

The Role of Large Language Models in Musicology: Are We Ready to Trust the Machines?

Pedro Ramoneda^{1*} Emilia Parada-Cabaleiro² Benno Weck¹ Xavier Serra¹

¹Music Technology Group, Universitat Pompeu Fabra, Barcelona, Spain

²Department of Music Pedagogy, Nuremberg University of Music, Germany

Abstract

In this work, we explore the use and reliability of Large Language Models (LLMs) in musicology. From a discussion with experts and students, we assess the current acceptance and concerns regarding this, nowadays ubiquitous, technology. We aim to go one step further, proposing a semi-automatic method to create an initial benchmark using retrieval-augmented generation models and multiple-choice question generation, validated by human experts. Our evaluation on 400 human-validated questions shows that current vanilla LLMs are less reliable than retrieval-augmented generation from music dictionaries. This paper suggests that the potential of LLMs in musicology requires musicology driven research that can specialize LLMs by including accurate and reliable domain knowledge.

1 Introduction

In recent years, research on Large Language Models (LLMs) has led to notable advancements within the text generation domain (Wei et al., 2022a; Minaee et al., 2024). This is the result of training large models on vast non-domain-specific data (Gao et al., 2020; Hoffmann et al., 2022). Well-known families of models include Llama (AI@Meta, 2024) or GPT (Achiam et al., 2023), which can generate coherent and contextually relevant text, making them valuable tools in numerous applications and professions such as healthcare (Thirunavukarasu et al., 2023), journalism (Petridis et al., 2023), customer support (Kolasani, 2023) or education (Kasneci et al., 2023).

Despite their potential, LLMs' so-called hallucinations (Alkaissi and McFarlane, 2023), i. e., the lack of confidence and accuracy in the text they generate, prevents the use of this technology in most arts and humanities research tasks (Rane,

Musicologist: What's the historical context of this music piece?

LLM: It's by Beethoven in 2025! Aliens helped!

Musicologist: I'm not using THIS anymore.

Figure 1: Fictitious interaction illustrating why LLMs' hallucinations might prevent musicologists' trust.

2023; Lozić and Štular, 2023; Rane and Choudhary, 2024). Issues include a lack of contextual understanding, bias perpetuation (Gallegos et al., 2024), and ethical concerns such as generating misleading content (Weidinger et al., 2021). The lack of credible source attribution (Rashkin et al., 2023) almost render them nugatory for fields like literature, history (Walters and Wilder, 2023), and law (Weiser, 2024). However, LLMs can aid research through a variety of tasks, such as, translation, text analysis, data organization, historical context retrieval, or summarization. In this regard, interdisciplinary research involving the use and further development of LLMs within the humanities should be carried out. This will enable to constructively address existing risks and concerns while developing LLMs' full potential, by this delivering their benefits across disciplines.

In this work, we focus on musicology, a field where the impact of LLMs still needs to be explored. Musicology, the scholarly study of music, spans from historical research to theoretical analysis (Harap, 1937; Duckles et al., 2020). Our research mainly focuses on the former, an area which might be greatly supported by LLMs, e. g., by breaking language barriers, enhancing information retrieval, or supporting teaching and learning. However, reliable sources, such as music-specialized lexica, monographies, and research articles, are often, unlike in more technical disciplines, not open-access, which prevents LLMs to

* Corresponding author: pedro.ramoneda@upf.edu

access high quality information. This knowledge deprivation further increases the risk of LLMs to hallucinate, which often leads to non-reliable text generation in musicology related topics.

Through a pilot-survey involving experts and students from the field of musicology, we gather initial insights into the acceptance and trustworthiness of LLMs in domain-related tasks, and its potential impact for music professionals. Subsequently, we propose a methodology to measure to which extent such models possess domain expertise in the field of musicology, by this assessing their practical value for the discipline. We adopt a Multiple-Choice Question Generation (Liu et al., 2024) approach to semi-automatically construct a benchmark leveraging recent advancements in retrieval-augmented generation models (Lewis et al., 2020). To automatically generate high-quality questions, we provide the generation model with domain-knowledge from *The New Grove Dictionary of Music and Musicians* (Sadie and Tyrrell, 2001), an established and reliable source. The final benchmark, made up by 400 question-answer pairs validated by a human expert, is evaluated on several open-source models. This dual approach—survey and benchmark—provides a comprehensive understanding of the challenges and potential solutions for meaningful integration of LLMs in musicology.

