



HAL
open science

An Evaluation Framework for Attributed Information Retrieval using Large Language Models

Hanane Djeddal, Pierre Erbacher, Raouf Toukal, Laure Soulier, Karen Pinel-Sauvagnat, Sophia Katrenko, Lynda Tamine

► **To cite this version:**

Hanane Djeddal, Pierre Erbacher, Raouf Toukal, Laure Soulier, Karen Pinel-Sauvagnat, et al.. An Evaluation Framework for Attributed Information Retrieval using Large Language Models. CiKM'24 - 33rd ACM International Conference on Information and Knowledge Management, ACM, Oct 2024, Boise, Idaho, United States. 10.1145/3627673.3679172 . hal-04720446

HAL Id: hal-04720446

<https://hal.science/hal-04720446v1>

Submitted on 3 Oct 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

An Evaluation Framework for Attributed Information Retrieval using Large Language Models

Hanane Djeddal
hanane.djeddal@irit.fr
Université Paul Sabatier, IRIT
Toulouse, France
Ecovadis
Paris, France

Pierre Erbacher
Raouf Toukal
Laure Soulier
pierre.erbacher@isir.upmc.fr
raouf.toukal@isir.upmc.fr
laure.soulier@isir.upmc.fr
Sorbonne Université, CNRS, ISIR
Paris, France

Karen Pinel-Sauvagnat
Karen.Sauvagnat@irit.fr
Université Paul Sabatier, IRIT
Toulouse, France

Sophia Katrenko
skatrenko@ecovadis.com
Ecovadis
Paris, France

Lynda Tamine
lechani@irit.fr
Université Paul Sabatier, IRIT
Toulouse, France

Abstract

With the growing success of Large Language models (LLMs) in information-seeking scenarios, search engines are now adopting generative approaches to provide answers along with in-line citations as attribution. While existing work focuses mainly on attributed question answering, in this paper, we target information-seeking scenarios which are often more challenging due to the open-ended nature of the queries and the size of the label space in terms of the diversity of candidate-attributed answers per query. We propose a reproducible framework to evaluate and benchmark attributed information seeking, using any backbone LLM, and different architectural designs: (1) Generate (2) Retrieve then Generate, and (3) Generate then Retrieve. Experiments using HAGRID, an attributed information-seeking dataset, show the impact of different scenarios on both the correctness and attributability of answers.

CCS Concepts

• Information systems → Information retrieval.

Keywords

Large Language Models, Answer generation, Attributed information seeking

ACM Reference Format:

Hanane Djeddal, Pierre Erbacher, Raouf Toukal, Laure Soulier, Karen Pinel-Sauvagnat, Sophia Katrenko, and Lynda Tamine. 2024. An Evaluation Framework for Attributed Information Retrieval using Large Language Models. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management (CIKM '24)*, October 21–25, 2024, Boise, ID, USA. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3627673.3679172>

1 Introduction

Large Language Models (LLMs) [1, 5] have shown great potential in information-seeking scenarios. This has led to a paradigm shift from document ranking to natural language answer generation to address the user’s query. While this new approach helps overcome several challenges previously faced by traditional search solutions, it also brings new ones to light. One such challenge stems from LLMs’ tendency to generate unfaithful text often referred to as “hallucination” [21]. This problem becomes critical in information-seeking and knowledge-intensive scenarios where users expect contextually rich, up-to-date, and reliable responses to their queries [16]. Attribution has emerged as a new solution where texts are generated by LLMs with supporting sources called *attributions*, allowing users to fact-check the model’s claims [3, 9]. Current works on attributed search often leverage existing Question Answering (QA) datasets to benchmark and evaluate the performance of different attribution methods [3, 8, 9]. However, attributed information seeking is more challenging due to the open-ended nature of the query and the scale of the label space of the query and its answer, as well as the answer and its attribution [3]. A query can have multiple possible answers, supported by multiple possible sources. These aspects are important to consider for the evaluation of attributable information-seeking approaches, but QA frameworks do not inherently account for them [3, 9]. While there has been a continuous rise in the number of contributions in this field [2, 12, 16], there is still a critical lack of open-source frameworks useful for designing, developing, and evaluating attributable information-seeking systems. This paper tackles this gap with the hope of fostering research in this area. While our proposed framework can be used for any dataset targeting an information-seeking task, we perform experiments and report results obtained using the recently designed HAGRID dataset [15], establishing its baselines following three architectures commonly used for attribution: 1) Generate (G) [5, 23, 24] where the model is prompted to generate an answer with or without its attribution; 2) Retrieve then Generate (RTG) [2, 4, 10] where the system first retrieves a list of documents relevant to the user query and then generates an answer based on the query and the relevant



This work is licensed under a Creative Commons Attribution International 4.0 License.




