

Question Suggestion for Conversational Shopping Assistants Using Product Metadata

Nikhita Vedula
Amazon, Seattle
veduln@amazon.com

Oleg Rokhlenko
Amazon, Seattle
olegro@amazon.com

Shervin Malmasi
Amazon, Seattle
malmasi@amazon.com

ABSTRACT

Digital assistants have become ubiquitous in e-commerce applications, following the recent advancements in Information Retrieval (IR), Natural Language Processing (NLP) and Generative Artificial Intelligence (AI). However, customers are often unsure or unaware of how to effectively converse with these assistants to meet their shopping needs. In this work, we emphasize the importance of providing customers a fast, easy to use, and natural way to interact with conversational shopping assistants. We propose a framework that employs Large Language Models (LLMs) to automatically generate contextual, useful, answerable, fluent and diverse questions about products, via in-context learning and supervised fine-tuning. Recommending these questions to customers as helpful suggestions or hints to both start and continue a conversation can result in a smoother and faster shopping experience with reduced conversation overhead and friction. We perform extensive offline evaluations, and discuss in detail about potential customer impact, and the type, length and latency of our generated product questions if incorporated into a real-world shopping assistant.

CCS CONCEPTS

• **Computing methodologies** → **Natural language generation**; **Natural language processing**; *Supervised learning*.

KEYWORDS

Conversational Shopping Assistants; Product Question Suggestion

ACM Reference Format:

Nikhita Vedula, Oleg Rokhlenko, and Shervin Malmasi. 2024. Question Suggestion for Conversational Shopping Assistants Using Product Metadata. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '24)*, July 14–18, 2024, Washington, DC, USA. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3626772.3661371>

1 INTRODUCTION AND BACKGROUND

Driven by the recent progress in IR, NLP and generative AI, there has been a surge in the deployment of intelligent virtual assistants and chat bots employing both voice and text, in major e-commerce services like Amazon, Walmart and Shopify. Shopping assistants

can recommend products, summarize reviews and catalog metadata, compare products, and answer questions about product features, orders and deliveries. The adoption of such conversational assistants by customers to meet their diverse shopping needs hinges on accessible, convenient and customer-friendly application interfaces [2]. To create sustainable and satisfying interactions between customers and shopping assistants, we must help customers adapt to them, by expediting learning curves and minimizing the need for customer trial and experimentation.

One way of achieving this is to provide customers with a useful set of *question suggestions* or *question hints*, that can guide customers in initiating and advancing their shopping conversations with assistants. These questions can also inspire them to ask the assistant their own questions based on their needs. The automatic Question Generation (QG) problem is well studied across various domains [7, 11, 19]. However, a shopping assistant suggesting irrelevant, unanswerable or redundant questions can severely impair customer trust and the overall buying experience. In this work, we introduce shoppers to the product-based question answering (QA) functionality of shopping assistants. We propose an LLM-based approach leveraging in-context learning (ICL) [14] and supervised fine-tuning (SFT), to *automatically generate question suggestions* grounded in catalog-derived product metadata and buyer reviews.

Further, we identify and propose several shopping-specific criteria that should be satisfied by our generated product questions, to make interactions with shopping assistants more natural, intuitive and optimal for customers (defined in Section 3.1). We associate and pair each generated question with its corresponding product context information containing the answer to the question, ensuring that our generated questions are *answerable*. Since this context is derived from the catalog or reviews, it is human authored and has the added benefit of avoiding automatic answer generation [8, 15, 19], which suffers from a risk of hallucination [13, 18]. Questions paired with their relevant answers or context can also be useful in downstream applications such as creating an automatic bank of frequently asked questions (FAQs) [3, 10], building a retrieval index for a retrieval augmented generation (RAG) LLM [5], or fine-tuning a model for e-commerce based QA [4].