2 Pilot-survey: LLMs in musicology

We conducted a survey targeting professionals related to musicology. The survey included questions to identify the respondent’s domain of study (e.g., musicology, composition, music pedagogy, music performance), the highest level of music education completed or being pursued, and their familiarity with technologies known as LLMs such as ChatGPT. Additionally, the survey inquired about the frequency of interactions with LLMs, particularly in the context of musical topics like Music Theory and Music History. Participants were asked to rate the trustworthiness and usefulness of LLMs for these subjects, as well as to consider its revolutionary impact on the field of musicology. Lastly, the survey explored the possible consequences of LLMs on music professionals, both presently and in the future.

A total of 33 participants, having or pursuing a Bachelor’s degree in music, completed the survey: 20 students, 7 lecturers, 11 researchers, and 8 mu-

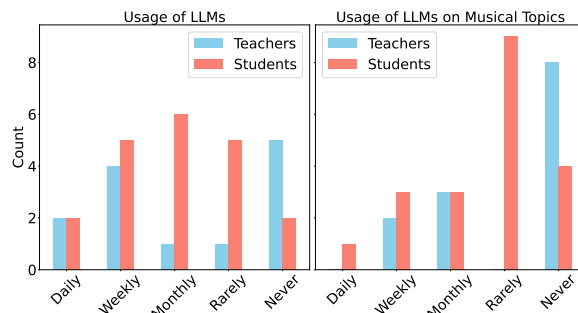


Figure 2: Survey’s answers about the usage of LLMs in general (left) and on music topics (right)

sic educators (multiple areas can be selected). In terms of discipline, the respondents are distributed as the following: 22 Musicology and Related Studies, 10 Music Performance, 3 Music Pedagogy, 2 Composition, 1 Conduction, and 1 Music Therapy. While only one participant (from the field of musicology) had not heard about LLMs before, in terms of the participants’ frequency of use and trustworthiness, a noticeable gap between students and teachers¹ can be observed.

Figure 2 illustrates how often students and teachers interact with LLMs in general and about music topics. Teachers frequently or not at all, while most students use it weekly or monthly. However, both groups tend to not use LLMs for music-related topics. Participants’ judgement of LLMs’ trustworthiness is depicted in Figure 3 and the trend of ratings is similar across both groups. Additionally, confidence in LLMs is slightly higher for Music History than for Music Theory, indicating a nuanced perception of their reliability in different musicology subfields.

Most of the participants (78%) agreed that LLMs might revolutionize the field of musicology (cf. Figure 4, left). While the anticipated potential consequences of LLMs for the field are varied, professional transformation seems to be the most prominent (20 votes), as illustrated in the histogram (cf. Figure 4, right). In conclusion, despite limited current usage and trust, experts anticipate a significant future impact of LLMs on musicology, motivating current research on the topic.

3 Musicology Benchmark: TrustMus

This section outlines our strategy for evaluating how much LLMs hallucinate in musicology. It summarizes the creation of the human-validated

¹For simplicity, with ‘teachers’ we refer to all the participants who did not identify themselves as student.

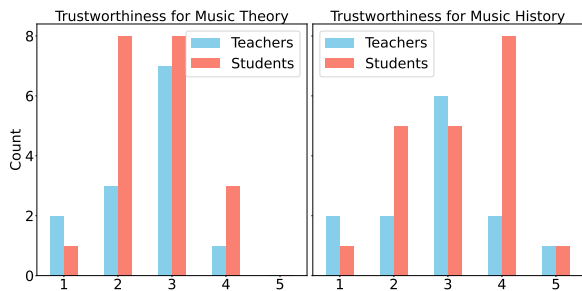


Figure 3: Survey’s answers about the usage of LLMs in general (left) and on music topics (right).