Figure 1: LLM-based scenarios for IR with attribution.

Table 1: Statistics on HAGRID dataset. #Q is the number of queries. #A the number of answers. #Inform is the number of informative answers. #Attrib is the number of Attributable answers. avg # P and avg # Cit are resp. the average number of relevant passages per query and citations by answer.

Split	#Q	#A	#Inform	#Attrib	avg #P	avg #Cit
Train	1922	3214	3214	754	2.73	2
Dev	716	1318	1157	826	2.91	2.5

documents, and 3) Generate then retrieve (**GTR**) [8, 13] where attribution is added in a post-generation stage. While our framework is also agnostic to backbone LLMs used, we report experiment results using Zephyr_{3b}¹, a 3 billion parameter finetuned version of mistral [14]. Regarding evaluation, we measure both the correctness of the answer and the quality of the attribution using a set of different metrics reported in the literature [3, 8, 9], relying on either human-built golds or automatically-built references. In the latter case, we use evaluation metrics based on Natural Language Inference (NLI) models [26] which has been found to correlate well with human judgments [3] making attribution evaluation and benchmarking possible without the need for gold citation annotation.

The contributions of this paper are as follows: 1) we provide an open-source framework that implements key architectures using any backbone LLM, and provides several automatic evaluation metrics for attributed information seeking using different references either automatically built-up or human-assessed; 2) we perform a systematic analysis of the different proposed architectures and key parameters and components, including the impact of retrieval on generation; 3) as of the time of this writing and to our knowledge, this is the first benchmarking of the dataset HAGRID, that could be easily used with new datasets adequate for attributable generative information seeking scenarios. All implementations of scenarios and metrics as well as the generated outputs are available online².

2 Experimental design

2.1 Task formulation

The task of attributed information seeking is formalized as follows: given a user query q_0 and a corpus of document passages D , the goal is to generate an attributed answer A^{att} in response to the query q_0 . A^{att} is a set of n distinct statements (i.e. sentences) a_i , where each statement a_i has in line citations referring to a set $C_i = \{c_{i1}, c_{i2}, \dots\}$ of documents that support the statement a_i , with $c_{ij} \in D$. Following these notations, an attributed answer is represented as $A^{att} = \{(a_1, C_1), (a_2, C_2), \dots, (a_n, C_n)\}$, while an answer without attribution is represented as $A = \{a_1, a_2, \dots, a_n\}$.

2.2 Dataset

We use the HAGRID dataset [15], built on top of the MIRACL English subset [29] where each query and its human-assessed relevant documents are fed to an LLM (GPT-3.5 turbo) to generate answers. Human annotators evaluate the LLM-generated answers according to two criteria: 1) *informativeness* refers to the correctness of the answer statements regarding the query; 2) *attributability* refers to the correspondence between the answer statements and the gold citations. Statistics on the dataset are shown in Table 1. Unlike other datasets used for attributed generative question-answering [3, 9], the HAGRID dataset has three main peculiarities that motivate us to use it for designing our benchmark: 1) its core dataset MIRACL targets an ad-hoc information retrieval task; 2) it provides a gold set of answers with attributions and human annotations w.r.t to both informativeness of answer statements a_i and their attribution quality; and 3) despite its great usefulness for the IR and NLP communities to evaluate attributable generative information-seeking models, it is not so far used for a baseline evaluation with LLM-based approaches.

2.3 Scenarios

Here, we describe the three main architectures used for answer generation with attribution in literature [3, 9, 18] depicted in Figure 1. We report our experiment results using Zephyr_{3b} but results with other LLMs are provided with the code demo³.

(1) **Generate (G)**: this scenario directly generates an answer A using an LLM in response to an input query q_0 in a closed-book fashion. Since the model solely depends on its pre-training data, we stick to the simple case of generating an answer A without citations. (2) **Retrieve Then Generate (RTG)**: this scenario first retrieves a list $P = \{d_1, \dots, d_k\}$ of k candidate supporting documents to the query q_0 from the corpus D , which are then fed to the LLM with q_0 to generate an attributed answer A^{att} . We investigate two scenarios:

- *Vanilla retrieval (vanilla)*: we use a two-stage ranker based on BM25 [22] and MonoT5 [19] to retrieve the list P of relevant documents to query q_0 used to generate the attributed answer A^{att} .

