

# CITECHECK: Towards Accurate Citation Faithfulness Detection

Anonymous ACL submission

## Abstract

Citation faithfulness detection is critical for enhancing retrieval-augmented generation (RAG) systems, yet large-scale Chinese datasets for this task are scarce. Existing methods face prohibitive costs due to the need for manually annotated negative samples. To address this, we introduce the first large-scale Chinese dataset CITECHECK for citation faithfulness detection, constructed via a cost-effective approach using two-stage manual annotation. This method balances positive and negative samples while significantly reducing annotation expenses. CITECHECK comprises training and test splits. Experiments demonstrate that: (1) the test samples are highly challenging, with even state-of-the-art LLMs failing to achieve high accuracy; and (2) training data augmented with LLM-generated negative samples enables smaller models to attain strong performance using parameter-efficient fine-tuning. CITECHECK provides a robust foundation for advancing citation faithfulness detection in Chinese RAG systems. The dataset is publicly available to facilitate research.

## 1 Introduction

Large Language Models (LLMs) are prone to generating factual errors through hallucinations when answering real-world questions. Retrieval-augmented generation (RAG) systems (Lewis et al., 2020; Guu et al., 2020; Borgeaud et al., 2022) address this limitation by leveraging external information retrieval to ground LLM responses in verifiable sources. Recent advancements extend RAG systems to generate text with inline citations (Gao et al., 2023b), enabling users to validate the reliability of generated content by cross-referencing cited documents. However, studies reveal a critical weakness in these systems: citation faithfulness. A substantial portion of generated text may lack proper support from the cited references (Liu et al., 2023; Hu et al., 2024b), undermining the

trustworthiness and verification capability of RAG outputs. This challenge necessitates accurate citation faithfulness detection—determining whether cited passages genuinely support their associated claims—as a fundamental requirement for improving RAG reliability.

Developing robust citation faithfulness detection methods requires large-scale, high-quality datasets. While English benchmarks have emerged (Yue et al., 2023), Chinese datasets remain notably absent. Constructing such resources presents unique challenges: realistic negative samples (unsupported citations) must originate from strong RAG systems to ensure practical usage, yet these systems rarely produce such errors. For instance, a RAG system with a 10% error rate would require annotating approximately 70,000 samples to collect 7,000 negative examples—a prohibitively expensive endeavor. This tension between dataset quality and construction cost demands innovative solutions for efficient data curation without compromising sample integrity.

To bridge this gap, we introduce CITECHECK, the first large-scale Chinese dataset for citation faithfulness detection. Our approach combines 11,307 knowledge-intensive questions with a novel two-stage annotation framework that reduces labeling costs while preserving data quality. CITECHECK comprises two distinct components designed to address both detection difficulty and training efficacy.

The development and test sets each contain 500 positive (supported) and 500 negative (unsupported) samples totaling 2,000 unmodified RAG outputs. Experimental analysis demonstrates these original samples pose significant challenges, with state-of-the-art LLMs achieving limited detection accuracy. The training set includes 9,796 samples (4,898 positive/negative pairs) where negative instances are generated through LLM-based document modification rather than relying solely on rare

RAG errors. Despite this augmentation, parameter-efficient fine-tuning on 7B-8B parameter models yields strong detection performance, confirming the preserved quality of modified negative samples.

Our contributions are threefold: CITECHECK establishes the first comprehensive benchmark for Chinese citation faithfulness detection; (2) We propose an efficient data augmentation strategy that reduces annotation costs by 86% compared to conventional approaches; (3) Extensive experiments validate the dataset’s quality and utility, showing that models trained on our augmented data effectively generalize to challenging real-world samples. This work advances reliable RAG development by providing essential resources and methodologies for building verifiable, citation-grounded LLM applications in Chinese.

## 2 Dataset Construction

### 2.1 Question Collection

We collect Chinese questions from the sources: **WebText** (Xu, 2019): A large-scale Chinese community question-answering dataset spanning diverse topics.

**WebCPM** (Qin et al., 2023): A Chinese long-form question-answering dataset focused on interactive web search contexts.

**Zhihu-KOL** (Wang, 2023): A high-quality question-answering dataset derived from Zhihu, a prominent Chinese QA platform.

**RGB** (Chen et al., 2024): A bilingual question-answering dataset based on news reports.

**TrickQA**: Questions with ambiguous, incorrect, or unverifiable premises (see Appendix A for details).

After collecting these questions, we input them into an open-sourced RAG system to simulate real-world question-answering scenarios and analyze how the system processes and responds to these diverse inputs. The RAG system retrieves five external documents and generates responses. Statements in the answers are annotated with citation marks (1–5), indicating alignment with information from the corresponding documents. On average, each statement spans 33.4 tokens, while each document averages 177.3 tokens. An original sample is formed by pairing a labeled statement with its cited documents, represented as a tuple (question, answer, statement, cited documents).

### 2.2 Data Augmentation

The goal of data augmentation is to create negative samples of high quality by making minor modifications to the cited documents in the original samples. Given the use of an industrial RAG system, the number of negative samples in the original samples is estimated to be insufficient. To construct a balanced training set, as well as a label-balanced dev set and test set for evaluation, successfully augmented negative samples can be used. The modified documents should not be inconsistent or incoherent, so as not to provide the trained model with a false basis for judging the negative samples.

We use GPT-4o (OpenAI et al., 2024) for data augmentation. After providing the original sample to the LLM, it is asked to perform the following steps in sequence:

**Segments Identification**: Find all key segments in the cited document that directly support the information in the statement.

**Segments Grouping**: Group the key segments by the information they support, with each group containing key segments that support the same or related information in the statement.

**Segments Modification**: Select a group of key segments and modify them so that they do not support the corresponding information in the statement.

The modification changes only the portion that relates to the supported information in the statement. This maintains logical flow and non-contradictory information within the key segments, and keep the key segments logical in the context of the document and non-contradictory to other information in the document. If there is more than one key segment in a group, the information in all of them should be consistent after the modification.

For each sample, the LLM is asked to try two methods of modification:

**Content Revision**: Alter specific details within a key segment without introducing direct contradictions to the original information.

**Structure Preservation**: Remove information from a key segment while ensuring the overall coherence and integrity of the segment remain intact.

After completing the LLM augmentation, each original sample is accompanied by the LLM-labeled key segment information and corresponds to the two augmented samples generated by the LLM using the two modification methods. The cost is 0.026\$ per sample. See Appendix B for more details of the augmentation.

Question	特斯拉在中国的纯电动汽车销量占比是多少? What is Tesla's share of all-electric car sales in China?	乘坐飞机的时候托运一个行李箱，再带一个20寸的箱子，带上飞机的箱子会被称重吗？ When I check a suitcase on an airplane and bring a 20" case, will the case be weighed when I bring it on the plane?	为什么当你把一幅画切成几块时，每一块都只有一部分，但是当你把一张全息图切成几块时，每一块都有完整的图像？ Why is it that when you cut a painting into pieces, each piece has only a part of it, but when you cut a hologram into pieces, each piece has the full image?	如果闪电劈中混凝土大楼之类的，就算在房间里处于隔缘状态，也会受伤害么？ If lightning strikes a concrete building or something like that, can you get hurt even if you're in a room that's insulated?
Answer (Statement)	特斯拉在2023年上半年纯电动 汽车市场的占有率为21.7%。 Tesla's share of the all-electric car market in the first half of 2023 is 21.7 percent.	是的，乘坐飞机时... 每件托 运行李的重量不能超过 <u>50公斤</u> <u>，而每位旅客随身携带物品的</u> <u>重量以5公斤为限。</u> 此外，每 位旅客...  The weight of each piece of checked baggage cannot exceed 50 kilograms, while the weight of each passenger's carry-on is limited to 5 kilograms.	这是因为... 全息的原理在于 将动量空间通过全息技术记录 在实空间里，因此， <u>如果全息</u> <u>图缺失一部分，那部分实空间</u> <u>的信息就丢失了，像就不完整</u> 了。即使... If a portion of the hologram is missing, that portion of real space is lost and the likeness is incomplete.	闪电击中混凝土大楼... 在澳 大利亚新南威尔士州纽卡斯尔 的一个案例中，闪电击中了一 座大楼，虽然 <u>大楼的结构保持</u> <u>良好</u> ，但是巨大的爆裂声和震 动可能会对内部的人造成伤害 。因此，... The building is structurally sound.
Cited Documents	[1] 【2023上半年】全球纯电动 汽车销量出炉...】...据该报 道，特斯拉在纯电动汽车市场 期间占据21.7%的份额。...  First half of 2023 / Tesla held a 21.7% share of the all-electric car market during the period.	[1] 办理 <u>托运行李</u> 对行李物品 规定如下：... <u>每件行李物品重</u> <u>量不能超过50公斤。</u> ... Check-in baggage / The weight of each baggage item can not ex- ceed 50 kilograms. [2] 随身携带物品的重量， <u>每</u> <u>位旅客以5公斤为限。</u> ... The weight of carry-on items is limited to 5 kg per passenger.	[1] ...如果普通照片 <u>缺失一部</u> <u>分</u> ， <u>那部分实空间的信息就丢</u> <u>失了，像就不完整了。</u> 全息照 片如果 <u>缺失一部分</u> ，同样会造 成信息的 <u>缺失</u> ，但是... If a portion of an ordinary photo- graph is missing, that portion of real space is lost and the likeness is incomplete.	[1] ...澳大利亚新南威尔士州 纽卡斯尔...可清楚看到闪电击 中大楼的场面，同时可听到巨 大的爆裂声。据悉，闪电所击 中的大楼为一处健身房。...
Label	positive	positive	negative	negative
Note	supported by a single document	supported by multiple documents	contradictory information	unmentioned information

