	<b>0.75</b>	0.496	0.604	0.585	<b>0.744</b>	<b>0.639</b>	0.538
	GPT-4o	<b>0.746</b>	0.387	0.659	0.649	<b>0.624</b>	<b>0.654</b>	0.604	0.528	<b>0.688</b>
	GPT-4o-mini	0.698	<b>0.413</b>	0.593	<b>0.717</b>	0.607	0.642	0.5	0.517	<b>0.688</b>
Few-Shot	Gemma-2 (27B)	0.645	0.293	0.692	0.554	0.581	0.573	0.708	0.592	0.438
	LLaMA-3.1 (70B)	0.722	0.361	0.714	0.637	0.64	0.64	0.688	0.637	0.562
	Qwen-2.5 (72B)	0.755	0.449	0.771	0.505	0.646	0.636	<b>0.708</b>	<b>0.682</b>	0.5
	GPT-4o	<b>0.817</b>	<b>0.615</b>	0.719	0.679	<b>0.716</b>	<b>0.765</b>	0.542	0.617	<b>0.688</b>
	GPT-4o-mini	0.770	0.497	<b>0.786</b>	<b>0.706</b>	0.697	0.741	0.417	0.398	<b>0.688</b>

Table 15: Advanced baseline results on CAQA. The bold and underlined results show the best.