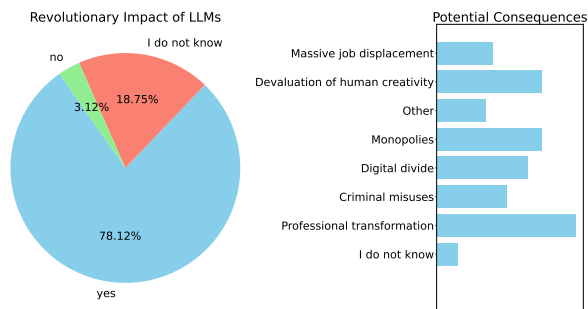


Figure 4: Survey answers about the revolutionary impact (right) and potential consequences (left) of LLMs.

multiple-choice benchmark *TrustMus*, i. e., a collection of reliable questions related to various musical topics and concepts, and analyzes the models’ performance on the benchmark.

Following previous works (Li et al., 2023), multiple-choice questions are generated after extracting relevant information from a text source: here, *The Grove Dictionary Online* (Sadie and Tyrrell, 2001).² In order to identify the most relevant articles within the text source, we used a PageRank-like algorithm (Hagberg et al., 2008).

To accelerate the creation of *TrustMus*, we designed a workflow inspired by recent works (Yan et al., 2024; Jeong et al., 2024; Asai et al., 2023; Dhuliawala et al., 2023), as shown in Figure 5. First, we generated five questions from each article, each with four possible answer options, using a fine-tuned LLM for retrieval-augmented generation (RAG) (Liu et al., 2024), resulting in 7 500 questions. Second, we discard questions that did not have relation with musicology or a unique and unambiguous answer, by prompting the same LLM to decide based on the article, eliminating 2 632 questions. Next, we attempted to answer the remaining questions using a RAG-like model that we term *Llama Professor* by giving the arti-

²The Grove Dictionary is a copyrighted work. Using its content for generating questions is under fair use for research purposes. The EU Directive on Copyright in the Digital Single Market allows text and data mining for research purposes.

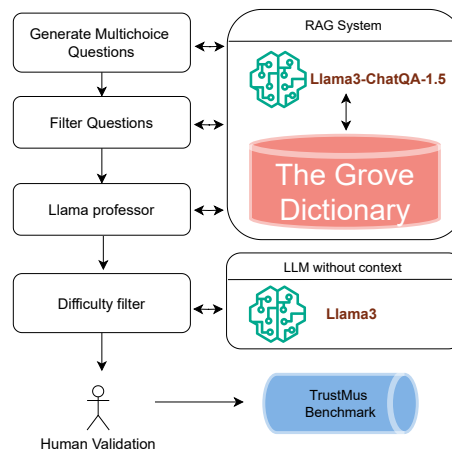


Figure 5: Language chain for generating the multiple-choice questions.

cle as context to the LLM. Questions for which Llama Professor chooses the wrong answer option are considered ambiguous or unusable and are thus removed, resulting in 3 285 valid questions. All previous prompts used the Chain of Thought (CoT) method to enhance the model’s reasoning skills (Wei et al., 2022b). Before human intervention, we attempted to answer the questions with llama3-8B (AI@Meta, 2024) without RAG and in one shot, i. e., without the chain of thought (cf. the difficulty filter in Figure 5), which lead an accuracy of 67.4%. Thus, arguably simple questions are eliminated, resulting in 1 081 domain-ones.

The resulting set of questions was automatically classified with a CoT prompt into four classes, according to their topic: People (Ppl); Instruments and Technology (I&T); Genres, Forms, and Theory (Thr); Culture and history (C&H). An expert human annotator validated questions until 100 valid ones per class were identified (on average, 17% of those assessed were discarded).³

4 Results and Discussion

4.1 Human validation insights

Some examples of hallucinations of Llama3 without RAG and CoT, the difficulty filter, are as follows: *What does the natural sign (\natural) do in music notation? A) Raises a note by one semitone, B) Raises a note by two semitones, C) Lowers a note by one semitone, D) Cancels a previous sharp or flat.* The correct answer is D, but Llama3 chose A, which any musician should know is incorrect.

