



CORAL: Benchmarking Multi-turn Conversational Retrieval-Augmented Generation

Yiruo Cheng¹, Kelong Mao¹, Ziliang Zhao¹, Guanting Dong¹, Hongjin Qian²,
Yongkang Wu³, Tetsuya Sakai⁴, Ji-Rong Wen¹, Zhicheng Dou^{1*}

¹Gaoling School of Artificial Intelligence, Renmin University of China

²Beijing Academy of Artificial Intelligence

³Huawei Poisson Lab

⁴Waseda University, Tokyo, Japan

{chengyr, mk1, dou}@ruc.edu.cn

Abstract

Retrieval-Augmented Generation (RAG) has become a powerful paradigm for enhancing large language models (LLMs) through external knowledge retrieval. Despite its widespread attention, existing academic research predominantly focuses on single-turn RAG, leaving a significant gap in addressing the complexities of multi-turn conversations found in real-world applications. To bridge this gap, we introduce CORAL, a large-scale benchmark designed to assess RAG systems in realistic multi-turn conversational settings. CORAL includes diverse information-seeking conversations automatically derived from Wikipedia and tackles key challenges such as open-domain coverage, knowledge intensity, free-form responses, and topic shifts. It supports three core tasks of conversational RAG: passage retrieval, response generation, and citation labeling. We propose a unified framework to standardize various conversational RAG methods and conduct a comprehensive evaluation of these methods on CORAL, demonstrating substantial opportunities for improving existing approaches. Our dataset and code are available at <https://github.com/Ariya12138/CORAL>.

1 Introduction

Retrieval-Augmented Generation (RAG) has emerged as a promising approach in question answering, leveraging large language models (LLMs) alongside external knowledge retrieval to enhance the quality and accuracy of generated responses (Lewis et al., 2020; Guu et al., 2020; Huang and Huang, 2024). While RAG has gained traction both in academia and industry, a notable gap exists between how it is studied in academic settings and how it is implemented in real-world systems. Academic research predominantly focuses on single-turn interactions (Tan et al., 2024b; Jin et al., 2024; Wang et al., 2024; Dong et al.,

2024a; Zhu et al., 2024; Li et al., 2025), whereas most industrial RAG systems (OpenAI, 2022; Anthropic, 2023; AI, 2023; Team, 2023) are designed to handle multi-turn conversations. In practice, multi-turn conversation is the norm, where RAG systems must dynamically adapt to evolving context and user intent across multiple turns.

The shift from single-turn to multi-turn conversations introduces unique challenges for RAG. In multi-turn settings, systems must deal with redundant or irrelevant information from prior interactions and cope with abrupt topic shifts (Ye et al., 2024; Adlakha et al., 2022). This complexity can degrade the retrieval and generation quality, especially as conversation histories grow, exacerbating the “long context problem” (Ratner et al., 2023; Ding et al., 2024). These issues highlight the need for dedicated research into multi-turn conversational RAG to address the realities of interactive, ongoing dialogue.

However, progress in this area is severely hindered by the lack of a comprehensive benchmark designed to evaluate conversational RAG systems. To align with the diverse and complex real-world applications of conversational RAG systems, we identify several critical features such a benchmark should satisfy: (1) open-domain coverage, allowing the system to handle questions from a wide range of topics; (2) knowledge-intensiveness, challenging systems to retrieve and generate responses that require deep, contextual knowledge; (3) free-form response generation, ensuring that models can produce detailed, contextually rich answers; (4) handling of topic shifts, evaluating the system’s ability to manage sudden changes in dialogue context without carrying over irrelevant information from previous turns; and (5) citation labeling, promoting transparency by requiring the system to cite the sources of the information it retrieves.

Unfortunately, no existing dataset satisfies all of these features. Although there are datasets for

*Corresponding author.

Dataset	Open-domain	Knowledge-Intensive	Free-form Response	Topic Shift	Citation Labeling
CORAL (ours)	✓	✓	✓	✓	✓
Wizard of Wikipedia (Dinan et al., 2019)	✓	✓	✓	✗	✗
CoQA (Reddy et al., 2019)	✗	✓	✓	✗	✗
OR-QuAC (Qu et al., 2020)	✓	✓	✗	✗	✗
Doc2Dial (Feng et al., 2020)	✗	△	✓	✗	✗
TREC CAsT19 (Dalton et al., 2020b)	✓	✓	✗	✗	✗
TREC CAsT20 (Dalton et al., 2020a)	✓	✓	✓	✓	✗
TREC CAsT21 (Dalton et al., 2021)	✓	✓	✓	✓	✗
QReCC (Anantha et al., 2021)	✓	✓	✓	✗	✗
TREC CAsT22 (Owoicho et al., 2022)	✓	✓	✓	✓	✗
KETOD (Chen et al., 2022)	✓	△	✓	✗	✗
TopiOCQA (Adlakha et al., 2022)	✓	✓	✓	✓	✗
INSCIT (Wu et al., 2023)	✓	✓	✓	✗	✗
TREC iKAT23 (Aliannejadi et al., 2024)	✓	✓	✓	✗	✗
RAD-Bench (Kuo et al., 2024)	✓	✓	✓	✗	✗
MTRAG (Katsis et al., 2025)	✓	✓	✓	✗	✗

Table 1: Comparison of CORAL with other conversational search and conversational QA datasets. △ indicates that only a portion of the dataset satisfies the property.