- *Retrieval with query generation (query-gen)*: we propose a new variant of the RTG architecture where we leverage the LLM’s parametric memory to guide the retrieval towards documents that align with the model’s internal knowledge using query reformulations. We hope to bridge the gap between the model’s stored knowledge and externally retrieved content [16, 25, 27]. To do so, given the user query q_0 , the LLM first generates a list of subqueries $Q = \{q_1, \dots, q_m\}$ based on q_0 . For each generated query $q_t \in Q$ as well as the user query q_0 , we retrieve k candidate relevant documents $P_t = \{d_{t1}, d_{t2}, \dots, d_{tk}\}$ to the query q_t , for $t = 0, \dots, m$, using the retrieval model used in the *vanilla* scenario. Then, we aggregate

¹<https://huggingface.co/stabilityai/stablelm-zephyr-3b>

²<https://github.com/hanane-djeddal/Attributed-IR>

³<https://github.com/hanane-djeddal/Attributed-IR>

using a data fusion technique (Cf. Section 2.4) the different lists of documents P_t with $t = 0, \dots, m$ to produce one final list of candidate documents $P = \{d_1, \dots, d_k\}$ which is then fed to the LLM to generate a grounded attributed answer A^{att} .

(3) **Generate Then Retrieve (GTR)**: This scenario first generates an answer A without citations, which is then used to identify relevant documents that support it. More specifically, to generate inline citations along answer A , we split it into statements a_i for which the retriever returns a ranked list P_i of k relevant documents $P_i = \{d_{i1}, d_{i2}, \dots, d_{ik}\}$ that we consider as the associated citations to build the attributed answer A^{att} . Similar to *vanilla* and *query-gen* scenarios, we use BM25+MonoT5 for the retrieval.

2.4 Variants and baselines

To better understand the impact of retrieval on performances in the RTG scenario, we establish a baseline **RTG-gold** where the LLM generates the answer using the gold documents. To evaluate the quality and the impact of our new scenario of generating m queries (**RTG-query-gen**), we test different ranking aggregation methods from the literature: 1) **CombSum** [7], 2) **CombSumMNZ** [7] and 3) **PM2** [6]. We also define additional baseline methods: 4) **Sort**: sorting the documents from the various lists by their normalized MonoT5 score, and 5) **Rerank**: reranking the documents from the different lists using MonoT5. We also report the best possible performance for query-gen (**MAX**) by evaluating the retrieval results of each generated query separately and choosing the best score for each metric. For the answer generation, we use the documents issued from **Rerank**.

2.5 Evaluation metrics

We consider in our framework two main categories w.r.t the output structure including both answer statements and citations.

Answer correctness. Contrary to QA settings where answers can be short and exact, in information-seeking scenarios, answers tend to be long and can vary in form. Therefore we use different n-gram-based and semantic-based metrics: ROUGE [17], BLEU [20], and BertScore [28] to evaluate the accuracy of the answer based on the gold answer. Since a query can have multiple possible answers in the ground truth, we report for each metric, the best score calculated by comparing the generated answer with each candidate answers.

Citation quality. We apply different metrics to measure either the answer statements' attribution w.r.t. the associated citations in the answer or the gold citations.

• **AutoAIS** [8]: This metric was originally defined for the task of attributed QA where the citations are not in line with the answer but reported in a separate list associated with the answer as a whole. In our case, each statement a_i has its own set of citations C_i in line with the answer A^{att} . Thus, we adapt this metric:

$$AutoAIS_{cit}(A^{att}) = avg_{(a_i, C_i) \in A^{att}} [argmax_{c_{ij} \in C_i} (NLI(c_{ij}, a_i))] \quad (1)$$

where $NLI(premise, hypothesis)$ is the output of an NLI classifier (1 if the premise entails the hypothesis, and 0 otherwise). This first application of AutoAIS determines whether a statement is attributed to at least one of its associated citations.

Moreover, to investigate to which extent an answer A^{att} is attributable to the retrieved documents $P = \{d_1, \dots, d_k\}$ fed to the LLM in

the **RTG** scenario, we introduce a variant of the AutoAIS metric, $AutoAIS_{pssg}(A^{att})$, as follows:

$$AutoAIS_{pssg}(A^{att}) = avg_{a_i \in A} [argmax_{d_i \in P} (NLI(d_i, a_i))] \quad (2)$$

• **Citation Recall and Precision**: Facing the lack of gold citations, the framework ALCE [9] introduced a citation recall and citation precision using the automatic binary score (0,1) of an NLI classifier between the statements and the citations. Citation recall determines if a statement a_i is entirely supported by its cited documents by examining the NLI score between the statement and the concatenation of its cited documents denoted $concat(C_i)$ i.e. by computing $NLI(concat(C_i), a_i)$. Citation precision measures to which extent a citation c_{ij} is irrelevant by checking whether $NLI(c_{ij}, a_i) = 0$ and if it can be removed without affecting the attribution score of its statement a_i such that $NLI(concat(C_i \setminus c_{ij}), a_i) = 1$. We use the implementation provided in the ALCE framework⁴ and refer to them using NLI_prec and NLI_rec in the rest of the paper.