Another type of limitation of LLMs in the context of musicology, is the need of the models for

³@: <https://zenodo.org/records/13644330>

Model	Quant	TrustMus	Rank	Ppl	I&T	Thr	C&H	LB	Rank
gpt-4o-2024-05-13 (Achiam et al., 2023)	API	58.75	1	60.0	44.0	61.0	70.0	58.38	1
mixtral-8x7b-instruct-v0.1 (Jiang et al., 2024)	✓	40.5	2	41.0	30.0	43.0	48.0	37.24	4
gpt-3.5-turbo-0125 (Achiam et al., 2023)	API	39.75	3	39.0	25.0	43.0	52.0	37.97	5
meta-llama-3-70b-instruct (AI@Meta, 2024)	✓	37.75	4	41.0	23.0	44.0	43.0	42.76	2
qwen2-72b-instruct (Bai et al., 2023)	✓	35.5	5	39.0	27.0	37.0	39.0	41.43	3
qwen2-7b-instruct (Bai et al., 2023)	✗	34.0	6	29.0	43.0	36.0	41.0	25.92	9
phi-3-medium-4k-instruct (Abdin et al., 2024)	✓	32.75	7	32.0	27.0	38.0	34.0	33.46	6
meta-llama-3-8b-instruct (AI@Meta, 2024)	✗	32.75	8	43.0	22.0	31.0	35.0	31.05	7
phi-3-small-128k-instruct (Abdin et al., 2024)	✗	31.5	9	20.0	29.0	41.0	36.0	28.12	8
Llama Professor (RAG)	✗	100.0	-	100.0	100.0	100.0	100.0	-	-

Table 1: Benchmark results (accuracy) on the 400 validated questions (TrustMus) and per category: People (Ppl); Instruments and Technology (I&T); Genres, Forms, and Theory (Thr); Culture and History (C&H). Whether the models are quantized (Quant), their rank, and LiveBench average score (LB) excluding math ranking is also given.

interpreting the information. This can be illustrated by how Llama3 handled the article about *Adagio* in the *Grove Dictionary Online*, which summarizes the evolution of the term over centuries. In this regard, when interrogating Llama3 about the term as described by Rousseau, the model refers to the modern definition.

4.2 TrustMus evaluation

Table 1 presents the benchmark results for various models evaluated on TrustMus.⁴ The models tested include the best open source performing models in LiveBench – LB (White et al., 2024) excluding coding and math categories, i. e., a benchmark for LLMs without contamination and reduced biases containing non-musicology knowledge. Due to its’ leading performance, results of OpenAI’s GPT models are also given for comparison. Models with less than 8B parameters were deployed in a computer with two RTX 2080ti GPUs with 16-bit precision, the largest models in a Colab A100 GPU with 4-bit quantization, and the GPT models through their official API.

The model `gpt-4o-2024-05-13` clearly outperforms others with an accuracy of 58.75% (cf. TrustMus score in Table 1), excelling in the categories Ppl, Thr, and C&H. This is not surprising as it is the leading model in LB as well, with a score of 58.38%. However, comparing the LB and TrustMus rankings reveals important differences about how the models perform in terms of general and in domain-specific knowledge. For instance, unlike in LB, the model `mixtral-8x22b-instruct-v0.1` performs well in our benchmark, ranking second with a score of 40.5%. It is important to note the simi-

lar performance between `qwen` with 72B and 7B, the latter being the best performing of the ‘small’ LLMs in TrustMus while showing the worst performance in LB. We also aim to acknowledge that comparing `meta-llama3` models with others is not entirely fair, as the benchmark was automatically generated by selecting questions from their specific blind spots, as detailed in Section 3.

The lower performance of open-source models compared to LLama professor (RAG) (Liu et al., 2024) highlights the importance of reliable domain-specific knowledge for musicology-related applications. This indicates considerable improvement possibilities with the potential of increasing the trustworthiness of LLMs in the field.