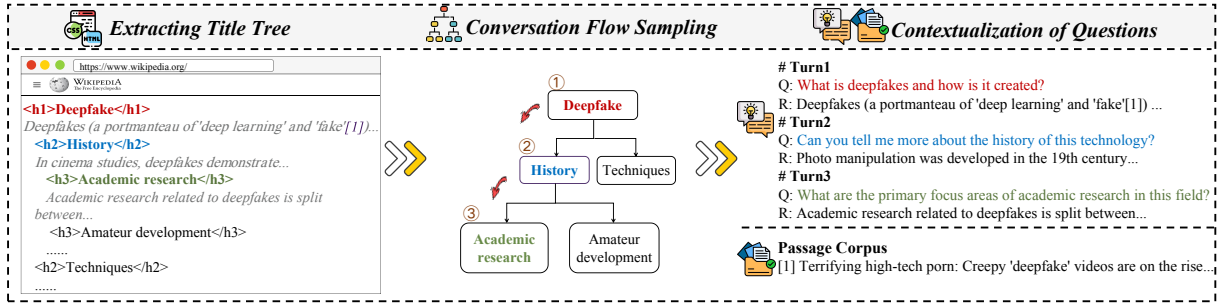
related tasks, such as conversational search and question answering, they do not adequately address the unique challenges of benchmarking multi-turn conversational RAG systems that align with all of the above features. For instance, datasets in conversational search like TREC CAsT (Dalton et al., 2020b,a, 2021; Owoicho et al., 2022) primarily focus on retrieval tasks and lack the capacity to assess a system’s generative abilities in producing free-form answers. Likewise, conversational QA datasets, such as QReCC (Anantha et al., 2021) and TopiOCQA (Adlakha et al., 2022), predominantly offer short, factual answers, falling short of reflecting the nuanced, long-form responses often required in practical applications. We summarize the limitations of existing related datasets in Table 1. These limitations highlight the pressing need for a more comprehensive benchmark that meets the full spectrum of requirements for evaluating conversational RAG.

In this paper, we introduce a large-scale multi-turn **CO**nversational **R**etrieval-**A**ugmented **G**eneration **L**anguage **B**enchmark (CORAL) that fulfills the above critical features to systematically evaluate and advance conversational RAG systems. In general, CORAL is derived from English Wikipedia web pages, containing a total of 8,000 diverse information-seeking conversations. We propose a novel approach to automatically convert Wikipedia content into conversational formats, with each conversation generated through tailored sampling from either a single page or multiple related pages. Specifically, as depicted in Figure 1, we treat the (sub)titles of Wikipedia pages as the source of questions, using the corresponding

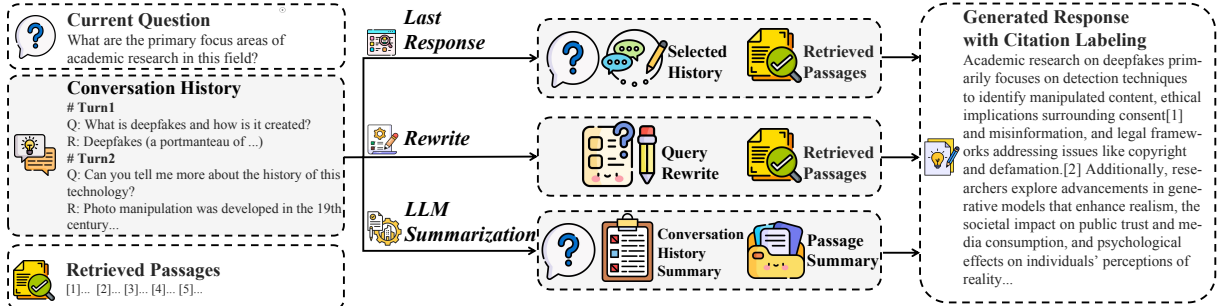
human-written Wikipedia content serving as high-quality free-form responses. The content itself is originally well-cited and includes related passages for retrieval, making Wikipedia a particularly suitable source for constructing conversational RAG datasets. To ensure coherent and diverse conversation flow, we design four sampling strategies based on the natural hierarchical properties of Wikipedia pages to first create the conversation flow. We then utilize powerful LLMs (e.g., GPT-4), to refine the original Wikipedia titles into well-formed conversational queries by incorporating contextual dependencies, such as co-reference and omission, resulting in the final conversation.

In CORAL, we evaluate conversational RAG systems across three essential tasks: (1) *Conversational Passage Retrieval*, which assesses the system’s ability to retrieve the relevant information from a large document set based on multi-turn context; (2) *Response Generation*, which tests the system’s capacity to generate accurate, contextually rich answers; and (3) *Citation Labeling*, which ensures that the generated responses are transparent and grounded by requiring correct attribution of sources. These tasks are fundamental for measuring the core capabilities of conversational RAG systems in real-world, multi-turn settings.

Additionally, we present a unified framework that standardizes various conversational RAG baselines and conduct a comprehensive evaluation of them on the CORAL benchmark. We find that the fine-tuned open-source LLM outperforms the commercial closed-source LLM in the retrieval stage, and shortening the input length to filter noise can not only maintain response quality but also improve



(a) Overview of the CORAL dataset construction process



(b) Overview of three conversation compression strategies

Figure 1: Part (a) is an overview of the CORAL dataset construction process. The red arrows show the sampled conversation flow, with numerical labels on the nodes indicating the round of the sampled conversation turns. The content under each sampled (sub)title serves as the conversational response in CORAL. Part (b) is the three conversation compression strategies in conversational RAG.

citation labeling accuracy.

In summary, the contributions of our work are:

(1) We present an automatic and novel approach for constructing large-scale conversational RAG benchmarks from Wikipedia, leveraging its hierarchical structure and high-quality content to create diverse information-seeking conversations.

(2) The CORAL benchmark itself is unique in its comprehensive coverage of critical features, filling a significant gap in the evaluation of conversational RAG systems.

(3) We develop a unified framework for standardizing and evaluating various conversational RAG baselines, facilitating systematic comparison and advancement in this rapidly evolving field.

2 Related Work

2.1 Retrieval-Augmented Generation

Existing RAG studies (Tan et al., 2024a; Dong et al., 2024b; Li et al., 2024b) primarily focus on the optimization of individual components. The rewriter (Li et al., 2024a; Wang et al., 2023a; Baek et al., 2024) module interprets and reconstructs user queries to align them more effectively with the search process. The reranker (Ma et al., 2023; Dong et al., 2024c; Xu et al., 2024) module inde-

pendently adjusts the ordering of retrieved documents based on their relevance. The post-retrieval processing (Xu et al., 2023; Yang et al., 2023; Wang et al., 2023b; Jiang et al., 2024; Jin et al., 2024) module then reduces the volume of these documents, stripping away non-essential information to focus solely on the content critical for generating precise responses.

While these works have significantly advanced single-turn RAG systems, gaps remain in multi-turn conversational settings (Ye et al., 2024). We address this by proposing a new benchmark and framework specifically for conversational RAG.