• **Citation Overlap**: Since HAGRID provides gold answers A^{att} with gold citations C_{gold} , we perform a more direct evaluation of citation attributability by computing for each query the overlap between the gold citations C_{gold} and the whole citations in the generated answer C_{gen} . We define citation overlap precision and recall for answer A^{att} to query q as:

$$Overlap(A^{att}) = \frac{|C_{gold} \cap C_{gen}|}{|X|} \quad (3)$$

Where $C_{gen} = \bigcup_{(a_i, c_i) \in A^{att}} C_i$, $|\cdot|$ is the cardinality of the set. X equals to C_{gen} or C_{gold} for computing the overlap precision or overlap recall, resp..

3 Experiments

3.1 Effectiveness results

We report the main results of the different scenarios (cf. Section 2.3) in Table 2 using the evaluation metrics presented in Section 2.5. For **RTG-vanilla** and **RTG-query-gen**, we fix the number of supporting documents to 2 for a fair comparison with the **RTG-gold** which relies on 2.7 gold documents on average (cf. Table 1). For the **GTR** scenario, we retrieve 1 document per statement to be close to the attributed answers in the ground truth including 2.5 citations per answer on average. From a general point of view, we can see that **RTG** scenarios have the best overall performances which confirms the benefit of retrieval-augmented answer generation [11]. This can also be seen through the lower results obtained by the **GTR** scenario suggesting that using attribution at a post-generation stage can propagate the hallucinations in the generated answer making it hard to identify supporting documents. Moreover, **RTG-gold** sets a strong upper bound, showcasing the importance of the quality of the retrieved documents. Among the evaluated **RTG** variants, **RTG-query-gen** has the best performance in citation quality regarding AutoAIS and NLI-based metrics. It also obtains higher results than the **RTG-gold** scenario regarding NLI-based metrics. This confirms our intuition that query reformulation helps in exhibiting parametric knowledge in LLMs towards information-seeking needs.

⁴<https://github.com/princeton-nlp/ALCE/tree/main>

Table 2: Main results for G, RTG, and GTR scenarios using the HAGRID dataset. *Gold answer* reports the citation quality in the gold answer using automatic metrics. Since the scenario G does not include citations, only answer correctness is evaluated. The generated simple answer in scenario GTR and G are the same, thus the answer correctness measures are equal. The best performance measures are in bold, and the second best is underlined.

	Correctness						Citations						
	BLEU	ROUGE-L			BertScore			Overlap		AutoAIS		NLI prec/rec	
		Prec.	Rec.	F-score	Prec.	Rec.	F-score	Prec.	Rec.	Cit.	Pssg.	Prec.	Rec.
Gold answer	-	-	-	-	-	-	-	-	87.97	89.21	83.65	79.80	
G	11.06	30.41	46.08	31.58	87.88	90.02	88.87	-	-	-	-	-	
RTG - gold	28.22	44.00	63.81	46.72	90.02	93.36	91.69	75.29	68.89	42.81	80.67	56.55	<u>42.31</u>
RTG - vanilla	<u>18.44</u>	<u>33.83</u>	<u>56.40</u>	<u>36.65</u>	<u>87.94</u>	<u>91.52</u>	<u>89.63</u>	<u>36.17</u>	<u>32.69</u>	41.86	78.95	<u>57.90</u>	41.63
RTG - query-gen	18.33	33.58	56.13	36.43	87.89	<u>91.55</u>	89.62	35.89	32.46	<u>42.68</u>	<u>80.10</u>	59.59	42.48
GTR	11.06	30.41	46.08	31.58	87.88	90.02	88.87	45.53	30.53	26.69	<u>26.69</u>	26.65	26.66

Figure 2: Citations and Correctness metrics for varying values of the number of supporting documents in RTG-user-query scenario




Table 3: Retrieval results for the RTG scenario. Best performances are in bold and the second best is underlined.