Table 1: Sample examples of the dataset. For the answer and cited documents we show only part of the content. We underline the selected statement in the answer. We mark in red and blue the key information associated with the label in the statement and the cited documents. We provide English translations of the questions and key information.

## 2.3 Two-stage Manual Annotation

The original samples need to be manually labeled as positive or negative samples before they can be used to form the dataset (examples are shown in Table 1). In the LLM augmentation phase, although we try to guide the LLM to augment negative samples with qualified quality, the LLM may generate some samples that do not meet the requirements. Therefore, the augmented samples also need to be manually labeled for compliance before they can be used to form the dataset. The goal of the two-stage manual annotation is to complete the manual annotation needed above.

In the first stage, the annotators (from the professional data annotation institution in China) need to label whether the original sample is a positive or negative sample, i.e., to determine whether the sum of the information provided by the cited documents fully supports the statement. In order to reduce the difficulty of labeling, the information of key segments labeled by LLM will be provided to the annotators as a reference. However, since the LLM labeling is not always accurate, if the annotators are unable to make a judgment after reading the key segments, they still need to read other parts of the documents to make a judgment. In this stage, the number of negative samples identified by the annotation is 1,006, with a negative sample rate of

about 9%. We randomly selected 2,000 samples (1,000 negative and 1,000 positive) and split them equally to create the development and test sets. The augmented samples corresponding to the positive samples in the remaining original samples will be labeled in the second stage.

In the second stage, the annotators need to determine whether an augmented sample is of acceptable quality and whether it is a negative sample. In order to reduce the difficulty of labeling, we show the annotator a comparison of the documents before and after the modification in the form of modification traces. Among the augmented samples that the annotators determine to be negative samples of acceptable quality, we select 2,449 samples that use the modification methods of changing information and deleting information respectively, totaling 4,898 samples. These augmented negative samples together with the 4,898 positive samples in the original samples identified by the first stage of annotation constitute the training set. The two-stage manual annotation costs 0.5\$ per sample. See Appendix C for instructions for annotators.

## 3 Experiments

In our experiments, we evaluate the dataset using two approaches. First, we assess the zero-shot performance of state-of-the-art LLMs on the de-

<b>Dev</b>	<b>Acc</b>	<b>Acc<sub>p</sub></b>	<b>Acc<sub>n</sub></b>
GPT-4o	<b>83.7</b>	97.0	70.4
Qwen2.5-Plus	81.6	97.0	66.2
DeepSeek-v3	69.4	99.2	39.6
Llama-3.1-8B	<b>91.4</b>	91.6	91.2
Mistral-7B	89.5	91.2	87.8
Qwen2.5-7B	91.2	95.0	87.4
<b>Test</b>	<b>Acc</b>	<b>Acc<sub>p</sub></b>	<b>Acc<sub>n</sub></b>
GPT-4o	<b>83.9</b>	96.2	71.6
Qwen2.5-Plus	81.2	94.8	67.6
DeepSeek-v3	69.4	99.4	39.4
Llama-3.1-8B	<b>90.6</b>	90.4	90.8
Mistral-7B	89.8	92.0	87.6
Qwen2.5-7B	88.5	90.4	86.6

Table 2: Results of experiments on the dev set and the test set. We report overall accuracy (**Acc**), accuracy on positive samples (**Acc<sub>p</sub>**), and accuracy on negative samples (**Acc<sub>n</sub>**) in percentage form.

development and test sets. This aims to highlight the challenge posed by the test samples. Second, due to resource constraints, we conduct parameter-efficient fine-tuning on smaller models using the training data. This focuses on demonstrating the effectiveness of the training samples.

### 3.1 Settings

State-of-the-art LLMs that we use for zero-shot performance tests include GPT-4o (OpenAI et al., 2024), Qwen2.5-Plus (Qwen et al., 2024), and DeepSeek-v3 (DeepSeek-AI et al., 2024). We provide the sample to the LLMs and ask for their judgment. The relatively small language models we use for training include Llama-3.1-8B (Grattafiori et al., 2024), Mistral-7B (Jiang et al., 2023), and Qwen2.5-7B (Qwen et al., 2024). The parameter-efficient fine-tuning method we use is LoRA (Hu et al., 2022). See Appendix D for training details. We use accuracy as the metric. Since there are equal numbers of positive and negative samples, the accuracy is equivalent to the commonly used balanced accuracy (Luo et al., 2023), which is the average of the accuracy on positive and negative samples. We also report the accuracy of positive and negative samples separately.

### 3.2 Results

Table 2 reveals significant differences in performance between LLMs tested under zero-shot conditions and smaller models fine-tuned with parameter-efficient methods. Among the zero-shot LLMs,

GPT-4o achieved the highest overall accuracy, outperforming Qwen2.5-Plus and DeepSeek-v3. However, even GPT-4o struggled with negative samples, achieving only 70.4% accuracy on the dev set and 71.6% on the test set. This limitation highlights a persistent challenge in distinguishing negative cases, which significantly impacts overall accuracy. DeepSeek-v3, while demonstrating near-perfect accuracy on positive samples, performed poorly on negative samples (39.6% dev, 39.4% test), indicating a clear trade-off between the two categories.

In contrast, smaller models fine-tuned with the training set achieved remarkable improvements, particularly in handling negative samples. Llama-3.1-8B stood out as the top performer, achieving 91.4% accuracy on the dev set and 90.6% on the test set, while maintaining a strong balance between positive and negative samples. These results suggest that the training data effectively addressed the challenges posed by negative samples, enabling the fine-tuned models to achieve significantly higher overall accuracy. Overall, the results underscore the effectiveness of fine-tuning in improving model robustness, particularly for negative samples. The dataset’s training data appears to play a crucial role in enhancing model performance, as evidenced by the fine-tuned models’ ability to achieve high accuracy across both positive and negative samples. These insights suggest that tailored training strategies and targeted fine-tuning can significantly enhance model capabilities, even for smaller models.

## 4 Conclusion

In this work, we propose the first large-scale Chinese dataset CITECHECK for citation faithfulness detection. To solve the high-cost problem caused by the lack of negative samples when constructing the dataset using strong RAG systems, we propose the method of data augmentation with two-stage manual annotation. This method allows us to construct a dataset with a balanced number of positive and negative samples at a relatively low cost and guarantees the quality of the dataset. We conduct experiments and validate the quality of the dataset in two aspects: (1) the test samples consisting of the original samples are challenging for detection, and (2) the training samples consisting of the original positive samples and the augmented negative samples can be effectively applied for training.

## Limitations

The main limitation of the dataset is the availability of only binary judgment labels (positive or negative). We do not manually label which part of the statement in the negative sample is unsupported, nor do we manually label the evidence in the documents that the statement in the positive sample is supported. However, key segments labeling and modifications in the LLM augmentation phase are available, which compensates for the limitation to some extent.

The main limitation of the experiments is the lack of more experiments on other test sets for the model obtained from training to show the generalization performance. This limitation comes from the lack of relevant Chinese datasets. We will continue to track the relevant Chinese datasets proposed and conduct experiments.

## Ethics Statement

We comply with the license to use language models for scientific research purposes only. Questions are collected with the permission of the license of open-source datasets or with the consent of the relevant users. The datasets we construct will also be open source for scientific research purposes. We conduct checks to minimize potential risk issues with datasets, including personal privacy concerns and harmful content.

The AI assistant we use in our work is Copilot (for simple code completion).

## References

Sebastian Borgeaud, Arthur Mensch, Jordan Hoffmann, Trevor Cai, Eliza Rutherford, Katie Millican, George 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 W. Rae, Erich Elsen, and Laurent Sifre. 2022. [Improving language models by retrieving from trillions of tokens](#). In *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, pages 2206–2240. PMLR.