2.2 Conversational Search

Conversational search enables users to interact with retrieval systems through multi-turn dialogues (Mo et al., 2024a). Two main approaches are conversational query rewriting (CQR) and conversational dense retrieval (CDR). CQR transforms context-dependent queries into fully rewritten versions for ad-hoc retrieval, focusing on selecting relevant tokens from the conversation history (Voskarides et al., 2020; Kumar and Callan, 2020; Lin et al., 2021b) or using LLMs to generate rewrites (Lin et al., 2020; Yu et al., 2020; Vakulenko et al., 2021; Wu et al., 2022). CDR jointly encodes conversation

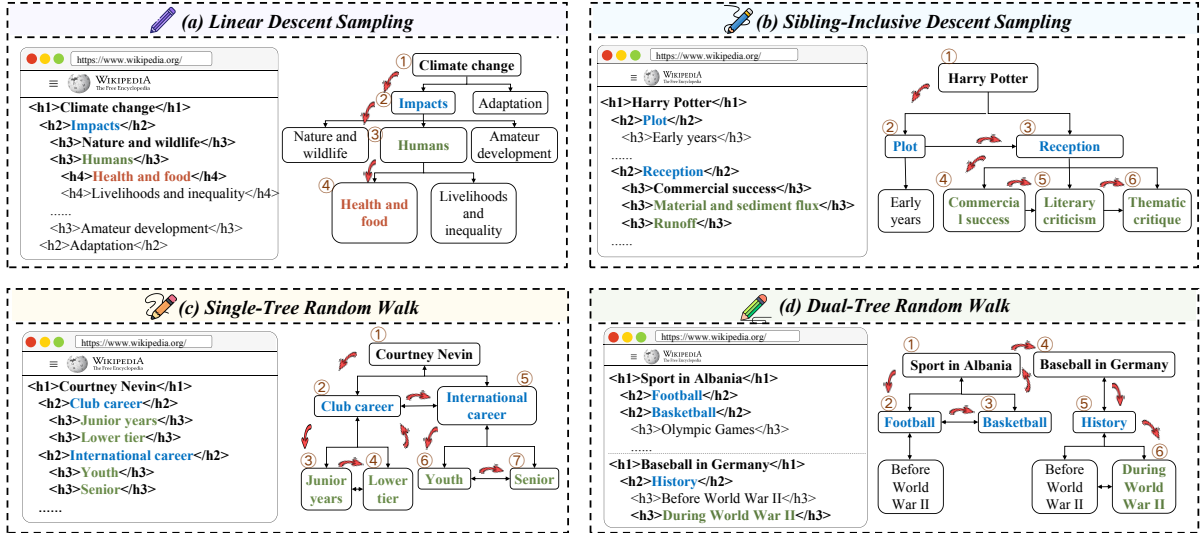


Figure 2: Illustration of the four sampling strategies. The red arrows show the sampled conversation flow, with numerical labels on the nodes indicating the round of the sampled conversation turns.

history and the current query for end-to-end dense retrieval (Yu et al., 2021; Mao et al., 2024).

Challenges like limited training data are addressed through data augmentation (Lin et al., 2021a; Mao et al., 2022b; Dai et al., 2022; Jin et al., 2023; Chen et al., 2024; Mo et al., 2024c), and context denoising (Mao et al., 2022a; Mo et al., 2023; Mao et al., 2023b; Mo et al., 2024b) improves retrieval by filtering irrelevant conversation history. However, a benchmark is still needed for evaluating response generation and citation labeling.

3 CORAL

3.1 Data Source

We choose Wikipedia as our data source for the following reasons, which align with the critical features in Table 1. (1) Wikipedia pages are well-structured and enriched by global volunteers, covering a broad range of topics; (2) The logically interconnected titles provide a strong foundation basis for generating diverse queries, with each representing a distinct intent. (3) The human-authored content under each title includes references that not only allow for free-form responses with precise citation labeling but also serve as the golden retrieval evidence for their respective titles.

However, the content may include noisy text, and reference pages are often too long for effective retrieval. We follow previous work (Qian et al., 2023) to clean the Wikipedia pages. Specifically, for content, we remove Wikipedia templates, special symbols, and other invalid text. For references,

we first split the reference pages into smaller passages. Then, we exclude passages shorter than 16 words or with a non-English token ratio exceeding 0.3, and then calculate term recall to identify suitable passages. After these refinements, we generate a clean set of 20,000 high-quality pages for subsequent conversations.

3.2 The CORAL Dataset Construction

We transform one or more related Wikipedia web pages into information-seeking conversations through a three-stage approach.

3.2.1 Extracting Title Trees

First, we extract all subheadings (i.e., titles) from the raw HTML of the Wikipedia pages. These subheadings create a natural hierarchy for the content, enabling us to construct a title tree, where the page title (H1-level heading) serves as the root. Subsequent headings (e.g., H2 to H6) divide the content into progressively detailed sections, with each level corresponding to a node’s depth in the tree. The directional links between nodes will dictate the flow of the generated conversations. Besides, to enhance the complexity and diversity of conversations, we also adjust the depth, breadth, multi-subtopic exploration, and topic shifts during the construction of these title trees.

3.2.2 Conversation Flow Sampling

To generate coherent and diverse conversations, we implement the following four sampling strategies based on the extracted title trees:

	LDS		SIDS		STRW		DTRW	
	Train	Test	Train	Test	Train	Test	Train	Test
# Conversation	1800	200	1800	200	1800	200	1800	200
# Turns	5934	651	16082	1727	18165	1949	19411	2153
# Turns / Conversation	3.30	3.26	8.93	8.64	10.09	9.75	10.78	10.77
# Tokens / Question	13.70	13.89	12.62	12.64	12.72	12.88	14.15	14.75
# Tokens / Response	233.81	147.16	242.54	155.54	243.34	191.60	300.47	259.72
# Positive passages/ Turn	3.25	2.03	2.64	1.73	3.01	2.12	3.98	3.50

Table 2: Data statistics of four different conversation structures.

(1) *Linear Descent Sampling (LDS)*: This strategy begins at the root node and permits movement only from parent nodes to their child nodes. LDS serves as the most basic sampling path, emulating the progressive logic typical of real conversational information-seeking scenarios. As illustrated in Figure 2(a), the title tree starts with the overall theme of climate change and progressively narrows down to specific impacts associated with this global issue. Following the red arrow, the focus shifts to the human aspects, particularly examining how climate change affects human health and food security. This structure exemplifies a gradual deepening of the query intent as the conversation unfolds.