Setting	Agg	@1		@10		nDCG
		Precision	Recall	Precision	Recall	
user query	-	<u>48.04</u>	23.29	16.45	63.31	53.21
query-gen	sort	36.17	17.87	13.84	57.03	43.88
	combSum	47.07	22.37	16.17	62.45	51.39
	combSum-MNZ	45.39	21.22	16.28	61.98	50.49
	Rerank	46.79	22.83	<u>16.33</u>	63.99	53.21
	PM2	48.60	23.16	<u>16.33</u>	63.43	52.83
	MAX	66.34	32.42	18.66	<u>72.69</u>	63.58

We also note the difference between citation quality and faithfulness to input documents. $AutoAIS_{cit}$ scores can be lower than half the scores of $AutoAIS_{Pssg}$ across different scenarios. This means that while input documents can correctly support the sentences in the answer, they are not always correctly cited.

3.2 Complementary analysis

3.2.1 Impact of the number of supporting documents. Figure 2 shows the performances of the RTG-vanilla scenario when varying the number of supporting documents. Please note that trends are similar for the RTG-query-gen scenario. Regarding the correctness metrics (Figure 2(a)), the number of documents has a slight impact, with a small tendency towards the decrease of ROUGE metrics. Focusing on the citation metrics (Figure 2(b)), the number of documents seems to have little impact on $AutoAIS_{Pssg}$, contrary to the other metrics which overall decline when increasing the number of documents. This suggests that the model may struggle to cite

correctly when increasing the number of input documents, which may stem from a tendency to over-cite (increase in NLI precision without necessarily improving recall). We also note that the recall of citations-overlap peaks at around 2 documents which corresponds to the average number of gold documents in the dataset.

3.2.2 Impact of Retrieval. Using the generated queries in the RTG-query-gen scenario can improve the retrieval scenario, especially in citation quality. We, therefore, analyze in Table 3 the impact of the different aggregation approaches over the generated queries on the retrieval effectiveness. We can see that the re-ranking aggregation method (**Rerank**) seems to give the best retrieval score @10 though not improving much from the vanilla scenario. The best obtainable results using **query-gen** scenario can go up to 72% recall @10, leaving room to investigate other aggregation methods.

3.2.3 Correlation between NLI and Annotations. We evaluate the gold answers provided in HAGRID using the automatic citation metrics and examine how they correlate with the provided human annotation of attributability. We find that the Pearson correlation coefficient between $AutoAIS_{cit}$ and the attribution annotation is pretty modest at 0.36, similarly with citation precision and recall at 0.35 and 0.36 respectively. This is an interesting finding as previous work [3] found a strong correlation between human annotations and AutoAIS when using QA dataset for attribution, which testify to the complexity of attributed information-seeking and its distinction from attributed QA.

4 Conclusion

In this work, we propose an extensible open-source framework to benchmark attributable information-seeking using representative LLM-based approaches. We consider a set of up-to-date SOTA evaluation metrics to measure both the correctness of the generated answer and the quality of the in-line citations. Our work provides valuable baseline models and insights from empirical results and trends allowing to strengthen the maturity of research that targets faithful generative information-seeking.

5 Acknowledgments

This work partially benefited from support from the French National Research Agency (Project GUIDANCE, ANR-23-IAS1-0003)