Jiawei Chen, Hongyu Lin, Xianpei Han, and Le Sun. 2024. [Benchmarking large language models in retrieval-augmented generation](#). In *Thirty-Eighth AAAI Conference on Artificial Intelligence, AAAI*

*2024, Thirty-Sixth Conference on Innovative Applications of Artificial Intelligence, IAAI 2024, Fourteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2014, February 20-27, 2024, Vancouver, Canada*, pages 17754–17762. AAAI Press.

DeepSeek-AI, Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, Damai Dai, Daya Guo, Dejian Yang, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, Guangbo Hao, Guanting Chen, Guowei Li, H. Zhang, Han Bao, Hanwei Xu, Haocheng Wang, Haowei Zhang, Honghui Ding, Huajian Xin, Huazuo Gao, Hui Li, Hui Qu, J. L. Cai, Jian Liang, Jianzhong Guo, Jiaqi Ni, Jiashi Li, Jiawei Wang, Jin Chen, Jingchang Chen, Jingyang Yuan, Junjie Qiu, Junlong Li, Junxiao Song, Kai Dong, Kai Hu, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang, Lecong Zhang, Lei Xu, Leyi Xia, Liang Zhao, Litong Wang, Liyue Zhang, Meng Li, Miaojun Wang, Mingchuan Zhang, Minghua Zhang, Minghui Tang, Mingming Li, Ning Tian, Panpan Huang, Peiyi Wang, Peng Zhang, Qiancheng Wang, Qihao Zhu, Qinyu Chen, Qiushi Du, R. J. Chen, R. L. Jin, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang, Runxin Xu, Ruoyu Zhang, Ruyi Chen, S. S. Li, Shanghao Lu, Shangyan Zhou, Shanhua Chen, Shaoqing Wu, Shengfeng Ye, Shengfeng Ye, Shirong Ma, Shiyu Wang, Shuang Zhou, Shuiping Yu, Shunfeng Zhou, Shuting Pan, T. Wang, Tao Yun, Tian Pei, Tianyu Sun, W. L. Xiao, Wangding Zeng, Wanjia Zhao, Wei An, Wen Liu, Wenfeng Liang, Wenjun Gao, Wenqin Yu, Wentao Zhang, X. Q. Li, Xiangyue Jin, Xianzu Wang, Xiao Bi, Xiaodong Liu, Xiaohan Wang, Xiaojin Shen, Xiaokang Chen, Xiaokang Zhang, Xiaosha Chen, Xiaotao Nie, Xiaowen Sun, Xiaoxiang Wang, Xin Cheng, Xin Liu, Xin Xie, Xingchao Liu, Xingkai Yu, Xinnan Song, Xinxia Shan, Xinyi Zhou, Xinyu Yang, Xinyuan Li, Xuecheng Su, Xuheng Lin, Y. K. Li, Y. Q. Wang, Y. X. Wei, Y. X. Zhu, Yang Zhang, Yanhong Xu, Yanhong Xu, Yanping Huang, Yao Li, Yao Zhao, Yaofeng Sun, Yaohui Li, Yaohui Wang, Yi Yu, Yi Zheng, Yichao Zhang, Yifan Shi, Yiliang Xiong, Ying He, Ying Tang, Yishi Piao, Yisong Wang, Yixuan Tan, Yiyang Ma, Yiyuan Liu, Yongqiang Guo, Yu Wu, Yuan Ou, Yuchen Zhu, Yuduan Wang, Yue Gong, Yuheng Zou, Yujia He, Yukun Zha, Yunfan Xiong, Yunxian Ma, Yuting Yan, Yuxiang Luo, Yuxiang You, Yuxuan Liu, Yuyang Zhou, Z. F. Wu, Z. Z. Ren, Zehui Ren, Zhangli Sha, Zhe Fu, Zhean Xu, Zhen Huang, Zhen Zhang, Zhenda Xie, Zhengyan Zhang, Zhewen Hao, Zhibin Gou, Zhicheng Ma, Zhi-gang Yan, Zhihong Shao, Zhipeng Xu, Zhiyu Wu, Zhongyu Zhang, Zhuoshu Li, Zihui Gu, Zijia Zhu, Zijun Liu, Zilin Li, Ziwei Xie, Ziyang Song, Ziyi Gao, and Zizheng Pan. 2024. [Deepseek-v3 technical report](#). Preprint, arXiv:2412.19437.

Luyu Gao, Zhuyun Dai, Panupong Pasupat, Anthony Chen, Arun Tejasvi Chaganty, Yicheng Fan, Vincent Y. Zhao, Ni Lao, Hongrae Lee, Da-Cheng Juan, and Kelvin Guu. 2023a. [RARR: researching and revising what language models say, using language](#)

428	models.	In <i>Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)</i> , ACL 2023, Toronto, Canada, July 9-14, 2023, pages 16477–16508. Association for Computational Linguistics.	490
429			491
430			492
431			493
432			494
433	Tianyu Gao, Howard Yen, Jiatong Yu, and Danqi Chen.		495
434	2023b. Enabling large language models to generate text with citations.	In <i>Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023</i> , pages 6465–6488. Association for Computational Linguistics.	496
435			497
436			498
437			499
438			500
439			501
440	Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri,		502
441	Abhinav Pandey, Abhishek Kadian, Ahmad Al-		503
442	Dahle, Aiesha Letman, Akhil Mathur, Alan Schel-		504
443	ten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh		505
444	Goyal, Anthony Hartshorn, Aobo Yang, Archi Mi-		506
445	tra, Archie Sravankumar, Artem Korenev, Arthur		507
446	Hinsvark, Arun Rao, Aston Zhang, Aurelien Ro-		508
447	driguez, Austen Gregerson, Ava Spataru, Baptiste		509
448	Roziere, Bethany Biron, Binh Tang, Bobbie Chern,		510
449	Charlotte Caucheteux, Chaya Nayak, Chloe Bi,		511
450	Chris Marra, Chris McConnell, Christian Keller,		512
451	Christophe Touret, Chunyang Wu, Corinne Wong,		513
452	Cristian Canton Ferrer, Cyrus Nikolaidis, Damien Al-		514
453	lonsius, Daniel Song, Danielle Pintz, Danny Livshits,		515
454	Danny Wyatt, David Esiobu, Dhruv Choudhary,		516
455	Dhruv Mahajan, Diego Garcia-Olano, Diego Perino,		517
456	Dieuwke Hupkes, Egor Lakomkin, Ehab AlBadawy,		518
457	Elina Lobanova, Emily Dinan, Eric Michael Smith,		519
458	Filip Radenovic, Francisco Guzmán, Frank Zhang,		520
459	Gabriel Synnaeve, Gabrielle Lee, Georgia Lewis An-		521
460	derson, Govind Thattai, Graeme Nail, Gregoire Mi-		522
461	alon, Guan Pang, Guillem Cucurell, Hailey Nguyen,		523
462	Hannah Korevaar, Hu Xu, Hugo Touvron, Iliyan		524
463	Zarov, Imanol Arrieta Ibarra, Isabel Kloumann, Is-		525
464	han Misra, Ivan Evtimov, Jack Zhang, Jade Copet,		526
465	Jaewon Lee, Jan Geffert, Jana Vranes, Jason Park,		527
466	Jay Mahadeokar, Jeet Shah, Jelmer van der Linde,		528
467	Jennifer Billock, Jenny Hong, Jenya Lee, Jeremy Fu,		529
468	Jianfeng Chi, Jianyu Huang, Jiawen Liu, Jie Wang,		530
469	Jiecao Yu, Joanna Bitton, Joe Spisak, Jongsoo Park,		531
470	Joseph Rocca, Joshua Johnstun, Joshua Saxe, Jun-		532
471	teng Jia, Kalyan Vasudevan Alwala, Karthik Prasad,		533
472	Kartikeya Upasani, Kate Plawiak, Ke Li, Kenneth		534
473	Heafield, Kevin Stone, Khalid El-Arini, Krithika Iyer,		535
474	Kshitiz Malik, Kuenley Chiu, Kunal Bhalla, Kushal		536
475	Lakhotia, Lauren Rantala-Yeary, Laurens van der		537
476	Maaten, Lawrence Chen, Liang Tan, Liz Jenkins,		538
477	Louis Martin, Lovish Madaan, Lubo Malo, Lukas		539
478	Blecher, Lukas Landzaat, Luke de Oliveira, Madeline		540
479	Muzzi, Mahesh Pasupuleti, Mannat Singh, Manohar		541
480	Paluri, Marcin Kardas, Maria Tsimpoukelli, Mathew		542
481	Oldham, Mathieu Rita, Maya Pavlova, Melanie Kam-		543
482	badur, Mike Lewis, Min Si, Mitesh Kumar Singh,		544
483	Mona Hassan, Naman Goyal, Narjes Torabi, Niko-		545
484	lay Bashlykov, Nikolay Bogoychev, Niladri Chatterji,		546
485	Ning Zhang, Olivier Duchenne, Onur Çelebi, Patrick		547
486	Alrassy, Pengchuan Zhang, Pengwei Li, Petar Va-		548
487	sic, Peter Weng, Prajjwal Bhargava, Pratik Dubal,		549
488	Praveen Krishnan, Punit Singh Koura, Puxin Xu,		550
489	Qing He, Qingxiao Dong, Ragavan Srinivasan, Raj		551
			552
			553