(2) *Sibling-Inclusive Descent Sampling (SIDS)*: This strategy builds on LDS by introducing directional links between sibling nodes. This feature is essential because conversational processes often encompass both in-depth and parallel explorations of related subtopics. As shown in Figure 2(b), when discussing the reception of Harry Potter, the subsequent three rounds of dialogue analyze it from three distinct perspectives: commercial success, literary criticism, and thematic critique. This enhancement enriches the breadth of discussions within the conversation structure.

(3) *Single-Tree Random Walk (STRW)*: This strategy further enhances SIDS by incorporating interconnections among sibling nodes as well as between parent and child nodes. Essentially, it forms a directed graph with bidirectional edges. As illustrated in Figure 2(c), after an in-depth exploration of Courtney Nevin’s club career, the focus shifts to her international career.

(4) *Dual-Tree Random Walk (DTRW)*: It mimics the topic shifts that occur in real conversational scenarios, allowing for greater flexibility. It enables transitions between two different but somewhat related trees, which are retrieved using the root title

as a query and employing the BM25 algorithm. As illustrated in Figure 2(d), the conversation shifts from sports in Albania to baseball in Germany.

3.2.3 Contextualization of Questions

As introduced in Section 3.1, we treat the subtitles as the sources of questions, with their corresponding contents serving as the responses. In this final stage, we contextualize the keyword subtitles into conversational questions to enhance the realism of the conversation.

Specifically, for each turn, we first create a keyword chain that includes the current node and all its ancestor nodes. This keyword chain, along with the response of the current node, is then used to prompt GPT-4¹ to rewrite the original keyword title into a natural language question. We then continue to prompt GPT-4 to further contextualize the question into a conversational format by incorporating linguistic phenomena such as ellipses, references, and omissions (Dalton et al., 2020b), which are prevalent in real conversational scenarios. The prompt details are provided in Appendix A.

3.3 The Final Dataset Format and Statistics

The key statistics of CORAL are summarized in Table 2. Our dataset consists of 8,000 conversations with the four types introduced in Section 3.2.2. These 8,000 conversations are evenly distributed across four distinct structural types, with each type containing 2,000 conversations. Specifically, the LDS conversation type includes 3 to 6 turns per conversation. For the remaining types—SIDS, STRW, and DTRW—each category consists of 1,600 sets of conversations with 6 to 10 turns, along with an additional 400 sets featuring 11 to 20 turns per conversation. The design of the longer conversation

¹gpt-4-turbo-2024-04-09 from <https://openai.com/api>

Category	Method	MRR	MAP	NDCG@3	Recall@20	Recall@100
CDR Models	Conv-ANCE-Q	19.8	28.6	20.5	39.1	51.0
	KD-ANCE-Q	22.6	33.1	24.5	38.5	48.0
	Conv-ANCE-C	20.5	29.6	21.1	39.8	53.4
	KD-ANCE-C	23.2	33.6	24.9	40.3	49.6
CQR Models	LLM4CS (GPT-3.5)	21.2	31.1	23.0	35.5	44.4
	Qwen2.5-1.5B	16.3	23.8	17.2	31.0	39.2
	Qwen2.5-1.5B-SFT	23.1	33.6	25.1	39.4	48.6

Table 3: Retrieval performance comparisons. The best performance is **bold**. *Conv-ANCE-Q* denotes the Conv-ANCE is trained on the QReCC dataset and *Conv-ANCE-C* denotes the Conv-ANCE is trained on CORAL training dataset.

intends to simulate real-world challenges encountered in conversational scenarios, such as redundant information and the long context problem.

Our final dataset format is as follows: A conversation $\mathbf{C} = \{(q_i, r_i)\}_{i=1}^n$ comprised of n turns. q_i is a contextualized query of the i -th turn generated in Section 3.2.3, and r_i is the i -th turn golden response, which is the cleaned plain text under the corresponding (sub-)title in the HTML. The supporting web pages for r_i , listed in the HTML Reference Section, can be processed as described in Section 3.1 to serve as the golden relevant passages $P_i^+ = \{p_{i,1}, p_{i,2}, \dots\}$. On average, each conversation turn has 3.17 related passages, and the average golden response length is 255 tokens. Finally, we obtain a passage corpus \mathcal{P} , which contains 200K passages from all the golden references P_i^+ .

3.4 Evaluation Tasks

CORAL mainly supports three fundamental conversational RAG tasks:

(1) *Conversational Passage Retrieval*: This task evaluates a system’s capability to extract relevant information from extensive document collections, considering the context of multi-turn conversations. Formally, given the k -th question q_k and the corresponding conversation history $H_k = \{q_i, r_i\}_{i=1}^{k-1}$, where q_i and r_i respectively denote the question and response of the i -th turn, the retriever \mathbb{R} aims to retrieve the relevant passages P_k from the passage corpus \mathcal{P} . We use MRR, MAP, NDCG@3, Recall@20 and Recall@100 as retrieval evaluation metrics.

(2) *Response Generation*: This task challenges the system’s ability to produce accurate, detailed, and contextually appropriate answers. Given the k -th question q_k , the corresponding conversation history H_k , and the relevant passages P_k , the generator \mathbb{G} needs to generate an informative response to answer the question. We use rule-

based metrics BLEU-1 (Papineni et al., 2002), and ROUGE-L (Lin, 2004) to evaluate the response quality compared with r_k . Given the lengthier responses in our benchmark, we additionally utilize the model-based evaluation method proposed in RichRAG (Wang et al., 2024).

(3) *Citation Labeling*: This task evaluates the method’s ability to accurately attribute information sources within the generated responses. Following ALCE (Gao et al., 2023), the generated response r_k consists of n statements s_1, s_2, \dots, s_n . Each statement s_i cites a list of passages $C_i = \{c_{i,1}, c_{i,2}, \dots\}$, where $c_{i,j} \in P_k$. We adopt Citation Recall and Citation Precision defined in ALCE (Gao et al., 2023) to evaluate the accuracy of citation labeling.

4 Conversational RAG Framework

A conversational RAG system typically comprises a retriever and a generator to handle the current user query q_k , the conversation history H_k , and the retrieved passages P_k . As the conversation progresses, both the growing conversation history and the noisy retrieved passages can negatively impact the system’s efficiency and effectiveness, making it harder to generate accurate responses. To solve the problem, we propose a simple compression framework to efficiently manage these inputs. Specifically, we introduce a conversation compression function f to compress the conversation, and then use the compressed contents as the real inputs of retrievers and LLM generators. In addition to conversation compression, we also apply post-retrieval results compression. Following existing approaches (Xu et al., 2023), we simply take LLMs as the compression function f_p , leaving the exploration of more compression methods in future work.