Finally, we extensively evaluate our approach, and discuss its potential impact on shopping assistants in Sections 4 and 5. We expect to obtain a consistent alignment between the offline and online results and performance trends, as well as several valuable insights about real customer interactions with shopping assistants. Providing automatic question suggestions to customers that are a good mix of both broad and specific questions can effectively streamline the conversation, which can in turn lead to low customer dissatisfaction rate. The latency of real-time question generation by LLMs can be reduced by mechanisms such as caching specific outputs or

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGIR '24, July 14–18, 2024, Washington, DC, USA.

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-0431-4/24/07
<https://doi.org/10.1145/3626772.3661371>

delivering the generated tokens in a streaming fashion [16]. While shorter questions are easier to display (especially on mobile device interfaces), longer questions can serve as useful summaries for multiple key product aspects at a time, saving customers from clicking or typing out a series of questions.

Overall, our proposed approach reduces the number of steps required for a customer to reach their shopping goal, contributing to an improved overall shopping experience, fostering customer satisfaction and engagement. We show examples of questions generated by our approach and their paired answers in Table 2.

2 PROBLEM STATEMENT AND DATASET

We formally describe our problem of generating *conversational product-based question suggestions*. Every product in an e-commerce catalog is associated with multiple sources of textual content about it, coming from both the product sellers or manufacturers (product descriptions and metadata) as well as buyers of the product (reviews). Our objective is to leverage this rich content to create a model capable of generating one or more product questions grounded in the input product context, that potential customers may ask shopping assistants about the product.

We use the publicly available Amazon Reviews Dataset [12] as our information source. It contains at least 200 million buyer reviews for tens of thousands of products spanning more than 20 product categories, as well as product catalog information including the title, category, descriptions, price, brand and other product-specific attributes. Since not all customer reviews are informative or of high quality, we only use reviews that have received a high number of *helpfulness* votes from other customers, as well as reviews written by reviewers who are a part of Amazon’s Vine¹ program.

3 OUR APPROACH

Our proposed LLM-based question generation approach takes textual content about a product as input (derived from a product review or a catalog description from Section 2), and generates one or more question suggestions for customers from that content.

3.1 Product Question Quality Criteria

We first identify and propose a set of crucial, multi-faceted quality criteria required for an effective and efficient conversation between a customer and shopping assistant. Our approach seeks to satisfy these for each generated product question.

(i) **Relevance:** The question should be applicable and appropriate with respect to the product under consideration and its features.

(ii) **Usefulness:** The product question (and its corresponding answer) should provide helpful information to customers, that can benefit them in deciding whether or not to purchase the product.

(iii) **Answerability:** The answer to the generated product question must be present in its input context (review or catalog snippet).

(iv) **Fluency:** The generated question should be grammatically correct, fluent, coherent and easily understandable in general.

(v) **Style:** The generated question should mimic a customer’s inquiry style. E.g., the question “Which colors do you prefer for this jacket?” is about a valid product aspect (*color*), but is framed as a clarification question that an assistant might ask a customer, rather

¹<https://www.amazon.com/vine/about>

Human: You are an intelligent shopping assistant helping customers shop for products. Use the **Product Info** given below to output the top product question, question type and its estimated customer interest score (1-10). A customer should be able to ask a salesperson this question.

Your question should be about broad features of the product, specific product aspects, compatibility with other products, comparisons with other products, or other types of important buying guide questions that can fulfill diverse customer needs.

Your question must adhere to the following criteria:

1. Your question must be relevant to the given Product Info
2. Your question should be useful in helping the customer decide whether to purchase the given product.
3. The answer to your question must be contained within the Product Info.
4. Your question should be grammatically correct, fluent and coherent.
5. Your question can either be short and concise, or long and verbose based on your judgment of the Product Info.
6. Avoid using personal, first person or second person pronouns in your output question.
7. Use anaphora like ‘it’ or ‘this’ to refer to the product when applicable.

Use the Product Info given below to output the top product question, question type and its estimated customer interest score (1-10). Output structured pipe-separated columnar data without any other text.

Product Info: {data}.

Assistant:

Figure 1: Prompt for ICL-based product question generation.

than a valid question that a customer may ask a shopping assistant. A better formulated question suggestion for a customer would be “Which colors are available for this jacket?”