554	nifer Chan, Jenny Zhen, Jeremy Reizenstein, Jeremy Teboul, Jessica Zhong, Jian Jin, Jingyi Yang, Joe Cummings, Jon Carvill, Jon Shepard, Jonathan McPhie, Jonathan Torres, Josh Ginsburg, Junjie Wang, Kai Wu, Kam Hou U, Karan Saxena, Kartikay Khan-delwal, Katayoun Zand, Kathy Matosich, Kaushik Veeraraghavan, Kelly Michelena, Keqian Li, Kiran Jagadeesh, Kun Huang, Kunal Chawla, Kyle Huang, Lailin Chen, Lakshya Garg, Lavender A, Leandro Silva, Lee Bell, Lei Zhang, Liangpeng Guo, Licheng Yu, Liron Moshkovich, Luca Wehrstedt, Madian Khabsa, Manav Avalani, Manish Bhatt, Martynas Mankus, Matan Hasson, Matthew Lennie, Matthias Reso, Maxim Groshev, Maxim Naumov, Maya Lathi, Meghan Keneally, Miao Liu, Michael L. Seltzer, Michal Valko, Michelle Restrepo, Mihir Patel, Mik Vyatskov, Mikayel Samvelyan, Mike Clark, Mike Macey, Mike Wang, Miquel Jubert Hermoso, Mo Metanat, Mohammad Rastegari, Munish Bansal, Nandhini Santhanam, Natascha Parks, Natasha White, Navyata Bawa, Nayan Singhal, Nick Egebo, Nicolas Usunier, Nikhil Mehta, Nikolay Pavlovich Laptev, Ning Dong, Norman Cheng, Oleg Chernoguz, Olivia Hart, Omkar Salpekar, Ozlem Kalinli, Parkin Kent, Parth Parekh, Paul Saab, Pavan Balaji, Pedro Rittner, Philip Bontrager, Pierre Roux, Piotr Dollar, Polina Zvyagina, Prashant Ratanchandani, Pritish Yuvaraj, Qian Liang, Rachad Alao, Rachel Rodriguez, Rafi Ayub, Raghotham Murthy, Raghu Nayani, Rahul Mitra, Rangaprabhu Parthasarathy, Raymond Li, Rebekkah Hogan, Robin Battey, Rocky Wang, Russ Howes, Ruty Rinott, Sachin Mehta, Sachin Siby, Sai Jayesh Bondu, Samyak Datta, Sara Chugh, Sara Hunt, Sargun Dhillon, Sasha Sidorov, Satadru Pan, Saurabh Mahajan, Saurabh Verma, Seiji Yamamoto, Sharadh Ramaswamy, Shaun Lindsay, Shaun Lindsay, Sheng Feng, Shenghao Lin, Shengxin Cindy Zha, Shishir Patil, Shiva Shankar, Shuqiang Zhang, Shuqiang Zhang, Sinong Wang, Sneha Agarwal, Soji Sajuyigbe, Soumith Chintala, Stephanie Max, Stephen Chen, Steve Kehoe, Steve Satterfield, Sudarshan Govindaprasad, Sumit Gupta, Summer Deng, Sungmin Cho, Sunny Virk, Suraj Subramanian, Sy Choudhury, Sydney Goldman, Tal Remez, Tamar Glaser, Tamara Best, Thilo Koehler, Thomas Robinson, Tianhe Li, Tianjun Zhang, Tim Matthews, Timothy Chou, Tzook Shaked, Varun Vontimitta, Victoria Ajayi, Victoria Montanez, Vijai Mohan, Vinay Satish Kumar, Vishal Mangla, Vlad Ionescu, Vlad Poenaru, Vlad Tiberiu Mihailescu, Vladimir Ivanov, Wei Li, Wencheng Wang, Wenwen Jiang, Wes Bouaziz, Will Constable, Xiaocheng Tang, Xiaojian Wu, Xiaolan Wang, Xilun Wu, Xinbo Gao, Yaniv Kleinman, Yanjun Chen, Ye Hu, Ye Jia, Ye Qi, Yenda Li, Yelin Zhang, Ying Zhang, Yossi Adi, Youngjin Nam, Yu, Wang, Yu Zhao, Yuchen Hao, Yundi Qian, Yunlu Li, Yuzi He, Zach Rait, Zachary DeVito, Zef Rosnbrick, Zhaoduo Wen, Zhenyu Yang, Zhiwei Zhao, and Zhiyu Ma. 2024. <b>The llama 3 herd of models.</b> Preprint, arXiv:2407.21783.	613
614	Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Ming-Wei Chang. 2020. <b>Retrieval augmented language model pre-training.</b> In <i>Proceedings of the</i>	615
616	<i>37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event, volume 119 of Proceedings of Machine Learning Research</i> , pages 3929–3938. PMLR.	616
617	Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. <b>Lora: Low-rank adaptation of large language models.</b> In <i>The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022</i> . OpenReview.net.	621
618	Xuming Hu, Junzhe Chen, Xiaochuan Li, Yufei Guo, Lijie Wen, Philip S. Yu, and Zhijiang Guo. 2024a. <b>Towards understanding factual knowledge of large language models.</b> In <i>The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024</i> . OpenReview.net.	622
619	Xuming Hu, Xiaochuan Li, Junzhe Chen, Yinghui Li, Yangning Li, Xiaoguang Li, Yasheng Wang, Qun Liu, Lijie Wen, Philip Yu, and Zhijiang Guo. 2024b. <b>Evaluating robustness of generative search engine on adversarial factoid questions.</b> In <i>Findings of the Association for Computational Linguistics: ACL 2024</i> , pages 10650–10671, Bangkok, Thailand. Association for Computational Linguistics.	623
620	Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, Lélio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. 2023. <b>Mistral 7b.</b> Preprint, arXiv:2310.06825.	624
621	Mojtaba Komeili, Kurt Shuster, and Jason Weston. 2022. <b>Internet-augmented dialogue generation.</b> In <i>Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2022, Dublin, Ireland, May 22-27, 2022</i> , pages 8460–8478. Association for Computational Linguistics.	625
622	Patrick S. H. Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. <b>Retrieval-augmented generation for knowledge-intensive NLP tasks.</b> In <i>Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual</i> .	623
623	Huanshuo Liu, Hao Zhang, Zhijiang Guo, Kuicai Dong, Xiangyang Li, Yi Quan Lee, Cong Zhang, and Yong Liu. 2024. <b>Crla: Adaptive retrieval-augmented generation via probe-guided control.</b> <i>CoRR</i> , abs/2405.18727.	624
624	Nelson F. Liu, Tianyi Zhang, and Percy Liang. 2023. <b>Evaluating verifiability in generative search engines.</b> In <i>Findings of the Association for Computational Linguistics: EMNLP 2023, Singapore, December 6-10, 2023</i> .	625
625	Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Ming-Wei Chang. 2020. <b>Retrieval augmented language model pre-training.</b> In <i>Proceedings of the</i>	626
626	<i>37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event, volume 119 of Proceedings of Machine Learning Research</i> , pages 3929–3938. PMLR.	627
627	Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. <b>Lora: Low-rank adaptation of large language models.</b> In <i>The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022</i> . OpenReview.net.	628
628	Xuming Hu, Junzhe Chen, Xiaochuan Li, Yufei Guo, Lijie Wen, Philip S. Yu, and Zhijiang Guo. 2024a. <b>Towards understanding factual knowledge of large language models.</b> In <i>The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024</i> . OpenReview.net.	629
629	Xuming Hu, Xiaochuan Li, Junzhe Chen, Yinghui Li, Yangning Li, Xiaoguang Li, Yasheng Wang, Qun Liu, Lijie Wen, Philip Yu, and Zhijiang Guo. 2024b. <b>Evaluating robustness of generative search engine on adversarial factoid questions.</b> In <i>Findings of the Association for Computational Linguistics: ACL 2024</i> , pages 10650–10671, Bangkok, Thailand. Association for Computational Linguistics.	630
630	Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, Lélio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. 2023. <b>Mistral 7b.</b> Preprint, arXiv:2310.06825.	631
631	Mojtaba Komeili, Kurt Shuster, and Jason Weston. 2022. <b>Internet-augmented dialogue generation.</b> In <i>Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2022, Dublin, Ireland, May 22-27, 2022</i> , pages 8460–8478. Association for Computational Linguistics.	632
632	Patrick S. H. Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. <b>Retrieval-augmented generation for knowledge-intensive NLP tasks.</b> In <i>Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual</i> .	633
633	Huanshuo Liu, Hao Zhang, Zhijiang Guo, Kuicai Dong, Xiangyang Li, Yi Quan Lee, Cong Zhang, and Yong Liu. 2024. <b>Crla: Adaptive retrieval-augmented generation via probe-guided control.</b> <i>CoRR</i> , abs/2405.18727.	634
634	Nelson F. Liu, Tianyi Zhang, and Percy Liang. 2023. <b>Evaluating verifiability in generative search engines.</b> In <i>Findings of the Association for Computational Linguistics: EMNLP 2023, Singapore, December 6-10, 2023</i> .	635
635	Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Ming-Wei Chang. 2020. <b>Retrieval augmented language model pre-training.</b> In <i>Proceedings of the</i>	636
636	<i>37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event, volume 119 of Proceedings of Machine Learning Research</i> , pages 3929–3938. PMLR.	637
637	Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. <b>Lora: Low-rank adaptation of large language models.</b> In <i>The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022</i> . OpenReview.net.	638
638	Xuming Hu, Junzhe Chen, Xiaochuan Li, Yufei Guo, Lijie Wen, Philip S. Yu, and Zhijiang Guo. 2024a. <b>Towards understanding factual knowledge of large language models.</b> In <i>The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024</i> . OpenReview.net.	639