Settings	Evaluators (Size)	Category					Complexity				
		Sup.	Par.	Con.	Irr.	Overall	S.	C.	I.	U.	
Zero-Shot	LLaMA-2 (7B)	0.423	0.121	0.057	0.170	0.279	0.286	0.249	0.282	0.260	
	LLaMA-2 (13B)	0.418	0.164	0.161	0.125	0.279	0.314	0.270	0.303	0.253	
	LLaMA-3 (8B)	0.467	0.120	0.072	0.007	0.296	0.304	0.271	0.283	0.259	
	LLaMA-3.1 (8B)	0.544	0.049	0.130	0.017	0.318	0.319	0.326	0.319	0.285	
	LLaMA-3 (70B)	0.746	0.104	0.653	0.592	0.525	0.645	0.279	0.305	0.578	
	LLaMA-3.1 (70B)	0.688	0.168	0.547	0.609	0.544	0.545	0.549	0.545	0.499	
	Mistral-v0.2 (7B)	0.456	0.178	0.191	0.153	0.305	0.315	0.281	0.294	0.265	
	Mistral-v0.3 (7B)	0.661	0.160	0.051	0.334	0.362	0.363	0.374	0.356	0.337	
	Ministral* (8B)	-	-	-	-	-	-	-	-	-	
	Mixtral-v1.0 (8x7B)	0.677	0.094	0.17	0.635	0.494	0.495	0.516	0.484	0.487	
	Vicuna (7B)	0.513	0.100	0.064	0.199	0.327	0.343	0.273	0.312	0.256	
	Vicuna (13B)	0.634	0.211	0.393	0.275	0.405	0.432	0.314	0.361	0.374	
	Phi-3-small (7B)	0.624	0.217	0.481	0.569	0.507	0.509	0.508	0.504	0.453	
	Phi-3-medium (14B)	0.627	0.148	0.383	0.291	0.406	0.406	0.413	0.407	0.391	
	Gemma-2 (9B)	0.667	0.28	0.498	0.624	0.556	0.557	0.572	0.552	0.508	
	Gemma-2 (27B)	0.653	0.184	0.569	0.646	0.566	0.566	0.579	<u>0.566</u>	0.537	
	Qwen-2.5 (7B)	0.696	0.241	0.617	0.404	0.556	0.557	0.572	0.551	0.537	
	Qwen-2.5 (14B)	0.68	0.132	<u>0.708</u>	<u>0.66</u>	<u>0.617</u>	0.640	<b>0.622</b>	<b>0.611</b>	0.547	
	Qwen-2.5 (72B)	0.629	0.266	0.701	0.471	0.571	0.593	0.583	0.565	0.53	
	GPT-3.5-turbo	0.583	0.017	0.598	0.512	0.497	0.555	0.321	0.363	0.363	
GPT-4	<b>0.771</b>	<b>0.456</b>	<b>0.745</b>	0.473	<b>0.630</b>	<b>0.685</b>	0.451	0.514	<b>0.616</b>		
GPT-4o	<u>0.769</u>	<u>0.445</u>	0.598	0.626	<b>0.630</b>	<u>0.676</u>	<u>0.591</u>	0.47	<u>0.588</u>		
GPT-4o-mini	0.718	0.297	0.632	<b>0.703</b>	0.616	0.672	0.473	0.444	0.559		
Few-Shot	LLaMA-2 (7B)	0.300	0.066	0.009	0.334	0.248	0.259	0.218	0.167	0.308	
	LLaMA-2 (13B)	0.419	0.199	0.167	0.089	0.272	0.274	0.271	0.233	0.267	
	LLaMA-3 (8B)	0.573	0.202	0.234	0.156	0.336	0.356	0.279	0.310	0.294	
	LLaMA-3.1 (8B)	0.631	0.101	0.307	0.157	0.353	0.354	0.375	0.343	0.362	
	LLaMA-3 (70B)	0.741	0.192	0.608	0.584	0.531	0.628	0.295	0.314	0.563	
	LLaMA-3.1 (70B)	0.698	0.259	0.713	0.665	0.627	0.626	<u>0.638</u>	<u>0.632</u>	<u>0.589</u>	
	Mistral-v0.2 (7B)	0.412	0.152	0.041	0.415	0.349	0.339	0.278	0.300	0.271	
	Mistral-v0.3 (7B)	0.571	0.072	0.252	0.134	0.342	0.343	0.351	0.336	0.326	
	Ministral (8B)	0.646	0.114	0.405	0.498	0.463	0.461	0.475	0.47	0.453	
	Mixtral-v1.0 (8x7B)	0.589	0.084	0.356	0.558	0.455	0.456	0.473	0.451	0.416	
	Vicuna (7B)	0.578	0.183	0.081	0.324	0.325	0.337	0.272	0.354	0.311	
	Vicuna (13B)	0.633	0.208	0.383	0.288	0.403	0.427	0.315	0.397	0.374	
	Phi-3-small (7B)	0.592	0.185	0.301	0.498	0.445	0.444	0.455	0.448	0.416	
	Phi-3-medium (14B)	0.626	0.122	0.095	0.447	0.417	0.417	0.445	0.415	0.385	
	Gemma-2 (9B)	0.705	0.390	0.568	0.593	0.572	0.571	0.586	0.578	0.526	
	Gemma-2 (27B)	0.646	0.231	0.67	0.572	0.57	0.57	0.585	0.569	0.511	
	Qwen-2.5 (7B)	0.676	0.198	0.677	0.545	0.57	0.572	0.582	0.563	0.555	
	Qwen-2.5 (14B)	0.699	0.257	<b>0.741</b>	<u>0.676</u>	0.646	0.680	<b>0.656</b>	<b>0.638</b>	<b>0.608</b>	
	Qwen-2.5 (72B)	0.721	0.4	<u>0.736</u>	0.503	0.617	0.635	0.608	0.626	0.592	
	GPT-3.5-turbo	0.602	0.031	0.340	0.604	0.467	0.512	0.324	0.384	0.368	
GPT-4	<b>0.794</b>	<b>0.520</b>	0.728	0.653	<b>0.680</b>	<b>0.745</b>	0.492	0.473	0.559		
GPT-4o	<u>0.783</u>	<u>0.507</u>	0.683	0.641	<u>0.664</u>	0.73	0.559	0.449	0.529		
GPT-4o-mini	0.763	0.435	0.705	<b>0.7</b>	0.657	<u>0.741</u>	0.43	0.404	0.588		
Fine-Tuning	LLaMA-2 (7B)	0.922	0.897	<b>0.944</b>	<b>0.933</b>	0.926	0.923	0.815	0.931	0.921	
	LLaMA-2 (13B)	0.929	0.907	0.938	0.923	0.925	0.954	0.824	<b>0.936</b>	0.939	
	LLaMA-3 (8B)	0.935	0.901	0.935	0.928	0.926	0.935	0.820	0.930	0.924	
	LLaMA-3.1 (8B)	<b>0.946</b>	0.919	<u>0.944</u>	<u>0.934</u>	<u>0.941</u>	<u>0.953</u>	<u>0.850</u>	<b>0.939</b>	<b>0.945</b>	
	Mistral-v0.2 (7B)	0.927	0.908	<b>0.944</b>	0.849	0.882	0.935	0.831	0.921	0.905	
	Mistral-v0.3 (7B)	<u>0.944</u>	<u>0.921</u>	<b>0.947</b>	<b>0.935</b>	<b>0.942</b>	<b>0.956</b>	<b>0.852</b>	<u>0.937</u>	<u>0.941</u>	
	Vicuna (7B)	0.937	0.907	0.940	0.906	0.932	<b>0.956</b>	0.823	<b>0.936</b>	0.939	
	Vicuna (13B)	0.942	<b>0.923</b>	0.939	0.923	<b>0.933</b>	0.950	0.847	0.935	0.940	
	<i>Existing Attribution Evaluators</i>										
	AUTOIS (11B)	0.609	-	-	-	-	-	-	-	-	-
ATTRSCORE (13B)	0.687	-	0.523	0.541	0.521	0.559	0.410	0.432	0.353	-	