References

- [1] OpenAI Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, S. Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mo Bavarian, Jeff Belgum, Irwan Bello, Jake Berdine, Gabriel Bernadett-Shapiro, Christopher Berner, Lenny Bogdonoff, Oleg Boiko, Madelaine Boyd, Anna-Luisa Brakman, Greg Brockman, Tim Brooks, Miles Brundage, Kevin Button, Trevor Cai, Rosie Campbell, Andrew Cann, Brittany Carey, Chelsea Carlson, Rory Carmichael, Brooke Chan, Che Chang, Fotis Chantzis, Derek Chen, Sully Chen, Ruby Chen, Jason Chen, Mark Chen, B. Chess, Chester Cho, Casey Chu, Hyung Won Chung, Dave Cummings, Jeremiah Currier, Yunxing Dai, Cory Decareaux, Thomas Degry, Noah Deutsch, Damien Deville, Arka Dhar, David Dohan, Steve Dowling, Sheila Dunning, Adrien Ecoffet, Atty Eleti, Tyna Eloundou, David Farhi, L. Fedus, Niko Felix, Sim'on Posada Fishman, Juston Forte, Isabella Fulford, Leo Gao, Elie Georges, C. Gibson, Vik Goel, Tarun Gogineni, Gabriel Goh, Raphael Gontijo-Lopes, Jonathan Gordon, Morgan Grafstein, S. Gray, Ryan Greene, Joshua Gross, S. Gu, Yufei Guo, Chris Hallacy, Jesse Han, Jeff Harris, Yuchen He, Mike Heaton, Johannes Heidecke, Chris Hesse, Alan Hickey, Wade Hickey, Peter Hoeschele, Brandon Houghton, Kenny Hsu, Shengli Hu, Xin Hu, Joost Huizinga, Shantanu Jain, Shawn Jain, Joanne Jang, Angela Jiang, Roger Jiang, Haozhun Jin, Denny Jin, Shino Jomoto, Billie Jonn, Heewoo Jun, Tomer Kaftan, Lukasz Kaiser, Ali Kamali, I. Kanitscheider, N. Keskar, Tabarak Khan, Logan Kilpatrick, Jong Wook Kim, Christina Kim, Yongjik Kim, Hendrik Kirchner, J. Kiros, Matthew Knight, Daniel Kokotajlo, Lukasz Kondraciuk, A. Kondrich, Aris Konstantinidis, Kyle Kosic, Gretchen Krueger, Vishal Kuo, Michael Lampe, Ikai Lan, Teddy Lee, J. Leike, Jade Leung, Daniel Levy, Chak Ming Li, Rachel Lim, Molly Lin, Stephanie Lin, Mateusz Litwin, Theresa Lopez, Ryan Lowe, Patricia Lue, A. Makanju, Kim Malfacini, Sam Manning, Todor Markov, Yaniv Markovski, Bianca Martin, Katie Mayer, Andrew Mayne, Bob McGrew, S. McKinney, C. McLeavey, Paul McMillan, Jake McNeil, David Medina, Aalok Mehta, Jacob Menick, Luke Metz, Andrey Mishchenko, Pamela Mishkin, Vinnie Monaco, Evan Morikawa, Daniel P. Mossing, Tong Mu, Mira Murati, O. Murk, David M'ely, Ashvin Nair, Reiichiro Nakano, Rajeev Nayak, Arvind Neelakantan, Richard Ngo, Hyeonwoo Noh, Ouyang Long, Cullen O'Keefe, J. Pachocki, Alex Paino, Joe Palermo, Ashley Pantuliano, Giambattista Parascandolo, Joel Parish, Emy Parparita, Alexandre Passos, Mikhail Pavlov, Andrew Peng, Adam Perelman, Filipe de Avila Belbute Peres, Michael Petrov, Henrique Pondé de Oliveira Pinto, Michael Pokorny, Michelle Pokrass, Vitchyr H. Pong, Tolly Powell, Alethea Power, Boris Power, Elizabeth Proehl, Raul Puri, Alec Radford, Jack Rae, Aditya Ramesh, Cameron Raymond, Francis Real, Andrea Rimbach, Carl Ross, Bob Rotsted, Henri Roussez, Nick Ryder, M. Saltarelli, Ted Sanders, Shibani Santurkar, Girish Sastry, Heather Schmidt, David Schnurr, John Schulman, Daniel Selsam, Kyla Sheppard, T. Sherbakov, Jessica Shieh, S. Shoker, Pranav Shyam, Szymon Sidor, Eric Sigler, Maddie Simens, Jordan Sitkin, Katarina Slama, Ian Sohl, Benjamin D. Sokolowsky, Yang Song, Natalie Staudacher, F. Such, Natalie Summers, I. Sutskever, Jie Tang, N. Tezak, Madeleine Thompson, Phil Tillet, Amin Tootoonchian, Elizabeth Tseng, Preston Tuggle, Nick Turley, Jerry Tworek, Juan Felipe Cer'on Uribe, Andrea Vallone, Arun Vijayvergiya, Chelsea Voss, Carroll L. Wainwright, Justin Jay Wang, Alvin Wang, Ben Wang, Jonathan Ward, Jason Wei, C. J. Weinmann, Akila Welihinda, P. Welinder, Jiayi Weng, Lillian Weng, Matt Wiethoff, Dave Willner, Clemens Winter, Samuel Wolrich, Hannah Wong, Lauren Workman, Sherwin Wu, Jeff Wu, Michael Wu, Kai Xiao, Tao Xu, Sarah Yoo, Kevin Yu, Qiming Yuan, Wojciech Zaremba, Rowan Zellers, Chong Zhang, Marvin Zhang, Shengjia Zhao, Tianhao Zheng, Juntang Zhuang, William Zhuk, and Barret Zoph. [n. d.]. GPT-4 Technical Report. <https://www.semanticscholar.org/paper/GPT-4-Technical-Report-Achiam-Adler/163b4d6a79a5b19af88b8585456363340d9efd04>
- [2] Akari Asai, Zeqiu Wu, Yizhong Wang, Avirup Sil, and Hannaneh Hajishirzi. 2023. Self-RAG: Learning to Retrieve, Generate, and Critique through Self-Reflection. arXiv:2310.11511 [cs.CL]
- [3] Bernd Bohnet, Vinh Tran, Pat Verga, Roei Aharoni, Daniel Andor, Livio Baldini Soares, Massimiliano Ciaramita, Jacob Eisenstein, Kuzman Ganchev, Jonathan Herzig, Kai Hui, Tom Kwiatkowski, Ji Ma, Jianmo Ni, Tal Schuster, Lierni Sestorain Saralegui, William Weston Cohen, Michael Collins, Dipanjan Das, Don Metzler, Slav Petrov, and Kellie Webster. 2022. Attributed Question Answering: Evaluation and Modeling for Attributed Large Language Models. <https://arxiv.org/abs/2212.08037>