(vi) **Diversity:** The set of questions that are generated per product should cover a wide variety of product features or aspects without being redundant or repetitive. The questions should be of varying lengths (both short and concise as well as long and verbose) and complexities (in terms of their shape and sentence structures), to cater to a multitude of customer requirements. Generating and recommending questions of several different types also aids customers at various stages of their shopping journeys. For instance, new customers may prefer broad or exploratory product questions, whereas customers midway during the shopping process may inquire about comparisons or specific product attributes.

We next describe the two LLM-based generation techniques that we leverage as part of our question generation approach.

3.2 Product Question Suggestion Generation

In-Context Learning: We harness the rich world knowledge possessed by LLMs, coupled with the question quality criteria described in Section 3.1 to construct a prompt (Figure 1) that can generate diverse product question suggestions for customers. In Section 4, we evaluate the performance of this prompt both in a zero-shot manner, and by enriching it with examples that the LLM can learn from in a few-shot manner.

Supervised Fine-tuning: We construct a high quality training dataset in a time-efficient manner with reduced human annotation effort, to fine-tune an LLM for our QG task. The dataset consists

of parallel pairs of product contexts (snippets of product reviews or catalog descriptions from Section 2), and the corresponding questions derived from those contexts. Note that the contexts can also be considered as the answers for the generated questions, since our approach by design ensures that the questions generated from the contexts are answerable by the contexts themselves.

We first collect an initial set of labeled question-context pairs via ICL, for products spanning a diverse set of features and categories. We manually inspect these data pairs to ensure that each product question satisfies the criteria described in Section 3.1, and if required, clean and/or rephrase the questions. This semi-automatic process results in the creation of a high quality supervised dataset of 1.8K question-context pairs. The questions cover a wide variety of product contexts, and range from short, simple questions to long questions specifying multiple aspects. We then fine-tune an LLM with this labeled data, to take a product context as input and generate a question about it.

Table 1: Evaluating product questions generated by our approach across several quality criteria using GPT-4.

Metric	ICL (zero-shot)	ICL (few-shot)	SFT
Relevance	0.80	0.86	0.81
Usefulness	0.72	0.73	0.75
Answerability	0.63	0.70	0.72
Fluency	0.98	0.99	0.95
Style	0.61	0.66	0.70

4 EXPERIMENTS AND RESULTS

The below design choices prioritize velocity for our offline pilot study, in terms of size, speed of training and inference, and performance on a manually curated validation set. However, as discussed in Section 5, our proposed approach is generalizable to any LLM.

Data: We construct a test dataset of 1,000 product contexts, spanning more than 20 product types [12], that we use to individually evaluate our proposed approach from Section 3.

Models: ICL questions are generated by the Claude-2 LLM.² For SFT, we train the 11B FLAN-T5-XXL [9] model with 1.8K pairs of product contexts and their corresponding question suggestions (from Section 3.2), for 8 epochs with an initial learning rate of $1e^{-5}$.

Evaluation Metrics: We evaluate the questions generated by our approach on each quality dimension defined in Section 3.1.

Automatic Evaluation: A diverse set of question suggestions are generated by our approach, with an average diversity $> 75\%$ across lists of questions, with respect to question length, lexical shape and product aspects mentioned. We use the GPT-4 LLM [1] to automatically evaluate the questions generated by our proposed approach, on each dimension described in Section 3.1. The results are displayed in Table 1. We observe that both the ICL and SFT variants of our approach achieve good performance on the relevance and fluency dimensions. There is a greater scope for improvement

in usefulness, answerability and question style, all of which achieve scores higher than 70% with our proposed approach. We also notice that more than 55% of the questions labeled as unanswerable were actually *partially answerable* by the provided context. The zero shot and few shot ICL variants perform comparably in case of usefulness and fluency metrics, and sometimes better than the SFT variant (for relevance and style), possibly because Claude is larger and better at following instructions than FLAN-T5-XXL [6, 17].