Table 16: Full results on CAQA. The bold and underlined results show the best and second best results respectively. Ministral behaves abnormally in the zero-shot setting and does not obey the instruction to select one of the four attribution categories.

that these two categories, which are not included in the previous attribution categories, are common and different in practical situations. In example 1, where the attribution category is *partially supportive*, most of the answer statement (highlighted in green) is mentioned in the citation, but the key information “*The Maryland Transportation Authority*” (highlighted in yellow) is not mentioned in the citation. This demonstrates the subtleties that can render an attribution insufficient. In example 2, which is categorised as *irrelevant*, the entirety of the answer statement is irrelevant to the citation. This exemplifies a clear case of irrelevant attribution. Notably, previous evaluators, AutoIS and AttrScore, are unable to accurately classify these cases. In contrast, Vicuna, an evaluator trained with our CAQA benchmark, successfully identifies the correct attribution categories. This underscores the effectiveness and practicality of employing the CAQA benchmark for developing attribution evaluators.

## G Analysis of Existing Attributed QA Systems

Following the work of Gao et al. (Gao et al., 2023b) we reproduce the attributed question answering system based on Vicuna-13B model, noted for its effectiveness in smaller language model configurations. Specifically, we provide the model with the top-3 retrieved passages and instruct the model to cite them accordingly. The retrieved passages and the instruction are consistent with the original implementation. Upon reviewing 234 instances of the system, our analysis revealed that: 44.4% of the instances accurately cited evidence supporting their answers, while 24.8% cited evidence that only partially supported the answers. Contradictory evidence was cited in 10.7% of cases, and 12.0% of the responses involved citations of irrelevant evidence. Additionally, 8.1% of the cases were categorized under other issues, including incomplete or unclear answers. The predominant challenges in incorrect attributions are identified as *partially supportive*, *contradictory*, and *irrelevant* citations, with *partially supportive* citations being the most common problem.

## H Human Annotation

The human annotation process for our study was conducted by the authors themselves, eliminating the need for external paid services. Three of our

annotators were asked to read these guidelines carefully. Only annotators with a thorough understanding of the guidelines and the task were allowed to participate in the manual evaluation. We ensured the reliability of the results by retaining only those annotations that were aligned across all three annotators. Annotation guidelines are shown in Fig. 6 and 7.

For the CAQA benchmark, we sampled 400 cases and provided them to three annotators. Among these, 2 examples were deemed unclear by more than one annotator and subsequently discarded. This resulted in 398 valid annotated examples. The Fleiss’ Kappa score for these annotations was 0.780, reflecting a substantial level of agreement. For the ALCE-FineGrained benchmark, we sampled 300 cases and provided them to three annotators. Of these, 24 examples were deemed unclear and were excluded, leaving 276 valid annotated examples. The Fleiss’ Kappa score for these annotations was 0.748, also indicating a substantial level of agreement.