- [4] Sebastian Borgeaud, Arthur Mensch, Jordan Hoffmann, Trevor Cai, Eliza Rutherford, Katie Millican, George Bm Van Den Driessche, Jean-Baptiste Lespiau, Bogdan Damoc, Aidan Clark, Diego De Las Casas, Aurelia Guy, Jacob Menick, Roman Ring, Tom Hennigan, Saffron Huang, Loren Maggiore, Chris Jones, Albin Cassirer, Andy Brock, Michela Paganini, Geoffrey Irving, Oriol Vinyals, Simon Osindero, Karen Simonyan, Jack Rae, Erich Elsen, and Laurent Sifre. 2022. Improving Language Models by Retrieval from Trillions of Tokens. In *Proceedings of the 39th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 162)*, Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato (Eds.). PMLR, 2206–2240. <https://proceedings.mlr.press/v162/borgeaud22a.html>
- [5] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language Models are Few-Shot Learners. In *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (Eds.), Vol. 33. Curran Associates, Inc., 1877–1901. https://proceedings.neurips.cc/paper_files/paper/2020/file/1457c0d6bfc4967418bfb8ac142f64a-Paper.pdf
- [6] Van Dang and W. Bruce Croft. 2012. Diversity by proportionality: an election-based approach to search result diversification. In *Proceedings of the 35th International ACM SIGIR Conference on Research and Development in Information Retrieval* (Portland, Oregon, USA) (SIGIR '12). Association for Computing Machinery, New York, NY, USA, 65–74. <https://doi.org/10.1145/2348283.2348296>
- [7] Edward A. Fox and Joseph A. Shaw. 1993. Combination of Multiple Searches. In *Text Retrieval Conference*. <https://api.semanticscholar.org/CorpusID:1309301>
- [8] Luyu Gao, Zhuyun Dai, Panupong Pasupat, Anthony Chen, Arun Tejasvi Chaganty, Yicheng Fan, Vincent Zhao, Ni Lao, Hongrae Lee, Da-Cheng Juan, and Kelvin Guu. 2023. 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)*, Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (Eds.). Association for Computational Linguistics, Toronto, Canada, 16477–16508. <https://doi.org/10.18653/v1/2023.acl-long.910>
- [9] Tianyu Gao, Howard Yen, Jiatong Yu, and Danqi Chen. 2023. Enabling Large Language Models to Generate Text with Citations. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, Houda Bouamor, Juan Pino, and Kalika Bali (Eds.). Association for Computational Linguistics, Singapore, 6465–6488. <https://doi.org/10.18653/v1/2023.emnlp-main.398>
- [10] Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Mingwei Chang. 2020. Retrieval Augmented Language Model Pre-Training. In *Proceedings of the 37th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 119)*, Hal Daumé III and Aarti Singh (Eds.). PMLR, 3929–3938. <https://proceedings.mlr.press/v119/guu20a.html>
- [11] Sebastian Hofstätter, Jiecao Chen, Karthik Raman, and Hamed Zamani. 2023. FiD-Light: Efficient and Effective Retrieval-Augmented Text Generation. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval* (<conf-loc>, <city>Taipei</city>, <country>Taiwan</country>, </conf-loc>) (SIGIR '23). Association for Computing Machinery, New York, NY, USA, 1437–1447. <https://doi.org/10.1145/3539618.3591687>
- [12] Nan Hu, Jiaoyan Chen, Yike Wu, Guilin Qi, Sheng Bi, Tongtong Wu, and Jeff Z. Pan. 2024. Benchmarking Large Language Models in Complex Question Answering Attribution using Knowledge Graphs. arXiv:2401.14640 [cs.CL]
- [13] Siqing Huo, Negar Arabzadeh, and Charles Clarke. 2023. Retrieving Supporting Evidence for Generative Question Answering. In *Proceedings of the Annual International ACM SIGIR Conference on Research and Development in Information Retrieval in the Asia Pacific Region* (<conf-loc>, <city>Beijing</city>, <country>China</country>, </conf-loc>) (SIGIR-AP '23). Association for Computing Machinery, New York, NY, USA, 11–20. <https://doi.org/10.1145/3624918.3625336>
- [14] Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, Léo Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. 2023. Mistral 7B. arXiv:2310.06825 [cs.CL]
- [15] Ehsan Kamaloo, Aref Jafari, Xinyu Zhang, Nandan Thakur, and Jimmy Lin. [n. d.]. HAGRID: A Human-LLM Collaborative Dataset for Generative Information-Seeking with Attribution. <https://doi.org/10.48550/arXiv.2307.16883> [cs]
- [16] Dongfang Li, Zetian Sun, Xinshuo Hu, Zhenyu Liu, Ziyang Chen, Baotian Hu, Aiguo Wu, and Min Zhang. [n. d.]. A Survey of Large Language Models Attribution. arXiv:2311.03731 [cs] <http://arxiv.org/abs/2311.03731>
- [17] Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. Text summarization branches out, 74–81.
- [18] Nelson Liu, Tianyi Zhang, and Percy Liang. 2023. Evaluating Verifiability in Generative Search Engines. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, Houda Bouamor, Juan Pino, and Kalika Bali (Eds.). Association for Computational Linguistics, Singapore, 7001–7025. <https://doi.org/10.18653/v1/2023.findings-emnlp.467>
- [19] Rodrigo Nogueira, Zhiying Jiang, Ronak Pradeep, and Jimmy Lin. 2020. Document Ranking with a Pretrained Sequence-to-Sequence Model. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, Trevor Cohn, Yulan He, and Yang Liu (Eds.). Association for Computational Linguistics, Online, 708–718. <https://doi.org/10.18653/v1/2020.findings-emnlp.63>
- [20] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the*