Formally, suppose $f(H_k)$ is the compressed conversation context, $P_k = \mathbb{R}(f(H_k), q_k)$ is list of passages retrieved by querying $f(H_k)$ with q_k ,

Category	# Tokens	Model	Generation		Citation Labeling	
			BLEU-1	ROUGE-L	Citation Recall	Citation Precision
Raw Context	2226	Qwen2.5-7B	22.2	13.1	3.1	18.1
		Mistral-7B	18.1	12.4	2.4	4.8
		Llama-3.1-8B	21.5	12.9	0.9	2.1
		Qwen2.5-7B-SFT	18.3	18.5	6.6	16.8
		Mistral-7B-SFT	23.7	20.1	4.6	11.1
		Llama-3.1-8B-SFT	24.2	19.7	4.2	9.8
Last Response	1474	Qwen2.5-7B	20.9	12.8	3.5	20.8
		Mistral-7B	18.1	12.3	2.7	4.5
		Llama-3.1-8B	20.4	12.7	1.3	3.1
		Qwen2.5-7B-SFT	23.9	16.5	10.4	24.8
		Mistral-7B-SFT	21.8	18.5	5.0	12.4
		Llama-3.1-8B-SFT	26.1	18.1	3.5	8.7
Rewrite	1236	Qwen2.5-7B	21.1	12.8	2.4	9.4
		Mistral-7B	18.8	12.3	2.5	3.8
		Llama-3.1-8B	18.8	12.4	1.7	3.1
		Qwen2.5-7B-SFT	18.9	16.4	7.4	16.8
		Mistral-7B-SFT	24.8	18.5	5.9	14.8
		Llama-3.1-8B-SFT	26.3	18.2	4.7	11.2
LLM Summarization	478	Qwen2.5-7B	21.0	12.7	2.9	13.0
		Mistral-7B	19.5	12.3	5.6	6.7
		Llama-3.1-8B	19.1	12.8	4.1	7.1
		Qwen2.5-7B-SFT	23.5	16.8	14.1	31.1
		Mistral-7B-SFT	16.9	17.1	8.3	19.8
		Llama-3.1-8B-SFT	18.7	16.5	4.5	10.7

Table 4: The comparison of different LLMs on response generation and citation labeling. # Tokens denotes the number of input tokens.

and $f_p(P_k)$ is the compressed results of the retrieval, the final generation task can be formulated as: $\mathbb{G}(q_k, f(H_k), f_p(P_k))$. The prompt for feeding $q_k, f(H_k), f_p(P_k)$ into the generator can be found in Appendix B. Various existing conversational RAG methods can be unified into our framework. In this work, we mainly investigate the following three methods for the conversation compression:

Last Response Strategy For the conversation history, we heuristically select all previous conversational questions $\{q_i\}_1^{k-1}$ and the last turn’s response r_{k-1} in the conversation history:

$$f_c^{\text{LR}}(H_k) = \{q_i\}_1^{k-1}, r_{k-1}. \quad (1)$$

Rewrite Strategy We adopt a conversational query rewriting model Rewrite() to transform the original query along with the conversation history into a standalone question rewrite \hat{q}_k :

$$f_c^{\text{RW}}(H_k) = \hat{q}_k = \text{Rewrite}(q_k; H_k). \quad (2)$$

In this strategy, $P_k = \mathbb{R}(\hat{q}_k)$ is list of passages retrieved by querying \hat{q}_k , and $f_p(P_k)$ is the compressed results of the retrieval, the final generation task can be formulated as: $\mathbb{G}(\hat{q}_k, f_p(P_k))$.

LLM Summarization Strategy Inspired by RECOMP (Xu et al., 2023), we use LLMs to generate abstractive summary of the conversation history:

$$f_c^{\text{SUM}}(H_k) = \text{LLM}(H_k). \quad (3)$$

The prompt is shown in Appendix C.

5 Experiments

In this section, we discuss the performance of conversational RAG on our benchmark, and provide a comprehensive analysis for each stage.

5.1 Evaluating Retrieval Performance

We concentrate on two main approaches in conversational search: conversational dense retrieval (CDR) and conversational query rewriting (CQR). For CDR, we use KD-ANCE and Conv-ANCE with ANCE as the base retriever. KD-ANCE (Yu et al., 2021) trains the session encoder by mimicking golden query embeddings, while Conv-ANCE (Karpukhin et al., 2020; Lin et al., 2021a) uses contrastive learning to train the session encoder, drawing it closer to relevant passages and further from irrelevant ones. For CQR, we utilize the LLM4CS (Mao et al., 2023a), which incorporates GPT-3.5, and an open-source LLM for

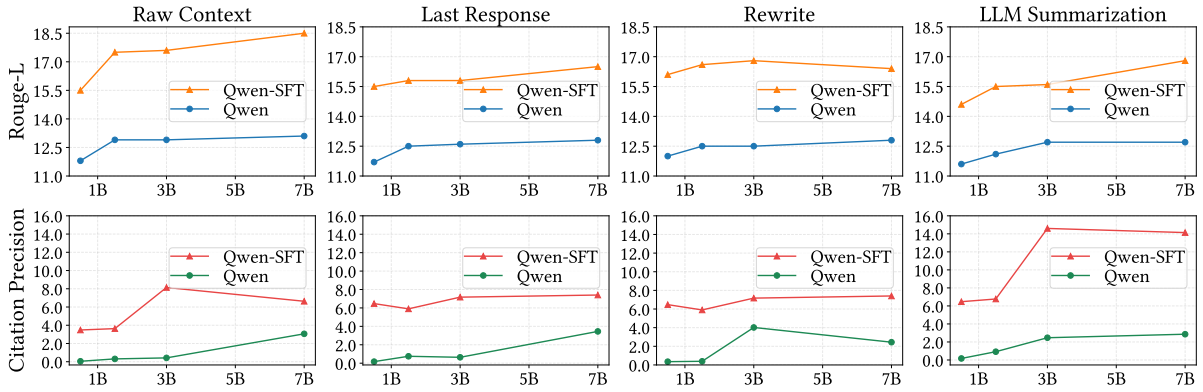


Figure 3: The scaling analysis of generation and citation labeling performance.

generating query rewrites respectively to enable a comparative analysis. Table 3 provides a detailed comparison between these two categories. We have the following observations:

(1) The performances of the CDR and CQR models are fairly comparable. Notably, the Qwen2.5-1.5B-SFT shows a strong competitive edge, not only surpasses the Qwen2.5-1.5B but also outperforms the advanced closed-source LLM GPT-3.5 across all evaluated metrics.