Additionally, we analyzed the labeling process and identified patterns of agreement and disagreement among annotators. Overall, the category with the highest agreement was ‘supportive’, followed by ‘irrelevant’, while the categories ‘partially supportive’ and ‘contradictory’ exhibited lower agreement. Specifically, annotators often confused ‘partially supportive’ with ‘irrelevant’, frequently misclassifying examples between these two categories. This confusion arose because distinguishing between them required identifying sub-facts (i.e., the smallest semantic units containing a subject, verb, and object) from answer statements and evidence. The process often required domain-specific knowledge, which annotators lacked, leading them to rely on co-occurring keywords rather than deeper semantic understanding. Similarly, annotators had difficulty identifying ‘contradictory’ due to a lack of domain knowledge. This led to misclassifications of ‘contradictory’ as either ‘partially supportive’ or ‘irrelevant’. Determining contradictions requires nuanced reasoning beyond surface-level overlaps, which posed challenges in the absence of sufficient domain context.

## I License

Our dataset will be distributed under the CC BY-SA 4.0 license.

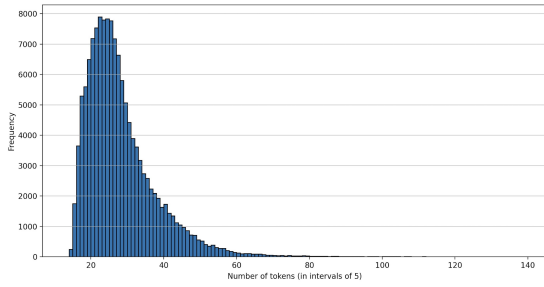


Figure 3: Histogram of the number of tokens across all citations in the CAQA benchmark training set.

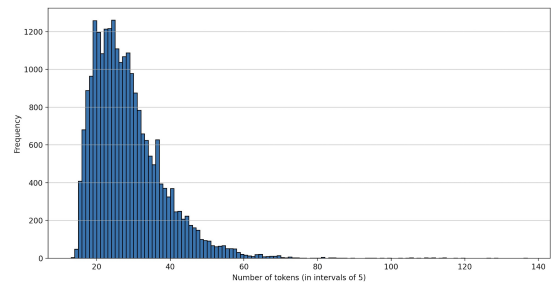


Figure 4: Histogram of the number of tokens across all citations in the CAQA benchmark test set.