Human Evaluation: We next perform a human evaluation study to examine whether the human judgments made on the quality of the generated product questions align with the automatic evaluation outcomes from GPT-4. We randomly sample and annotate 75 examples from the test set across each quality dimension. We observe general trends largely similar to those noted with automatic evaluation. We obtain a high overall agreement $> 75\%$ between the human and GPT-4 annotations across multiple dimensions. The break down of this agreement for each individual dimension is as follows: 88% agreement for relevance, 61.33% for usefulness, 81.33% for answerability, 90.66% for fluency, and 66% for style. We observe a very high agreement between humans’ and GPT-4’s judgments greater than 80% for the relatively objective metrics of relevance, fluency and answerability. The lower agreement between humans and GPT-4 for usefulness and style is possibly due to some subjectivity involved in interpreting what is a useful or stylistically valid question a customer may ask a shopping assistant.

Example Question Suggestions: Table 2 shows diverse examples of product contexts and suggestion questions we generated for them. The questions are relevant and fluent, with varied question types and lengths. The longer questions effectively summarize the input contexts (rows 1, 6). Some questions are about product features that are more obvious or less useful (rows 5, 6). Row 3 depicts a question that can be both relevant and useful to potential buyers, but is not derived from the context and cannot be answered by it.

5 DISCUSSION

Having validated our approach via offline automatic and human evaluations in Section 4, we now detail how it could have real-world contributions and practical impact. If deployed in a shopping assistant, our approach could be used to enable automatic question suggestions to be generated and shown to customers across all products in the catalog within a conversational experience. Various configurations of the customer experience can then be evaluated via A/B testing to determine the optimal configuration.

Offline-Online Evaluation Alignment: Based on our offline evaluation, we expect a good quality mix of different question types to be also generated online (e.g., generic, recommendations, comparisons, feature based as in Table 2), ranging from short and concise to long and complex. However, some generated questions may have stylistic issues, such as the inclusion of first person or second person pronouns. Some generated questions may also be preference elicitation questions that an assistant might ask a customer, rather than a question that a customer might ask an assistant.

Potential Practical Impact: As described in Section 1, incorporating a question suggestion feature within a commercial shopping

²<https://www.anthropic.com/news/claude-2>

Table 2: Examples of product contexts and their corresponding question suggestions generated by our proposed approach.

Context type	Product Context	Generated Questions	Question Quality
Review (Commercial toilet paper holder)	UGH! DUMB! I thought this screwed to the wall, not GLUED to the wall! ... A great way to destroy a huge portion of my wall and rip it off when I go to remove it or it falls off. And I'm going to assume that if I did peel it off and it didn't take part of my wall/paint with it, that it would leave sticky residue behind...	Q1: Does the toilet paper holder stick or screw to the wall, as advertised? Q2: Will removing the toilet paper holder leave behind sticky residue or cause wall damage?	Both questions are relevant, useful, answerable, fluent, diverse, and summarize the review well.
Review (Cat repellent mats)	Does NOT work for cats - false advertising! ... This thing is a joke, and cats are not even afraid of the beep anyway. I wish we could return it for a refund...	Q1: Are cats afraid of the beep emitted from the mat?	Relevant, useful, answerable, fluent and concise.
Catalog Description (Rongeurs)	Item Package Height: 20.32 Centimeters. Autoclavable. Medical Grade Steel. Unit Count: 1.0	Q1: Can the rongeurs be reused after cleaning and sterilization?	Relevant, useful, fluent, but unanswerable by the context
Review (Salted soy nuts)	I loved these because sometimes your looking for a salty snack that doesn't have so much fat in them like potato chips. This product satisfies the urge nicely	Q1: Are the salted soy nuts a healthy alternative to potato chips? Q2: Which snack has more fat, potato chips or the salted soy nuts?	Both questions are relevant, useful, answerable and fluent, but are not very diverse from each other
Catalog Description (Smart doorbell)	Built-in speaker/microphone. Talk to visitors anytime anywhere... Intelligent real-time monitoring via mobile phone... Video recording, picture-taking and screenshot, video playback and storage. Play videos anytime anywhere via mobile phone. When storage is full, the earliest recorded video will be overwritten.	Q1: What are key features of this doorbell? Q2: Can the camera enable mobile phone monitoring? Q3: Can the camera take pictures, record videos, and store them on a mobile device?	Q1, Q2, Q3 are relevant, fluent and answerable. Q1, Q3 are less useful to customers making a purchase decision. Q2, Q3 are not very diverse.
Review (Hair growth tonic)	This came in a day, very fast. Best ever, was sceptical at first but persevered. I have been using it for like a week and have seen that my eyelashes did look fuller.	Q1: How fast was the hair growth tonic delivered, and did it make eyelashes look fuller after use?	Relevant, fluent, slightly useful, summarizes the review well, answerable.