639	Xuming Hu, Xiaochuan Li, Junzhe Chen, Yinghui Li, Yangning Li, Xiaoguang Li, Yasheng Wang, Qun Liu, Lijie Wen, Philip Yu, and Zhijiang Guo. 2024b. <b>Evaluating robustness of generative search engine on adversarial factoid questions.</b> In <i>Findings of the Association for Computational Linguistics: ACL 2024</i> , pages 10650–10671, Bangkok, Thailand. Association for Computational Linguistics.	640
640	Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, Lélio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. 2023. <b>Mistral 7b.</b> Preprint, arXiv:2310.06825.	641
641	Mojtaba Komeili, Kurt Shuster, and Jason Weston. 2022. <b>Internet-augmented dialogue generation.</b> In <i>Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2022, Dublin, Ireland, May 22-27, 2022</i> , pages 8460–8478. Association for Computational Linguistics.	642
642	Patrick S. H. Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. <b>Retrieval-augmented generation for knowledge-intensive NLP tasks.</b> In <i>Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual</i> .	643
643	Huanshuo Liu, Hao Zhang, Zhijiang Guo, Kuicai Dong, Xiangyang Li, Yi Quan Lee, Cong Zhang, and Yong Liu. 2024. <b>Crla: Adaptive retrieval-augmented generation via probe-guided control.</b> <i>CoRR</i> , abs/2405.18727.	644
644	Nelson F. Liu, Tianyi Zhang, and Percy Liang. 2023. <b>Evaluating verifiability in generative search engines.</b> In <i>Findings of the Association for Computational Linguistics: EMNLP 2023, Singapore, December 6-10, 2023</i> .	645
645	Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Ming-Wei Chang. 2020. <b>Retrieval augmented language model pre-training.</b> In <i>Proceedings of the</i>	646
646	<i>37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event, volume 119 of Proceedings of Machine Learning Research</i> , pages 3929–3938. PMLR.	647
647	Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. <b>Lora: Low-rank adaptation of large language models.</b> In <i>The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022</i> . OpenReview.net.	648
648	Xuming Hu, Junzhe Chen, Xiaochuan Li, Yufei Guo, Lijie Wen, Philip S. Yu, and Zhijiang Guo. 2024a. <b>Towards understanding factual knowledge of large language models.</b> In <i>The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024</i> . OpenReview.net.	649
649	Xuming Hu, Xiaochuan Li, Junzhe Chen, Yinghui Li, Yangning Li, Xiaoguang Li, Yasheng Wang, Qun Liu, Lijie Wen, Philip Yu, and Zhijiang Guo. 2024b. <b>Evaluating robustness of generative search engine on adversarial factoid questions.</b> In <i>Findings of the Association for Computational Linguistics: ACL 2024</i> , pages 10650–10671, Bangkok, Thailand. Association for Computational Linguistics.	650
650	Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, Lélio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. 2023. <b>Mistral 7b.</b> Preprint, arXiv:2310.06825.	651
651	Mojtaba Komeili, Kurt Shuster, and Jason Weston. 2022. <b>Internet-augmented dialogue generation.</b> In <i>Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2022, Dublin, Ireland, May 22-27, 2022</i> , pages 8460–8478. Association for Computational Linguistics.	652
652	Patrick S. H. Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. <b>Retrieval-augmented generation for knowledge-intensive NLP tasks.</b> In <i>Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual</i> .	653
653	Huanshuo Liu, Hao Zhang, Zhijiang Guo, Kuicai Dong, Xiangyang Li, Yi Quan Lee, Cong Zhang, and Yong Liu. 2024. <b>Crla: Adaptive retrieval-augmented generation via probe-guided control.</b> <i>CoRR</i> , abs/2405.18727.	654
654	Nelson F. Liu, Tianyi Zhang, and Percy Liang. 2023. <b>Evaluating verifiability in generative search engines.</b> In <i>Findings of the Association for Computational Linguistics: EMNLP 2023, Singapore, December 6-10, 2023</i> .	655
655	Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Ming-Wei Chang. 2020. <b>Retrieval augmented language model pre-training.</b> In <i>Proceedings of the</i>	656
656	<i>37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event, volume 119 of Proceedings of Machine Learning Research</i> , pages 3929–3938. PMLR.	657
657	Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. <b>Lora: Low-rank adaptation of large language models.</b> In <i>The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022</i> . OpenReview.net.	658
658	Xuming Hu, Junzhe Chen, Xiaochuan Li, Yufei Guo, Lijie Wen, Philip S. Yu, and Zhijiang Guo. 2024a. <b>Towards understanding factual knowledge of large language models.</b> In <i>The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024</i> . OpenReview.net.	659
659	Xuming Hu, Xiaochuan Li, Junzhe Chen, Yinghui Li, Yangning Li, Xiaoguang Li, Yasheng Wang, Qun Liu, Lijie Wen, Philip Yu, and Zhijiang Guo. 2024b. <b>Evaluating robustness of generative search engine on adversarial factoid questions.</b> In <i>Findings of the Association for Computational Linguistics: ACL 2024</i> , pages 10650–10671, Bangkok, Thailand. Association for Computational Linguistics.	660
660	Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, Lélio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. 2023. <b>Mistral 7b.</b> Preprint, arXiv:2310.06825.	661
661	Mojtaba Komeili, Kurt Shuster, and Jason Weston. 2022. <b>Internet-augmented dialogue generation.</b> In <i>Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2022, Dublin, Ireland, May 22-27, 2022</i> , pages 8460–8478. Association for Computational Linguistics.	662
662	Patrick S. H. Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. <b>Retrieval-augmented generation for knowledge-intensive NLP tasks.</b> In <i>Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual</i> .	663
663	Huanshuo Liu, Hao Zhang, Zhijiang Guo, Kuicai Dong, Xiangyang Li, Yi Quan Lee, Cong Zhang, and Yong Liu. 2024. <b>Crla: Adaptive retrieval-augmented generation via probe-guided control.</b> <i>CoRR</i> , abs/2405.18727.	664
664	Nelson F. Liu, Tianyi Zhang, and Percy Liang. 2023. <b>Evaluating verifiability in generative search engines.</b> In <i>Findings of the Association for Computational Linguistics: EMNLP 2023, Singapore, December 6-10, 2023</i> .	665
665	Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Ming-Wei Chang. 2020. <b>Retrieval augmented language model pre-training.</b> In <i>Proceedings of the</i>	666
666	<i>37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event, volume 119 of Proceedings of Machine Learning Research</i> , pages 3929–3938. PMLR.	667
667	Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. <b>Lora: Low-rank adaptation of large language models.</b> In <i>The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022</i> . OpenReview.net.	668
668	Xuming Hu, Junzhe Chen, Xiaochuan Li, Yufei Guo, Lijie Wen, Philip S. Yu, and Zhijiang Guo. 2024a. <b>Towards understanding factual knowledge of large language models.</b> In <i>The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024</i> . OpenReview.net.	669
669	Xuming Hu, Xiaochuan Li, Junzhe Chen, Yinghui Li, Yangning Li, Xiaoguang Li, Yasheng Wang, Qun Liu, Lijie Wen, Philip Yu, and Zhijiang Guo. 2024b. <b>Evaluating robustness of generative search engine on adversarial factoid questions.</b> In <i>Findings of the Association for Computational Linguistics: ACL 2024</i> , pages 10650–10671, Bangkok, Thailand. Association for Computational Linguistics.	670
670	Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, Lélio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. 2023. <b>Mistral 7b.</b> Preprint, arXiv:2310.06825.	671
671	Mojtaba Komeili, Kurt Shuster, and Jason Weston. 2022. <b>Internet-augmented dialogue generation.</b> In <i>Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2022, Dublin, Ireland, May 22-27, 2022</i> , pages 8460–8478. Association for Computational Linguistics.	672
672	Patrick S. H. Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. <b>Retrieval-augmented generation for knowledge-intensive NLP tasks.</b> In <i>Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual</i> .	673
673	Huanshuo Liu, Hao Zhang, Zhijiang Guo, Kuicai Dong, Xiangyang Li, Yi Quan Lee, Cong Zhang, and Yong Liu. 2024. <b>Crla: Adaptive retrieval-augmented generation via probe-guided control.</b> <i>CoRR</i> , abs/2405.18727.	674