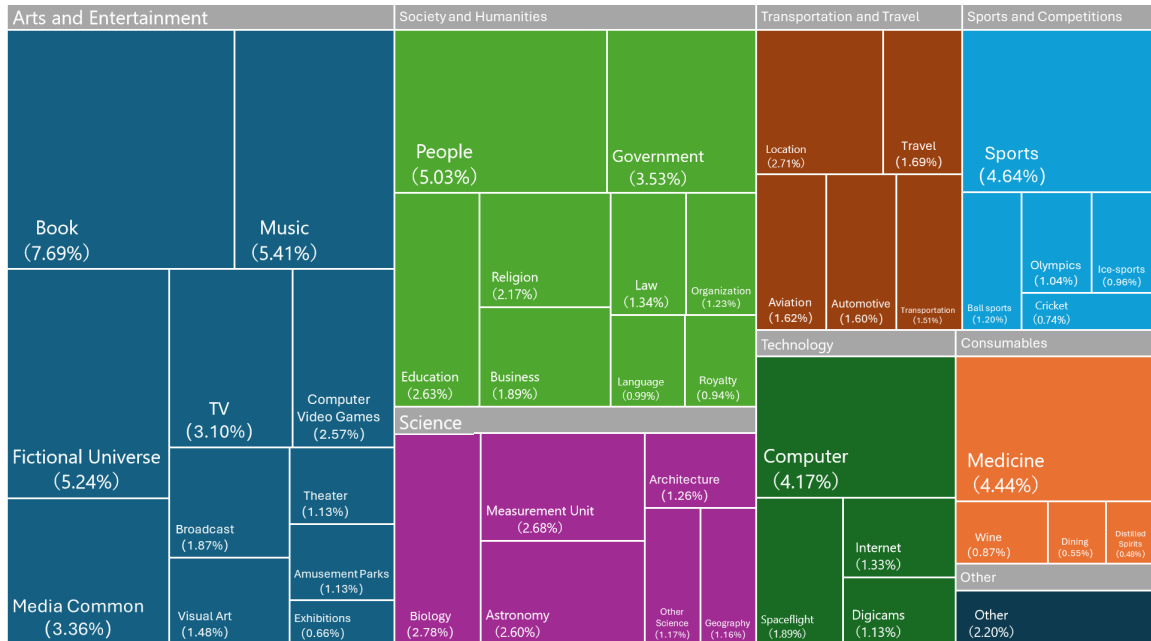


Figure 5: The distribution of examples across different domains in the CAQA benchmark.

Table 17: Examples of the four complexity types. Reasoning Graphs show the reasoning relationships between citations-answers. Green represents content associated with the answer, gray indicates excluded content, and orange indicates the common term connecting the citations.

Complexity	Examples	Reasoning Graphs
Single	<p><b>Question:</b> Which radio program episode appears in All Things Considered?</p> <p><b>Answer:</b> The radio program episode in which All Things Considered appears is Remorse: The 14 Stories of Eric Morse [1].</p> <p><b>Citations:</b> [1] <i>Remorse: The 14 Stories of Eric Morse is an episode of the radio program All Things Considered...</i></p>	
Union	<p><b>Question:</b> Which university did Rick Scott attend?</p> <p><b>Answer:</b> Rick Scott attended the University of Missouri–Kansas City and Southern Methodist University [1][2].</p> <p><b>Citations:</b> [1] <i>Rick Scott graduated from the University of Missouri–Kansas City ...</i> [2] <i>Rick Scott earned a juris doctor degree by working his way through Southern Methodist University, ...</i></p>	
Intersection	<p><b>Question:</b> The computer designer for Macintosh 128k and NeXT computer was whom?</p> <p><b>Answer:</b> The computer designer for Macintosh 128k and NeXT computer was Steve Jobs [1][2].</p> <p><b>Citations:</b> [1] <i>The computer designer for Macintosh 128k was Jerry Manock, who worked with Steve Jobs to develop the vertical body ...</i> [2] <i>...Several former Apple employees followed Jobs to NeXT, including Joanna Hoffman, Bud Tribble, George Crow, Rich Page...</i></p>	
Concatenation	<p><b>Question:</b> What are the official languages in the politician Mohammad Najibullah's country?</p> <p><b>Answer:</b> Pashto and Dari are the official languages in the politician Mohammad Najibullah's country. [1][2].</p> <p><b>Citations:</b> [1] <i>Mohammad Najibullah was the president of Afghanistan from 1986 to 1992 ...</i> [2] <i>Afghanistan s a multilingual country, where Pashto and Dari (a dialect of Persian) are the official languages with ...</i></p>	



**Example 1**

**Question:** The Maryland Transportation Authority is in charge of what landmarks?

**Answer Statement:** The Maryland Transportation Authority is in charge of providing funding for the construction of revenue-producing transportation facilities for the Maryland Department of Transportation, which have included improvements at the Port of Baltimore and the Baltimore-Washington International Airport

**Citation:** of a commercial enterprise. Its capital projects and operations are funded by tolls, concessions, investment income, and revenue bonds. In addition to its own toll facilities, the Authority finances construction of other revenue-producing transportation facilities for the Maryland Department of Transportation (MDOT). These projects have included improvements at the Port of Baltimore and the Baltimore-Washington International Airport. To provide construction funding, the Authority issues revenue bonds, which will be paid off over a period of years by tolls and other user fees generated by the facilities. The MDTA can issue either taxable or exempt bonds to finance large scale projects.