- 40th annual meeting of the Association for Computational Linguistics. 311–318.
- [21] Denis Peskoff and Brandon Stewart. 2023. Credible without Credit: Domain Experts Assess Generative Language Models. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (Eds.). Association for Computational Linguistics, Toronto, Canada, 427–438. <https://doi.org/10.18653/v1/2023.acl-short.37>
- [22] Stephen Robertson and Hugo Zaragoza. 2009. The Probabilistic Relevance Framework: BM25 and Beyond. *Foundations and Trends in Information Retrieval* 3 (01 2009), 333–389. <https://doi.org/10.1561/15000000019>
- [23] Zhiqing Sun, Xuezhi Wang, Yi Tay, Yiming Yang, and Denny Zhou. 2023. Recitation-Augmented Language Models. In *The Eleventh International Conference on Learning Representations*. <https://openreview.net/forum?id=-cqvvvb-NkI>
- [24] Orion Weller, Marc Marone, Nathaniel Weir, Dawn Lawrie, Daniel Khashabi, and Benjamin Van Durme. 2024. “According to . . .”: Prompting Language Models Improves Quoting from Pre-Training Data. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, Yvette Graham and Matthew Purver (Eds.). Association for Computational Linguistics, St. Julian’s, Malta, 2288–2301. <https://aclanthology.org/2024.eacl-long.140>
- [25] Kevin Wu, Eric Wu, and James Zou. 2024. How faithful are RAG models? Quantifying the tug-of-war between RAG and LLMs’ internal prior. *ArXiv abs/2404.10198* (2024). <https://api.semanticscholar.org/CorpusID:269157310>
- [26] Yuexiang Xie, Fei Sun, Yang Deng, Yaliang Li, and Bolin Ding. 2021. Factual Consistency Evaluation for Text Summarization via Counterfactual Estimation. In *Conference on Empirical Methods in Natural Language Processing*. <https://api.semanticscholar.org/CorpusID:237353254>
- [27] Shuo Zhang, Liangming Pan, Junzhou Zhao, and William Yang Wang. 2023. The Knowledge Alignment Problem: Bridging Human and External Knowledge for Large Language Models. <https://api.semanticscholar.org/CorpusID:265302382>
- [28] Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2019. Bertscore: Evaluating text generation with bert. *arXiv preprint* (2019).
- [29] Xinyu Zhang, Nandan Thakur, Odunayo Ogundepo, Ehsan Kamalloo, David Alfonso-Hermelo, Xiaoguang Li, Qun Liu, Mehdi Rezagholizadeh, and Jimmy Lin. 2022. Making a MIRACL: Multilingual Information Retrieval Across a Continuum of Languages. *arXiv:2210.09984* [cs.IR]