(2) KD-ANCE in the CDR category shows better results compared to Conv-ANCE. This may be attributed to the training methodologies: KD-ANCE possibly leverages golden rewrite data more effectively than Conv-ANCE, which uses in-batch negatives that may not be sufficiently challenging for optimal learning.




Figure 4: The GPT-4 evaluation score.

5.2 Evaluation Response Generation with Citation Labeling

We compare the raw context baselines with another three conversation compression strategies introduced in Section 4, selecting Qwen2.5, Mistral, and Llama as generators. We prompt the generator to generate the response along with the citations in the response. The generation and citation labeling performance is shown in Table 4, and the GPT-4 score is shown in Figure 4. We find that:

(1) By examining four methods of modeling conversation history, we observe an interesting trend: as the input is progressively condensed (from 2226 input tokens in the raw context to merely 478 input tokens in the LLM Summarization), the decrease in performance is surprisingly minimal, and there is even an observed improvement in terms of citation labeling. This suggests that some content within the dialogue history is irrelevant or redundant and can be removed without negatively impacting the model’s performance.

(2) Among three conversation compression strategies, the Rewrite with SFT exhibits superior performance, which could be attributed to the model’s enhanced capability to learn from the simplified question-answer pattern. Intriguingly, although the LLM Summarization strategy demonstrates weaker performance in response generation, it significantly enhances citation labeling. A possible explanation is that the summarization process effectively filters noise, thereby optimizing the content for generating more reliable responses.

5.3 Scaling Analysis on Model Parameters

We scale the generator’s parameters from 500M to 7B, as shown in Figure 3. We find that:

(1) There is a pronounced improvement in gen-

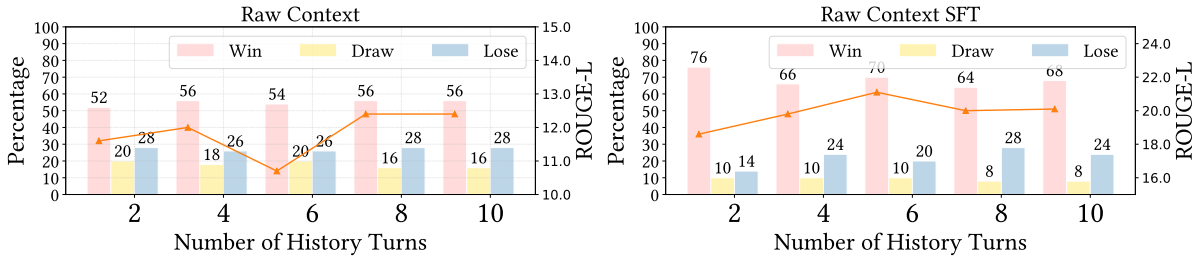


Figure 5: Generation results of different conversation history length. The curve in the figure represents the ROUGE-L score. The histogram shows the results of GPT-4 scores comparing model-generated responses with golden responses. *Win* indicates cases where model-generated responses outperform golden responses, *Draw* indicates cases where the two responses are considered equally good, and *Lose* indicates cases where the golden responses are considered better. The y-axis on the left represents the proportion of cases in the total number of cases.

eration as parameters increase from 500M to 1.5B, evidenced by a significant rise in ROUGE-L scores. However, beyond 3B parameters, the performance gains plateau, indicating diminishing returns with additional parameter scaling.

(2) Performance in citation labeling improves markedly as the parameter count extends from 3B to 7B. This suggests that a larger model capacity is beneficial for tasks that require extensive knowledge, such as accurate citation usage.

5.4 Quantitative Analysis on History Turns

To analyze the impact of conversation history length, we randomly select 50 conversations and vary the number of previous dialogue turns provided to the generator. This can be represented as $r_k^m = \mathbb{G}(q_k; H_k^m; P_k)$, where $k = 12$, $H_k^m = \{q_i, r_i\}_{k-m}^{k-1}$, and $m \in \{2, 4, 6, 8, 10\}$. Results are shown in Figure 5. We find that:

(1) After fine-tuning, the performance improves significantly, especially when using four history turns, resulting in a notable 55% improvement in the ROUGE-L. This demonstrates the effectiveness of SFT in modeling history.

(2) Before fine-tuning, response quality decreases with six turns of history compared to four, possibly due to the redundant information introduced by the longer history. However, after fine-tuning, response quality improves with six turns but declines with eight, suggesting a trade-off between richer information enriched by longer context and irrelevant information introduced by conversation history. These findings validate the challenges previously discussed in Section 1.

6 Conclusion

In this paper, we present an automatic approach using LLMs to construct large-scale, information-seeking conversations from Wikipedia pages. The resulting benchmark, CORAL, supports three fundamental tasks for evaluating conversational RAG systems. Additionally, we propose a unified framework to standardize various conversational RAG methods and conduct a comprehensive evaluation of these methods on CORAL. We envision CORAL as a valuable resource for advancing research in conversational RAG, fostering innovation, and improving real-world applications.

Limitations

Our work presents a conversational RAG benchmark named CORAL, which fills a notable void in assessing conversational RAG methods. In this benchmark, we examine the effects of compressing conversational history on answer generation, paving the way for future research in conversational RAG. However, since CORAL is built upon Wikipedia and existing LLMs are typically trained on corpora like Wikipedia and CommonCrawl, using these LLMs as generators could lead to contamination in the conversational RAG process due to the overlap in their training data. Additionally, the three conversation compression strategies employed in CORAL are somewhat basic, focusing solely on reducing the length of inputs rather than modeling the conversation history in a granular manner. Additionally, the use of the LLM Summarization strategy for compressing both conversation history and retrieved passages, while leveraging advanced models such as GPT-4, could lead to considerable expenses.