**AutoIS:** Supportive ✗

**AttrScore:** Irrelevant ✗

**Vicuna<sup>†</sup>:** Partially Supportive ✓

---

**Example 2**

**Question:** When did the last season of jersey shore air?

**Answer Statement:** The TV show Jersey Shore aired its final episode on December 20, 2012.

**Citation:** 8.56 million viewers, only to set another record with the airing of the fourth episode, which garnered 8.87 million viewers. On January 25, 2011, it was confirmed that the show had been renewed for a fourth season, to be filmed in Italy during the first half of 2011. The fourth season premiered August 4, 2011. MTV confirmed in June 2011 that the fifth season would return to Seaside Heights. Believed complications caused by Nicole Polizzi's pregnancy, and several cast members (including Polizzi, DeVecchio, and Farley) receiving spin-offs sparked talk about the future of the series past the fifth season, however

**AutoIS:** Supportive ✗

**AttrScore:** Contradictory ✗

**Vicuna<sup>†</sup>:** Irrelevant ✓

Table 18: Two examples of the results of the three attribution evaluators on ALCE-FineGrained. Content in yellow highlights portions of the answer statement not found in the citation, whereas green indicates content present in the citation.

You will see a question, the corresponding answer, and the cited reference. What you need to do is:

1. Read the question, the answer and the cited reference carefully.
2. You should judge whether the cited reference is *supportive*, *partially supportive*, *contradictory*, or *irrelevant* to answer of the question.

- **Supportive:** The cited reference includes facts that can fully support the answer.
- **Partially Supportive:** The cited reference lacks a part of the facts that are required to infer the answer.
- **Contradictory:** The cited reference includes facts that can infer a different answer.
- **Irrelevant:** The cited reference has no facts that can be used to infer the answer.

Here are some examples of the four categories:

### 1. Supportive

<p><b>Question:</b> Who is hosting the next world cup 2022?</p> <p><b>Answer:</b> The 2022 FIFA World Cup will be hosted by Qatar</p> <p><b>Reference:</b> Title: 2018 and 2022 FIFA World Cup bids. Content: FIFA's headquarters in Zurich. Russia was chosen to host the 2018 World Cup, and Qatar was chosen to host the 2022 World Cup. This made Russia the first Eastern European country to host the World Cup, while Qatar would be the first Middle Eastern country to host the World Cup. Blatter noted that the committee had decided to "go to new lands" and reflected a desire to "develop football" by bringing it to more countries. In each round a majority of twelve votes was needed. If no bid received 12 votes in a round, the bid with the fewest votes</p>
<p><b>Question:</b> Who lived to be the oldest person in the world?</p> <p><b>Answer:</b> The longest-lived human on record was Jeanne Calment, who lived to be 122 years and 164 days old</p> <p><b>Reference:</b> Title: Oldest people. Content: Oldest people This is a list of tables of the oldest people in the world in ordinal ranks. To avoid including false or unconfirmed claims of extreme old age, names here are restricted to those people whose ages have been validated by an international body that specifically deals in longevity research, such as the Gerontology Research Group (GRG) or "Guinness World Records" (GWR), and others who have otherwise been . According to this criterion, the longest human lifespan is that of Jeanne Calment of France (1875–1997), who lived to the age of 122 years, 164 days. She met Vincent van</p>

### 2. Partially Supportive

<p><b>Question:</b> What do you use to test for lipids?</p> <p><b>Answer:</b> To test for lipids, a blood sample is taken after a 12-hour fast, which is then used to measure a lipid profile through mass spectrometry, chromatography, or nuclear magnetic resonance</p> <p><b>Reference:</b> Title: Cholesterol. Content: and then every 3–12 months thereafter. A blood sample after 12-hour fasting is taken by a doctor, or a home cholesterol-monitoring device is used to measure a lipid profile, an approach used to estimate a person's lipoproteins, the vastly more important issue because lipoproteins have always been concordant with outcomes though the lipid profile is commonly discordant LDL Particle Number and Risk of Future Cardiovascular Disease in the Framingham Offspring Study. The lipid profile measures: (a) total cholesterol, (b) cholesterol associated with HDL (i.e. Higher Density {than water} Lipids-transported-within-proteins) particles ("which can regress arterial disease"), (c) triglycerides and (d) (by</p>
---

Figure 6: First page of the annotation guidelines.