assistant is a novel and useful real-world application, that millions of customers can adopt and engage with. Recommending a diverse mix of both broad and product-specific questions to customers can result in favorable outcomes such as increased engagement time with the assistant and clicks, since it is both easier and faster for customers to click on the suggested questions, rather than type their own. Question suggestions followed by their answers can impart new and beneficial information to customers about product aspects, that they may not have known or considered during their shopping journey, thus increasing customers' satisfaction with the shopping assistant. Mechanisms such as caching specific model outputs or processing and delivering the generated tokens in a streaming fashion can effectively reduce the latency of real-time question generation. Overall, there is immense potential in using LLM-based question suggestions to enrich the purchase experience of real world customers conversing with shopping assistants.

6 CONCLUSION AND FUTURE WORK

We focused on improving user adoption of shopping assistants using contextual question suggestions, to create a more convenient and seamless experience for customers. We proposed a framework to automatically generate relevant, useful, answerable and diverse product questions from salient product context, that can be suggested by a shopping assistant to customers as helpful conversation starters, continuations, or inspirations for their own questions.

Looking forward, we envision generative AI to significantly evolve into an even more integral part of the conversational shopping landscape, with digital assistants acting as personalized shopping guides for online buyers. In particular, the task of automatically generating beneficial and answerable suggestion questions that a customer may ask a shopping assistant is an essential and challenging problem in the e-commerce domain, to promote faster and smoother conversational interactions. There is scope for future improvement, especially to generate more useful and answerable product question suggestions in the desired style. One way of achieving this could be through the use of instruction fine-tuning [17] with our labeled data, a direction we leave for future work.

We also plan to extend our approach to incorporate multi-turn conversational history into the input context. We seek to generate personalized, customer-specific product question suggestions based on their past search or purchase behavior. We further aim to leverage customer interaction and/or feedback signals (e.g., clicks, likes) to control, modify and improve the questions being generated and presented to customers. Such developments will contribute towards boosting the overall dialogue quality, enriching customers' e-commerce experiences, and promoting deeper, enhanced digital interactions that are on-par with human shopping assistants.

ACKNOWLEDGEMENTS

We sincerely thank Eugene Agichtein, Besnik Fetahu, Giuseppe Castellucci, Jason Choi, Zhiyu Chen, Saar Kuzi and Elad Kravi for their valuable feedback and participation in the CBQG project. We are also grateful to the reviewers for their insightful comments.

PRESENTER BIOGRAPHY

Nikhita Vedula is an Applied Scientist at Amazon within the Shopping and Search organization. She received her B.S. degree from the National Institute of Technology, Nagpur, India in 2015, and her Ph.D. from the Ohio State University in 2020. Her research interests and contributions span the fields of natural language processing, conversational AI and information retrieval. She has published more than twenty research papers in leading peer-reviewed conferences such as SIGIR, WSDM, the Web conference, EMNLP and ACL. Her research has been recognized with one Best Paper award, one Best Paper Honorable Mention, and one Outstanding Paper award. She also regularly serves as a Program Committee member and reviewer of several prestigious conferences in her field such as SIGIR, ACL, the Web Conference, WSDM and CIKM.