674	2023, pages 7001–7025. Association for Computational Linguistics.	735
675		736
676	Zheheng Luo, Qianqian Xie, and Sophia Ananiadou.	737
677	2023. Chatgpt as a factual inconsistency evaluator for abstractive text summarization. <i>CoRR</i> ,	738
678	abs/2303.15621.	739
679		740
680	Joshua Maynez, Shashi Narayan, Bernd Bohnet, and	741
681	Ryan T. McDonald. 2020. On faithfulness and factu-	742
682	tality in abstractive summarization. In <i>Proceedings</i>	743
683	of the 58th Annual Meeting of the Association for	744
684	Computational Linguistics, ACL 2020, Online, July	745
685	5-10, 2020, pages 1906–1919. Association for Com-	746
686	putational Linguistics.	747
687	Reiichiro Nakano, Jacob Hilton, Suchir Balaji, Jeff Wu,	748
688	Long Ouyang, Christina Kim, Christopher Hesse,	749
689	Shantanu Jain, Vineet Kosaraju, William Saunders,	750
690	Xu Jiang, Karl Cobbe, Tyna Eloundou, Gretchen	751
691	Krueger, Kevin Button, Matthew Knight, Benjamin	752
692	Chess, and John Schulman. 2021. Webgpt: Browser-	753
693	assisted question-answering with human feedback.	754
694	<i>CoRR</i> , abs/2112.09332.	755
695	OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal,	756
696	Lama Ahmad, Ilge Akkaya, Florencia Leoni Ale-	757
697	man, Diogo Almeida, Janko Altenschmidt, Sam Alt-	758
698	man, Shyamal Anadkat, Red Avila, Igor Babuschkin,	759
699	Suchir Balaji, Valerie Balcom, Paul Baltescu, Haim-	760
700	ing Bao, Mohammad Bavarian, Jeff Belgum, Ir-	761
701	wan Bello, Jake Berdine, Gabriel Bernadett-Shapiro,	762
702	Christopher Berner, Lenny Bogdonoff, Oleg Boiko,	763
703	Madelaine Boyd, Anna-Luisa Brakman, Greg Brock-	764
704	man, Tim Brooks, Miles Brundage, Kevin Button,	765
705	Trevor Cai, Rosie Campbell, Andrew Cann, Brittany	766
706	Carey, Chelsea Carlson, Rory Carmichael, Brooke	767
707	Chan, Che Chang, Fotis Chantzis, Derek Chen, Sully	768
708	Chen, Ruby Chen, Jason Chen, Mark Chen, Ben	769
709	Chess, Chester Cho, Casey Chu, Hyung Won Chung,	770
710	Dave Cummings, Jeremiah Currier, Yunxing Dai,	771
711	Cory Decareaux, Thomas Degry, Noah Deutsch,	772
712	Damien Deville, Arka Dhar, David Dohan, Steve	773
713	Dowling, Sheila Dunning, Adrien Ecoffet, Atty Eleti,	774
714	Tyna Eloundou, David Farhi, Liam Fedus, Niko Felix,	775
715	Simón Posada Fishman, Juston Forte, Isabella Ful-	776
716	ford, Leo Gao, Elie Georges, Christian Gibson, Vik	777
717	Goel, Tarun Gogineni, Gabriel Goh, Rapha Gontijo-	778
718	Lopes, Jonathan Gordon, Morgan Grafstein, Scott	779
719	Gray, Ryan Greene, Joshua Gross, Shixiang Shane	780
720	Gu, Yufei Guo, Chris Hallacy, Jesse Han, Jeff Harris,	781
721	Yuchen He, Mike Heaton, Johannes Heidecke, Chris	782
722	Hesse, Alan Hickey, Wade Hickey, Peter Hoeschele,	
723	Brandon Houghton, Kenny Hsu, Shengli Hu, Xin	
724	Hu, Joost Huizinga, Shantanu Jain, Shawn Jain,	
725	Joanne Jang, Angela Jiang, Roger Jiang, Haozhun	
726	Jin, Denny Jin, Shino Jomoto, Billie Jonn, Hee-	
727	woo Jun, Tomer Kaftan, Łukasz Kaiser, Ali Ka-	
728	mali, Ingmar Kanitscheider, Nitish Shirish Keskar,	
729	Tabarak Khan, Logan Kilpatrick, Jong Wook Kim,	
730	Christina Kim, Yongjik Kim, Jan Hendrik Kirch-	
731	ner, Jamie Kiros, Matt Knight, Daniel Kokotajlo,	
732	Łukasz Kondraciuk, Andrew Kondrich, Aris Kon-	
733	stantinidis, Kyle Kosic, Gretchen Krueger, Vishal	
734	Kuo, Michael Lampe, Ikai Lan, Teddy Lee, Jan	
	Leike, Jade Leung, Daniel Levy, Chak Ming Li,	
	Rachel Lim, Molly Lin, Stephanie Lin, Mateusz	
	Litwin, Theresa Lopez, Ryan Lowe, Patricia Lue,	
	Anna Makanju, Kim Malfacini, Sam Manning, Todor	
	Markov, Yaniv Markovski, Bianca Martin, Katie	
	Mayer, Andrew Mayne, Bob McGrew, Scott Mayer	
	McKinney, Christine McLeavey, Paul McMillan,	
	Jake McNeil, David Medina, Aalok Mehta, Jacob	
	Menick, Luke Metz, Andrey Mishchenko, Pamela	
	Mishkin, Vinnie Monaco, Evan Morikawa, Daniel	
	Mossing, Tong Mu, Mira Murati, Oleg Murk, David	
	Mély, Ashvin Nair, Reiichiro Nakano, Rajeev Nayak,	
	Arvind Neelakantan, Richard Ngo, Hyeonwoo Noh,	
	Long Ouyang, Cullen O’Keefe, Jakub Pachocki, Alex	
	Paino, Joe Palermo, Ashley Pantuliano, Giambat-	
	tista Parascandolo, Joel Parish, Emy Parparita, Alex	
	Passos, Mikhail Pavlov, Andrew Peng, Adam Perelman,	
	Filipe de Avila Belbute Peres, Michael Petrov,	
	Henrique Ponde 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, Kendra Rimbach,	
	Carl Ross, Bob Rotsted, Henri Roussez, Nick Ry-	
	der, Mario Saltarelli, Ted Sanders, Shibani Santurkar,	
	Girish Sastry, Heather Schmidt, David Schnurr, John	
	Schulman, Daniel Selsam, Kyla Sheppard, Toki	
	Sherbakov, Jessica Shieh, Sarah Shoker, Pranav	
	Shyam, Szymon Sidor, Eric Sigler, Maddie Simens,	
	Jordan Sitkin, Katarina Slama, Ian Sohl, Benjamin	
	Sokolowsky, Yang Song, Natalie Staudacher, Fe-	
	lipe Petroski Such, Natalie Summers, Ilya Sutskever,	
	Jie Tang, Nikolas Tezak, Madeleine B. Thompson,	
	Phil Tillet, Amin Tootoonchian, Elizabeth Tseng,	
	Preston Tuggle, Nick Turley, Jerry Tworek, Juan Fe-	
	lipe Cerón Uribe, Andrea Vallone, Arun Vijayvergiya,	
	Chelsea Voss, Carroll Wainwright, Justin Jay Wang,	
	Alvin Wang, Ben Wang, Jonathan Ward, Jason Wei,	
	CJ Weinmann, Akila Welihinda, Peter Welinder, Ji-	
	ayi Weng, Lilian 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, Qim-	
	ing Yuan, Wojciech Zaremba, Rowan Zellers, Chong	
	Zhang, Marvin Zhang, Shengjia Zhao, Tianhao	
	Zheng, Juntang Zhuang, William Zhuk, and Bar-	
	ret Zoph. 2024. Gpt-4 technical report. Preprint,	
	arXiv:2303.08774.	
	Yujia Qin, Zihan Cai, Dian Jin, Lan Yan, Shihao	783
	Liang, Kunlun Zhu, Yankai Lin, Xu Han, Ning Ding,	784
	Huadong Wang, Ruobing Xie, Fanchao Qi, Zhiyuan	785
	Liu, Maosong Sun, and Jie Zhou. 2023. Webcpm:	786
	Interactive web search for chinese long-form ques-	787
	tion answering. In <i>Proceedings of the 61st Annual</i>	788
	<i>Meeting of the Association for Computational Lin-</i>	789
	<i>guistics (Volume 1: Long Papers), ACL 2023, Toronto,</i>	790
	<i>Canada, July 9-14, 2023</i> , pages 8968–8988. Associa-	791
	tion for Computational Linguistics.	792
	Qwen, :, An Yang, Baosong Yang, Beichen Zhang,	793
	Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li,	794
	Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin,	795
	Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang,	796