**Question:** Where was in the dark tv series filmed?  
**Answer:** In the Dark, a British crime drama series, was filmed in Manchester and Marsden  
**Reference:** Title: In the Dark (UK TV series). Content: of kidnapping two young girls. In the second two-parter, a heavily pregnant Helen is pulled into the dark side of urban Manchester as she deals with an unexpected tragedy. Filming for the series began in April 2017 in Manchester and Marsden. "The Daily Telegraph's Michael Hogan gave the first episode three stars out of five, noting that: ""In the Dark" did show promise and could yet come good. It was taut and tensely atmospheric with an intriguing premise which found its heroine caught in the middle between police and prime suspect." Reviewing the first episode, "The Guardian"s Sam Wollaston concluded

### 3. Contradictory

**Question:** When did Spain win their first world cup?  
**Answer:** Spain won their first FIFA World Cup in 1964, hosted in their home country  
**Reference:** Title: Spain national football team. Content: the silver medal. Spain qualified for their first FIFA World Cup in 1934, defeating Brazil in their first game and losing in a replay to the hosts and eventual champions Italy in the quarter-finals. The Spanish Civil War and World War II prevented Spain from playing any competitive matches between the 1934 World Cup and the 1950 edition's qualifiers. At the 1950 finals in Brazil, they topped their group to progress to the final round, then finished in fourth place. Until 2010, this had been Spain's highest finish in a FIFA World Cup finals, which had given them the name

**Question:** Who was the last person hanged in England?  
**Answer:** Peter Manuel was the last person to be hanged in the UK for killing a police officer  
**Reference:** Title: Harry Allen (executioner). Content: 1957 reduced the number of condemned criminals by 75%, from an average of 15 a year in the early 1950s to about four a year in the late 1950s. As Chief Executioner, on 11 July 1958 Allen hanged American-born Scottish serial killer Peter Manuel at Barlinnie prison, Glasgow. He also hanged Guenther Podola on 5 November 1959, a German-born petty thief, and the last man to be hanged in the UK for killing a police officer. His most controversial case was that of James Hanratty, hanged on 4 April 1962 at Bedford Prison for the "A6 murder" case. Efforts to

### 4. Irrelevant

**Question:** Who plays Patrick in 10 things I hate about you?  
**Answer:** Patrick is played by actor Heath Ledger in the 1999 film 10 Things I Hate About You  
**Reference:** Title: 10 Things I Hate About You. Content: assists by convincing Joey to pay Patrick to take out Kat, under the pretense that this will allow Joey to date Bianca. Patrick agrees to the deal, but Kat rebuffs his first few advances. Michael and Cameron help him by prying Bianca for information on Kat's likes and dislikes. Armed with this knowledge, Patrick begins to win Kat's interest. She goes to a party with him, which enables Bianca to go as well, much to Walter's dismay. At the party, Kat becomes upset when she sees Bianca with Joey,

**Question:** How many medals did Australia win in the 2000 Olympics?  
**Answer:** According to the information provided in the search results, Australia won a total of 58 medals at the 2000 Summer Olympics, with 14 gold, 26 silver, and 28 bronze  
**Reference:** Title: 2000 Summer Paralympics medal table. Content: The location and facilities were shared with the largest event, the 2000 Summer Olympics, which concluded on 1 October. The Games set records for athlete and country participation, tickets sold, hits to the official Games website, and medals on offer. A record of 122 countries (or 123 delegations including independent athletes from Timor-Leste) participated; 68 countries won medals, of which seven won a medal for the first time. A total of 1,657 medals were awarded during the Sydney games: 550 gold, 549 silver, and 558 bronze. Among these performances,

Figure 7: Second page of the annotation guidelines.