REFERENCES

- [1] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774* (2023).
- [2] Janarthanan Balakrishnan and Yogesh K Dwivedi. 2021. Conversational commerce: entering the next stage of AI-powered digital assistants. *Annals of Operations Research* (2021), 1–35.
- [3] Ankita Bihani, Jeffrey D Ullman, and Andreas Paepcke. 2018. FAQtor: Automatic FAQ generation using online forums. In *International Conference on Educational Data Mining*. 529–532.
- [4] Yang Deng, Wenxuan Zhang, Qian Yu, and Wai Lam. 2023. Product Question Answering in E-Commerce: A Survey. *arXiv preprint arXiv:2302.08092* (2023).
- [5] Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, and Haofen Wang. 2023. Retrieval-augmented generation for large language models: A survey. *arXiv preprint arXiv:2312.10997* (2023).
- [6] Yingqiang Ge, Wenyue Hua, Kai Mei, Juntao Tan, Shuyuan Xu, Zelong Li, Yongfeng Zhang, et al. 2024. Openagi: When llm meets domain experts. *Advances in Neural Information Processing Systems* 36 (2024).
- [7] Bilal Ghanem, Lauren Lutz Coleman, Julia Rivard Dexter, Spencer von der Ohe, and Alona Fyshe. 2022. Question Generation for Reading Comprehension Assessment by Modeling How and What to Ask. In *Findings of the Association for Computational Linguistics: ACL 2022*, Smaranda Muresan, Preslav Nakov, and Aline Villavicencio (Eds.). Association for Computational Linguistics, Dublin, Ireland, 2131–2146. <https://doi.org/10.18653/v1/2022.findings-acl.168>
- [8] Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. Large language models are zero-shot reasoners. *Advances in neural information processing systems* 35 (2022), 22199–22213.
- [9] Shayne Longpre, Le Hou, Tu Vu, Albert Webson, Hyung Won Chung, Yi Tay, Denny Zhou, Quoc V Le, Barret Zoph, Jason Wei, et al. 2023. The flan collection: Designing data and methods for effective instruction tuning. *arXiv preprint arXiv:2301.13688* (2023).
- [10] Yosi Mass, Boaz Carmeli, Haggai Roitman, and David Konopnicki. 2020. Unsupervised FAQ retrieval with question generation and BERT. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. 807–812.
- [11] Nikahat Mulla and Prachi Gharpure. 2023. Automatic question generation: a review of methodologies, datasets, evaluation metrics, and applications. *Progress in Artificial Intelligence* 12, 1 (2023), 1–32.
- [12] Jianmo Ni, Jiacheng Li, and Julian McAuley. 2019. Justifying recommendations using distantly-labeled reviews and fine-grained aspects. In *Proceedings of the 2019 conference on empirical methods in natural language processing and the 9th international joint conference on natural language processing (EMNLP-IJCNLP)*. 188–197.
- [13] Vipula Rawte, Amit Sheth, and Amitava Das. 2023. A survey of hallucination in large foundation models. *arXiv preprint arXiv:2309.05922* (2023).
- [14] Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, et al. 2022. Emergent abilities of large language models. *arXiv preprint arXiv:2206.07682* (2022).
- [15] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems* 35 (2022), 24824–24837.
- [16] Guangxuan Xiao, Yuandong Tian, Beidi Chen, Song Han, and Mike Lewis. 2023. Efficient streaming language models with attention sinks. *arXiv preprint arXiv:2309.17453* (2023).
- [17] Shengyu Zhang, Linfeng Dong, Xiaoya Li, Sen Zhang, Xiaofei Sun, Shuhe Wang, Jiwei Li, Runyi Hu, Tianwei Zhang, Fei Wu, et al. 2023. Instruction tuning for large language models: A survey. *arXiv preprint arXiv:2308.10792* (2023).
- [18] Yue Zhang, Yafu Li, Leyang Cui, Deng Cai, Lema Liu, Tingchen Fu, Xinting Huang, Enbo Zhao, Yu Zhang, Yulong Chen, et al. 2023. Siren’s song in the AI ocean: a survey on hallucination in large language models. *arXiv preprint arXiv:2309.01219* (2023).
- [19] Yuchen Zhuang, Yue Yu, Kuan Wang, Haotian Sun, and Chao Zhang. 2024. Toolqa: A dataset for llm question answering with external tools. *Advances in Neural Information Processing Systems* 36 (2024).