797 Jiaxi Yang, Jingren Zhou, Junyang Lin, Kai Dang,  
798 Keming Lu, Keqin Bao, Kexin Yang, Le Yu, Mei  
799 Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men,  
800 Runji Lin, Tianhao Li, Tingyu Xia, Xingzhang Ren,  
801 Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang,  
802 Yu Wan, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and  
803 Zihan Qiu. 2024. [Qwen2.5 technical report](#). *Preprint*,  
804 arXiv:2412.15115.

805 Hannah Rashkin, Vitaly Nikolaev, Matthew Lamm,  
806 Lora Aroyo, Michael Collins, Dipanjan Das, Slav  
807 Petrov, Gaurav Singh Tomar, Iulia Turc, and David  
808 Reitter. 2023. [Measuring attribution in natural lan-](#)  
809 [guage generation models](#). *Comput. Linguistics*,  
810 49(4):777–840.

811 Haochen Tan, Zhijiang Guo, Zhan Shi, Lu Xu, Zhili Liu,  
812 Yunlong Feng, Xiaoguang Li, Yasheng Wang, Lifeng  
813 Shang, Qun Liu, and Linqi Song. 2024. [Proxyqa: An](#)  
814 [alternative framework for evaluating long-form text](#)  
815 [generation with large language models](#). In *Proceed-*  
816 *ings of the 62nd Annual Meeting of the Association*  
817 *for Computational Linguistics (Volume 1: Long Pa-*  
818 *pers)*, ACL 2024, Bangkok, Thailand, August 11-16,  
819 2024, pages 6806–6827. Association for Compu-  
820 tational Linguistics.

821 Rui Wang. 2023. [Dataset: Zhihu-kol](#).

822 Bright Xu. 2019. [Nlp chinese corpus: Large scale chi-](#)  
823 [nese corpus for nlp](#).

824 Xiang Yue, Boshi Wang, Ziru Chen, Kai Zhang, Yu Su,  
825 and Huan Sun. 2023. [Automatic evaluation of attri-](#)  
826 [bution by large language models](#). In *Findings of the*  
827 *Association for Computational Linguistics: EMNLP*  
828 *2023, Singapore, December 6-10, 2023*, pages 4615–  
829 4635. Association for Computational Linguistics.

## 830 A Unexpected Questions

831 Real-world questions do not always have the correct premises. For example, in the question "水俣  
832 病的传染途径是什么? (What is the route of infec-  
833 tion for Minamata disease?)", Minamata disease  
834 is not an infectious disease. Taking this situation  
835 into account, we add a small number of human-  
836 written questions with incorrect premises and LLM-  
837 generated questions with hard-to-verify premises  
838 in the question collection phase. The number of  
839 these questions in the total number of questions is  
840 about 3%.

## 842 B Prompt for LLM Augmentation

843 See Table 3 for the prompt for LLM augmentation.  
844 Table 4 provides an English version.

## 845 C Instructions for Annotators

### 846 C.1 First Stage

847 In the first stage, we provide the annotators with  
848 the question, answer, statement, and cited doc-  
849 ments. What LLM considers to be key segments  
850 are highlighted in red in the cited documents (see  
851 Figure 1 for an example). We instruct the annota-  
852 tors to follow the process below:

853 (1) First look at the highlighted text. If the high-  
854 lighted text fully supports the statement, then the  
855 annotation is positive; if the highlighted text con-  
856 tradicts the statement, then the annotation is negative.

857 (2) If the annotation cannot be derived from the  
858 highlighted text, then look at the rest of the doc-  
859 ments to make the annotation. When the doc-  
860 ments fully support the statement, the label is  
861 positive, and when there is any information in the  
862 statement that contradicts the documents or infor-  
863 mation that is not mentioned in the documents, the  
864 label is negative.

### 865 C.2 Second Stage

866 In the second stage, we provide the annotator with  
867 the statement and the modified documents. In the  
868 documents, the modified parts are highlighted in  
869 green, where the dashed and crossed-out text is  
870 deleted and the rest is added (see Figure 1 for ex-  
871 amples).

872 For the annotation of whether the quality of the  
873 modification is acceptable, the annotators are in-  
874 structed to note that qualified modifications need to  
875 satisfy the following two requirements: (1) There  
876 are no contradictions within each modified doc-  
877 ument. (2) The modified key segments are fluent in

878 their own right and in the context of the document.  
879 The annotation for support is the same as the first  
880 stage, but based on the modified documents.

## 881 D Input and Training Details

882 We input the statement and the cited documents  
883 into the model and ask the model to determine  
884 whether the statement is fully supported by the doc-  
885 ments, outputting yes or no. For input, we label  
886 and concatenate the cited documents in order (as  
887 shown in Table 1). For training, we use the follow-  
888 ing settings: For training, we use the following  
889 settings: learning rate is 5e-4, number of epochs is  
890 10, scheduler is cosine scheduler, warmup ratio is  
891 0.03, batch size is 256, and LoRA setting is  $r = 8$ ,  
892  $a = 32$  and 0.1 dropout. We report the model  
893 performance for the epoch that achieves the best  
894 performance on the dev set.

## 895 E Related Works

896 Language models are known to produce hallu-  
897 cinations - statements that are inaccurate or un-  
898 founded (Maynez et al., 2020; Hu et al., 2024a).  
899 To address this limitation, recent research has fo-  
900 cused on augmenting LLMs with external tools  
901 such as retrievers (Guu et al., 2020; Borgeaud et al.,  
902 2022; Liu et al., 2024) and search engines (Nakano  
903 et al., 2021; Komeili et al., 2022; Tan et al., 2024).  
904 While this approach suggests that generated con-  
905 tent is supported by external references, the reli-  
906 ability of such attribution requires careful examina-  
907 tion. Recent studies have investigated the validity  
908 of these attributions. Liu et al. (2023) conducted  
909 human evaluations to assess the verifiability of re-  
910 sponds from generative search engines. Hu et al.  
911 (2024b) further investigate the reliability of such  
912 attributions when giving adversarial questions to  
913 RAG systems. Their findings revealed frequent  
914 occurrences of unsupported statements and inac-  
915 curate citations, highlighting the need for rigor-  
916 ous attribution verification (Rashkin et al., 2023).  
917 However, human evaluation processes are resource-  
918 intensive and time-consuming. To overcome these  
919 limitations, existing efforts (Gao et al., 2023a,b)  
920 proposed an automated approach using Natural  
921 Language Inference models to evaluate attribution  
922 accuracy. While several English-language bench-  
923 marks have been developed for this purpose (Yue  
924 et al., 2023), comparable resources in Chinese are  
925 notably lacking. Creating such datasets presents  
926 unique challenges, particularly in generating real-

这里有一段陈述和对应的一段参考文本。请按如下步骤完成任务，严格按我给出的格式进行输出：

(1) 找到参考文本中所有直接支撑陈述中信息的原始关键文段（可能有多个，每一处都要找到）。每行输出一个原始关键文段及其直接支撑的陈述中的信息，格式为“关键文段编号：关键文段（支撑陈述中的信息：支撑信息）”。