5 Conclusions

Our paper shows that while current usage and trust in LLMs in musicology are low, there is a strong expectation of future impact. However, LLMs are not yet at the required level for the field and do not meet the minimum quality, ethical and likely legal standards currently being discussed. Through the proposed semi-automatic benchmark, we present a first attempt to measure LLMs hallucinations on musicology-related tasks. This approach aims to facilitate the evaluation of future models, which promotes transparency and trustworthiness of the technology. Despite the effort, this initial experiments are insufficient. Besides a more thorough evaluation, there is the need to specialize current models for musicology-related tasks, while reducing their environmental footprint. Further research should focus on ensuring LLMs reliability to avoid misinformation, protecting user privacy and data security, and mitigating training data biases to promote responsible use in musicology. Collaboration between the technological, musicological, and content owner communities is essential for the proper development of this technology.

⁴Since we believe that open models are critical for transparency, reproducibility, and the advancement of knowledge, we use them in our research. We included ChatGPT in our comparison only because it is currently the most used LLM.

References

- Marah Abdin, Sam Ade Jacobs, Ammar Ahmad Awan, Jyoti Aneja, Ahmed Awadallah, Hany Awadalla, et al. 2024. Phi-3 technical report: A highly capable language model locally on your phone. *arXiv preprint arXiv:2404.14219*.
- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, et al. 2023. GPT-4 technical report. *arXiv preprint arXiv:2303.08774*.
- AI@Meta. 2024. [Llama 3 model card. Final report not published yet](#).
- Hussam Alkaiissi and Samy I McFarlane. 2023. Artificial hallucinations in ChatGPT: implications in scientific writing. *Cureus*, 15(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 preprint arXiv:2310.11511*.
- Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. 2023. Qwen technical report. *arXiv preprint arXiv:2309.16609*.
- Shehzaad Dhuliawala, Mojtaba Komeili, Jing Xu, Roberta Raileanu, Xian Li, Asli Celikyilmaz, and Jason Weston. 2023. Chain-of-verification reduces hallucination in large language models. *arXiv preprint arXiv:2309.11495*.
- Vincent Duckles, Jann Pasler, Glenn Stanley, Thomas Christensen, Barbara H. Haggh, Robert Balchin, et al. 2020. [Musicology](#). In *The Grove Music Online*. Oxford University Press.
- Isabel O Gallegos, Ryan A Rossi, Joe Barrow, Md Mehrab Tanjim, Sunghul Kim, Franck Dernoncourt, Tong Yu, Ruiyi Zhang, and Nesreen K Ahmed. 2024. Bias and fairness in large language models: A survey. *Computational Linguistics*.
- Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, et al. 2020. The pile: An 800gb dataset of diverse text for language modeling. *arXiv preprint arXiv:2101.00027*.
- Aric A. Hagberg, Daniel A. Schult, and Pieter J. Swart. 2008. Exploring network structure, dynamics, and function using networkx. In *Proceedings of the 7th Python in Science Conference (SciPy2008)*.
- Louis Harap. 1937. On the nature of musicology. *The Musical Quarterly*, 23(1).
- Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, et al. 2022. Training compute-optimal large language models. *arXiv preprint arXiv:2203.15556*.
- Soyeong Jeong, Jinheon Baek, Sukmin Cho, Sung Ju Hwang, and Jong C Park. 2024. Adaptive-RAG: Learning to adapt retrieval-augmented large language models through question complexity. *arXiv preprint arXiv:2403.14403*.
- Albert Q Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, et al. 2024. Mixtral of experts. *arXiv preprint arXiv:2401.04088*.
- Enkelejd Kasneci, Kathrin Seßler, Stefan Küchemann, Maria Bannert, Daryna Dementieva, Frank Fischer, et al. 2023. ChatGPT for good? on opportunities and challenges of large language models for education. *Learning and individual differences*, 103:102274.
- Saydulu Kolasani. 2023. Optimizing natural language processing, large language models (llms) for efficient customer service, and hyper-personalization to enable sustainable growth and revenue. *Transactions on Latest Trends in Artificial Intelligence*, 4(4).
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, et al. 2020. Retrieval-augmented generation for knowledge-intensive NLP tasks. In *Proceedings of Advances in Neural Information Processing Systems*, 33 (*NeurIPS 2020*).
- Junyi Li, Xiaoxue Cheng, Xin Zhao, Jian-Yun Nie, and Ji-Rong Wen. 2023. HaluEval: A large-scale hallucination evaluation benchmark for large language models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Zihan Liu, Wei Ping, Rajarshi Roy, Peng Xu, Mohammad Shoeybi, and Bryan Catanzaro. 2024. ChatQA: Building GPT-4 Level Conversational QA Models. *arXiv preprint arXiv:2401.10225*.
- Edisa Lozić and Benjamin Štular. 2023. Fluent but not factual: A comparative analysis of ChatGPT and other ai chatbots' proficiency and originality in scientific writing for humanities. *Future Internet*, 15(10).
- Shervin Minaee, Tomas Mikolov, Narjes Nikzad, Meysam Chenaghlu, Richard Socher, Xavier Amatriain, and Jianfeng Gao. 2024. Large language models: A survey. *arXiv preprint arXiv:2402.06196*.
- Savvas Petridis, Nicholas Diakopoulos, Kevin Crowston, Mark Hansen, Keren Henderson, Stan Jastrzebski, et al. 2023. Anglekindling: Supporting journalistic angle ideation with large language models. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI 23)*.