(2) 请将关键文段分组，每组包含的关键文段支撑陈述中的相同或相关的信息，输出一行分组结果，格式为“关键文段分组：第一组：（第一组关键文段编号），第二组：（第二组关键文段编号）...”。例如，陈述中有2个信息，关键文段1支撑信息1，关键文段2支撑信息2，关键文段3支撑信息1，那么输出“关键文段分组：第一组：（1, 3），第二组：（2）”

(3) 选择一组关键文段，对其中支撑陈述中信息的部分进行修改，满足以下要求：

- 修改应该使得关键文段无法完全支撑陈述中的对应信息。

- 修改应该保持关键文段的逻辑通顺、关键文段中的信息之间不矛盾。

- 修改之后的关键文段应该在参考文本的上下文语境中保持逻辑通顺，且与参考文本中的其他内容不矛盾。

- 只修改支撑陈述中信息的部分，其它部分保持不变。

- 如果一组中有多个关键文段，修改后它们的信息应该保持一致。

你需要尝试两种修改方法：

- 改变信息：将关键文段中的某一处信息修改为另外的信息。请不要进行与原信息产生直接冲突的修改。例如，原信息为“奥迪A7旗舰版的最高速度比上一代快”，合适的修改是“奥迪A7豪华版的最高速度比上一代快”，不合适的修改1是“奥迪A7旗舰版的最高速度比上一代慢”（使用反义词，与原信息直接冲突），不合适的修改2是“奥迪A7旗舰版的最高速度不比上一代快”（添加否定词，与原信息直接冲突）。

- 删除信息：将关键文段中的某一处信息删除。关键文段如果是完整的句子，删除信息后应该仍然是一个完整的句子。例如，原文段为“由于天气原因，项目推迟至3月15日启动”（完整的句子），合适的修改是“由于天气原因，项目推迟至3月启动”（仍然是完整的句子），不合适的修改是“由于天气原因”（不再是完整的句子）。

对每种方法，输出被修改的关键文段，并检查其逻辑通顺程度，给出一个1~10以内的整数作为评分（越高表示越通顺）。每行输出一个修改后的关键文段，格式为“方法-修改后的关键文段编号：修改后的关键文段（逻辑通顺程度：分数）”。

Table 3: The complete prompt for the LLM augmentation.

Here is a statement and a corresponding piece of reference text. Please complete the task as follows, strictly following the format I have given for the output:

(1) Find all the original key passages in the reference text that directly support the information in the statement (there may be more than one, find each one). Output one original key passage per line and the information in the statement it directly supports in the format “Key passage number: key passage (information in the supporting statement: supporting information)”.

(2) Please group key passages, each group contains key passages supporting the same or related information in the statement, output one line of the grouping results in the format of “Key passage grouping: Group 1: (first group of key passage numbers), Group 2: (second group of key passage numbers) ...”. For example, if there are 2 pieces of information in the statement, key paragraph 1 supports information 1, key paragraph 2 supports information 2, and key paragraph 3 supports information 1, then the output is “Key Paragraph Grouping: Group 1: (1, 3), Group 2: (2)”.

(3) Select a group of key text segments and modify the parts of them that support the information in the statement to meet the following requirements:

- The modification should make it impossible for the key passage to fully support the corresponding information in the statement.
- The modifications should maintain the logical flow of the key passages and no contradictions between the information in the key passages.
- The modification should keep the key paragraph logically coherent in the context of the reference text and not contradict the rest of the reference text.
- Modify only the parts that support the information in a statement, leaving the rest unchanged.
- If there is more than one key passage in a set, the information in them should remain consistent after revision.

You need to try two methods of modification:

- Changing the message: modifying the message in one part of the key paragraph to another. Do not make changes that directly conflict with the original information. For example, if the original message is “The Audi A7 Signature Edition has a faster top speed than its predecessor”, an appropriate change would be “The Audi A7 Luxury Edition has a faster top speed than its predecessor”, and an inappropriate change 1 would be “The Audi A7 Signature Edition has a slower top speed than its predecessor” (using an antonym, which is in direct conflict with the original message), and inappropriate modification 2 is ‘The top speed of the Audi A7 Signature Edition is not faster than the previous generation’ (adding a negative word, which is in direct conflict with the original message).
- Delete Information: Remove information from a place in a key paragraph. If the key paragraph is a complete sentence, it should still be a complete sentence after deleting the information. For example, if the original paragraph reads “Due to weather conditions, the project was delayed until March 15” (complete sentence), an appropriate change would be “Due to weather conditions, the project was delayed until March” (still a complete sentence), an inappropriate change would be “Due to the weather” (no longer a complete sentence).

For each method, output the key passage that was modified and check its logical fluency, giving an integer within 1 to 10 as a rating (higher means more fluent). Output one modified key passage per line in the format “method-modified key passage number: modified key passage (logical fluency: score)”.

Table 4: The complete prompt for the LLM augmentation (translated into English).

927           istic negative samples (unsupported citations). To  
928           address this gap, we introduce the first large-scale  
929           Chinese dataset for citation faithfulness detection,  
930           developed through a cost-effective two-stage man-  
931           ual annotation process.

<p><b>Question:</b>为什么太空旅行很重要？</p> <p><b>Answer:</b>太空旅行对于我们生活和未来有着重要的意义。首先，太空旅行可以带来新的科学和技术的发展。例如，NASA的“磁场塔”计划可以在人造磁场下模拟地球引力，为未来的太空生存提供可行的技术路径。其次，太空旅行可以使更多的人有机会涉足太空，并享受到其中的奇妙体验。此外，太空旅行还可以促进技术的创新和进步。最后，太空旅行还可以带来新的变革，不仅带来人类探索的新篇章，而且也将为我们生活的方方面面带来新的变革。</p> <p><b>Statement:</b>例如，NASA的“磁场塔”计划可以在人造磁场下模拟地球引力，为未来的太空生存提供可行的技术路径。</p> <p><b>Reference:</b></p> <p>1: 比如，NASA的“磁场塔”计划就是一项非常有前途的研究，它可以在人造磁场下模拟地球引力，从而为未来的太空生存提供可行的技术路径。最后，太空旅游也将面临一些挑战和风险。太空环境对人体的影响、太空辐射的危害、太空垃圾的处理等问题都需要得到解决。此外，太空旅游的成功也需要得到国际社会的支持和认可，有建立规范和标准的必要。总的来说，太空旅游是一项充满挑战和机遇的领域。未来，随着人类科技和探索的不断进展，太空旅游将会越来越先进、更加普及化。这将不仅带来人类探索的新篇章，而且也将为我们生活的方方面面带来新的变革。</p>	<p><b>Method:</b>改变信息</p> <p><b>Statement:</b>例如，NASA的“磁场塔”计划可以在人造磁场下模拟地球引力，从而为未来的太空生存提供可行的技术路径。</p> <p><b>Modified Reference:</b></p> <p>1: 比如，NASA的“磁场塔”计划就是一项非常有前途的研究，它可以在人造磁场下模拟地球引力，从而为未来的太空生存提供可行的技术路径。最后，太空旅游也将面临一些挑战和风险。太空环境对人体的影响、太空辐射的危害、太空垃圾的处理等问题都需要得到解决。此外，太空旅游的成功也需要得到国际社会的支持和认可，有建立规范和标准的必要。总的来说，太空旅游是一项充满挑战和机遇的领域。未来，随着人类科技和探索的不断进展，太空旅游将会越来越先进、更加普及化。这将不仅带来人类探索的新篇章，而且也将为我们生活的方方面面带来新的变革。</p>	<p><b>Method:</b>删除信息</p> <p><b>Statement:</b>例如，NASA的“磁场塔”计划可以在人造磁场下模拟地球引力，为未来的太空生存提供可行的技术路径。</p> <p><b>Modified Reference:</b></p> <p>1: 比如，NASA的“磁场塔”计划就是一项非常有前途的研究，它可以在人造磁场下模拟地球引力，从而为未来的太空生存提供可行的技术路径。最后，太空旅游也将面临一些挑战和风险。太空环境对人体的影响、太空辐射的危害、太空垃圾的处理等问题都需要得到解决。此外，太空旅游的成功也需要得到国际社会的支持和认可，有建立规范和标准的必要。总的来说，太空旅游是一项充满挑战和机遇的领域。未来，随着人类科技和探索的不断进展，太空旅游将会越来越先进、更加普及化。这将不仅带来人类探索的新篇章，而且也将为我们生活的方方面面带来新的变革。</p>
---	--	--

Figure 1: Examples of interfaces that provide samples to the annotators. The first figure shows an example of the first stage. The last two images show the second stage with the same sample modified (information changed/deleted).