- Nitin Rane. 2023. Role and challenges of ChatGPT and similar generative artificial intelligence in arts and humanities. *Available at SSRN 4603208*.
- Nitin Rane and Saurabh Choudhary. 2024. Role and challenges of ChatGPT, Google Bard, and similar generative Artificial Intelligence in Arts and Humanities. *Studies in Humanities and Education*, 5(1).
- Hannah Rashkin, Vitaly Nikolaev, Matthew Lamm, Lora Aroyo, Michael Collins, Dipanjan Das, Slav Petrov, Gaurav Singh Tomar, Iulia Turc, and David Reitter. 2023. Measuring attribution in natural language generation models. *Computational Linguistics*, 49(4).
- Stanley Sadie and John Tyrrell, editors. 2001. *The New Grove Dictionary of Music and Musicians*, 2nd edition. Macmillan Publishers, London. Grove Music Online. Edited by Deane Root. Accessed 05-05-2024. <http://www.oxfordmusiconline.com>.
- Arun James Thirunavukarasu, Darren Shu Jeng Ting, Kabilan Elangovan, Laura Gutierrez, Ting Fang Tan, and Daniel Shu Wei Ting. 2023. Large language models in medicine. *Nature Medicine*, 29(8).
- William H Walters and Esther Isabelle Wilder. 2023. Fabrication and errors in the bibliographic citations generated by chatgpt. *Scientific Reports*, 13(1):14045.
- Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, et al. 2022a. Emergent abilities of large language models. *Transactions Machine Learning Research*, 2022.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, et al. 2022b. Chain-of-Thought Prompting Elicits Reasoning in Large Language Models. In *Advances in Neural Information Processing Systems 35 (NeurIPS 2022)*.
- Laura Weidinger, John Mellor, Maribeth Rauh, Conor Griffin, Jonathan Uesato, Po-Sen Huang, et al. 2021. Ethical and social risks of harm from language models. *arXiv preprint arXiv:2112.04359*.
- Benjamin Weiser. 2024. [Here's what happens when your lawyer uses ChatGPT](#). *New York Times*. Accessed 05-05-2024.
- Colin White, Samuel Dooley, Manley Roberts, Arka Pal, Ben Feuer, Siddhartha Jain, et al. 2024. LiveBench: A Challenging, Contamination-Free LLM Benchmark. *arXiv preprint arXiv:2406.19314*.
- Shi-Qi Yan, Jia-Chen Gu, Yun Zhu, and Zhen-Hua Ling. 2024. Corrective retrieval augmented generation. *arXiv preprint arXiv:2401.15884*.