Acknowledgments

Zhicheng Dou is the corresponding author. This work was supported by National Natural Science Foundation of China No. 62272467, Beijing Natural Science Foundation No. L233008, Beijing Municipal Science and Technology Project No. Z231100010323009, the fund for building world-class universities (disciplines) of Renmin University of China. The work was partially done at the Engineering Research Center of Next-Generation Intelligent Search and Recommendation, MOE.

References

- Vaibhav Adlakha, Shehzaad Dhuliawala, Kaheer Suleman, Harm de Vries, and Siva Reddy. 2022. Topicqa: Open-domain conversational question answering with topic switching. *Transactions of the Association for Computational Linguistics*, 10:468–483.
- Moonshot AI. 2023. [Kimi chat](#).
- Mohammad Aliannejadi, Zahra Abbasiantaeb, Shubham Chatterjee, Jeffrey Dalton, and Leif Azzopardi. 2024. [TREC ikat 2023: A test collection for evaluating conversational and interactive knowledge assistants](#). In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2024, Washington DC, USA, July 14-18, 2024*, pages 819–829. ACM.
- Raviteja Anantha, Svitlana Vakulenko, Zhucheng Tu, Shayne Longpre, Stephen Pulman, and Srinivas Chappidi. 2021. [Open-domain question answering goes conversational via question rewriting](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2021, Online, June 6-11, 2021*, pages 520–534. Association for Computational Linguistics.
- Anthropic. 2023. [Introducing claude](#).
- Ingeol Baek, Jimin Lee, Joonho Yang, and Hwanhee Lee. 2024. [Crafting the path: Robust query rewriting for information retrieval](#). *CoRR*, abs/2407.12529.
- Haonan Chen, Zhicheng Dou, Kelong Mao, Jiongnan Liu, and Ziliang Zhao. 2024. [Generalizing conversational dense retrieval via llm-cognition data augmentation](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2024, Bangkok, Thailand, August 11-16, 2024*, pages 2700–2718. Association for Computational Linguistics.
- Zhiyu Chen, Bing Liu, Seungwhan Moon, Chinnadhurai Sankar, Paul A. Crook, and William Yang Wang. 2022. [KETOD: knowledge-enriched task-oriented dialogue](#). In *Findings of the Association for Computational Linguistics: NAACL 2022, Seattle, WA, United States, July 10-15, 2022*, pages 2581–2593. Association for Computational Linguistics.
- Zhuyun Dai, Arun Tejasvi Chaganty, Vincent Y. Zhao, Aida Amini, Qazi Mamunur Rashid, Mike Green, and Kelvin Guu. 2022. [Dialog inpainting: Turning documents into dialogs](#). In *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, pages 4558–4586. PMLR.
- Jeffrey Dalton, Chenyan Xiong, and Jamie Callan. 2020a. [Cast 2020: The conversational assistance track overview](#). In *Proceedings of the Twenty-Ninth Text REtrieval Conference, TREC 2020, Virtual Event [Gaithersburg, Maryland, USA], November 16-20, 2020*, volume 1266 of *NIST Special Publication*. National Institute of Standards and Technology (NIST).
- Jeffrey Dalton, Chenyan Xiong, and Jamie Callan. 2020b. [TREC cast 2019: The conversational assistance track overview](#). *CoRR*, abs/2003.13624.
- Jeffrey Dalton, Chenyan Xiong, and Jamie Callan. 2021. [TREC cast 2021: The conversational assistance track overview](#). In *Proceedings of the Thirtieth Text REtrieval Conference, TREC 2021, online, November 15-19, 2021*, volume 500-335 of *NIST Special Publication*. National Institute of Standards and Technology (NIST).
- Emily Dinan, Stephen Roller, Kurt Shuster, Angela Fan, Michael Auli, and Jason Weston. 2019. [Wizard of wikipedia: Knowledge-powered conversational agents](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Yiran Ding, Li Lina Zhang, Chengruidong Zhang, Yuanyuan Xu, Ning Shang, Jiahang Xu, Fan Yang, and Mao Yang. 2024. [Longrope: Extending LLM context window beyond 2 million tokens](#). In *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*. OpenReview.net.
- Guanting Dong, Xiaoshuai Song, Yutao Zhu, Runqi Qiao, Zhicheng Dou, and Ji-Rong Wen. 2024a. [Toward general instruction-following alignment for retrieval-augmented generation](#). *CoRR*, abs/2410.09584.
- Guanting Dong, Yutao Zhu, Chenghao Zhang, Zechen Wang, Zhicheng Dou, and Ji-Rong Wen. 2024b. [Understand what LLM needs: Dual preference alignment for retrieval-augmented generation](#). *CoRR*, abs/2406.18676.
- Guanting Dong, Yutao Zhu, Chenghao Zhang, Zechen Wang, Zhicheng Dou, and Ji-Rong Wen. 2024c. [Understand what LLM needs: Dual preference alignment for retrieval-augmented generation](#). *CoRR*, abs/2406.18676.

- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, Aurélien Rodriguez, Austen Gregerson, Ava Spataru, Baptiste Rozière, Bethany Biron, Binh Tang, Bobbie Chern, Charlotte Caucheteux, Chaya Nayak, Chloe Bi, Chris Marra, Chris McConnell, Christian Keller, Christophe Touret, Chunyang Wu, Corinne Wong, Cristian Canton Ferrer, Cyrus Nikolaidis, Damien Al-lonsius, Daniel Song, Danielle Pintz, Danny Livshits, David Esiobu, Dhruv Choudhary, Dhruv Mahajan, Diego Garcia-Olano, Diego Perino, Dieuwke Hupkes, Egor Lakomkin, Ehab AlBadawy, Elina Lobanova, Emily Dinan, Eric Michael Smith, Filip Radenovic, Frank Zhang, Gabriel Synnaeve, Gabrielle Lee, Georgia Lewis Anderson, Graeme Nail, Grégoire Mialon, Guan Pang, Guillem Cucurell, Hailey Nguyen, Hannah Korevaar, Hu Xu, Hugo Touvron, Iliyan Zarov, Imanol Arrieta Ibarra, Isabel M. Kloumann, Ishan Misra, Ivan Evtimov, Jade Copet, Jaewon Lee, Jan Geffert, Jana Vranes, Jason Park, Jay Mahadeokar, Jeet Shah, Jelmer van der Linde, Jennifer Billock, Jenny Hong, Jenya Lee, Jeremy Fu, Jianfeng Chi, Jianyu Huang, Jiawen Liu, Jie Wang, Jiecao Yu, Joanna Bitton, Joe Spisak, Jongsoo Park, Joseph Rocca, Joshua Johnstun, Joshua Saxe, Junteng Jia, Kalyan Vasuden Alwala, Kartikeya Upasani, Kate Plawiak, Ke Li, Kenneth Heafield, Kevin Stone, and et al. 2024. [The llama 3 herd of models](#). *CoRR*, abs/2407.21783.
- Song Feng, Hui Wan, R. Chulaka Gunasekara, Siva Sankalp Patel, Sachindra Joshi, and Luis A. Lastras. 2020. [doc2dial: A goal-oriented document-grounded dialogue dataset](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 8118–8128. Association for Computational Linguistics.
- Tianyu Gao, Howard Yen, Jiatong Yu, and Danqi Chen. 2023. [Enabling large language models to generate text with citations](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023*, pages 6465–6488. Association for Computational Linguistics.
- Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Ming-Wei Chang. 2020. [REALM: retrieval-augmented language model pre-training](#). *CoRR*, abs/2002.08909.
- Yizheng Huang and Jimmy Huang. 2024. [A survey on retrieval-augmented text generation for large language models](#). *CoRR*, abs/2404.10981.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de Las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, L  lio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timoth  e Lacroix, and William El Sayed. 2023. [Mistral 7b](#). *CoRR*, abs/2310.06825.
- Huiqiang Jiang, Qianhui Wu, Xufang Luo, Dongsheng Li, Chin-Yew Lin, Yuqing Yang, and Lili Qiu. 2024. [Longlmlingua: Accelerating and enhancing llms in long context scenarios via prompt compression](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2024, Bangkok, Thailand, August 11-16, 2024*, pages 1658–1677. Association for Computational Linguistics.
- Jiajie Jin, Yutao Zhu, Yujia Zhou, and Zhicheng Dou. 2024. [BIDER: bridging knowledge inconsistency for efficient retrieval-augmented llms via key supporting evidence](#). In *Findings of the Association for Computational Linguistics, ACL 2024, Bangkok, Thailand and virtual meeting, August 11-16, 2024*, pages 750–761. Association for Computational Linguistics.
- Zhuoran Jin, Pengfei Cao, Yubo Chen, Kang Liu, and Jun Zhao. 2023. [Instructor: Instructing unsupervised conversational dense retrieval with large language models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023, Singapore, December 6-10, 2023*, pages 6649–6675. Association for Computational Linguistics.
- Jeff Johnson, Matthijs Douze, and Herv   J  gou. 2019. Billion-scale similarity search with gpus. *IEEE Transactions on Big Data*, 7(3):535–547.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick S. H. Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. [Dense passage retrieval for open-domain question answering](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 6769–6781. Association for Computational Linguistics.
- Yannis Katsis, Sara Rosenthal, Kshitij Fadnis, Chulaka Gunasekara, Young-Suk Lee, Lucian Popa, Vraj Shah, Huaiyu Zhu, Danish Contractor, and Marina Danilevsky. 2025. [Mtrag: A multi-turn conversational benchmark for evaluating retrieval-augmented generation systems](#). *arXiv preprint arXiv:2501.03468*.
- Vaibhav Kumar and Jamie Callan. 2020. [Making information seeking easier: An improved pipeline for conversational search](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020, Online Event, 16-20 November 2020*, volume EMNLP 2020 of *Findings of ACL*, pages 3971–3980. Association for Computational Linguistics.
- Tzu-Lin Kuo, Feng-Ting Liao, Mu-Wei Hsieh, Fu-Chieh Chang, Po-Chun Hsu, and Da-Shan Shiu. 2024. [Rad-bench: Evaluating large language models capabilities in retrieval augmented dialogues](#). *CoRR*, abs/2409.12558.

- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph Gonzalez, Hao Zhang, and Ion Stoica. 2023. [Efficient memory management for large language model serving with pagedattention](#). In *Proceedings of the 29th Symposium on Operating Systems Principles, SOSP 2023, Koblenz, Germany, October 23-26, 2023*, pages 611–626. ACM.
- 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. [Retrieval-augmented generation for knowledge-intensive NLP tasks](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Minghan Li, Honglei Zhuang, Kai Hui, Zhen Qin, Jimmy Lin, Rolf Jagerman, Xuanhui Wang, and Michael Bendersky. 2024a. [Can query expansion improve generalization of strong cross-encoder rankers?](#) In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2024, Washington DC, USA, July 14-18, 2024*, pages 2321–2326. ACM.
- Xiaoxi Li, Guanting Dong, Jiajie Jin, Yuyao Zhang, Yujia Zhou, Yutao Zhu, Peitian Zhang, and Zhicheng Dou. 2025. [Search-o1: Agentic search-enhanced large reasoning models](#). *arXiv preprint arXiv:2501.05366*.
- Xiaoxi Li, Jiajie Jin, Yujia Zhou, Yongkang Wu, Zhonghua Li, Qi Ye, and Zhicheng Dou. 2024b. [Retrollm: Empowering large language models to retrieve fine-grained evidence within generation](#). *CoRR*, abs/2412.11919.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.
- Sheng-Chieh Lin, Jheng-Hong Yang, and Jimmy Lin. 2021a. [Contextualized query embeddings for conversational search](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021*, pages 1004–1015. Association for Computational Linguistics.
- Sheng-Chieh Lin, Jheng-Hong Yang, Rodrigo Nogueira, Ming-Feng Tsai, Chuan-Ju Wang, and Jimmy Lin. 2021b. [Multi-stage conversational passage retrieval: An approach to fusing term importance estimation and neural query rewriting](#). *ACM Transactions on Information Systems (TOIS)*, 39(4):1–29.
- Sheng-Chieh Lin, Jheng-Hong Yang, Rodrigo Frassetto Nogueira, Ming-Feng Tsai, Chuan-Ju Wang, and Jimmy Lin. 2020. [Conversational question reformulation via sequence-to-sequence architectures and pretrained language models](#). *CoRR*, abs/2004.01909.
- Yubo Ma, Yixin Cao, Yong Hong, and Aixin Sun. 2023. [Large language model is not a good few-shot information extractor, but a good reranker for hard samples!](#) In *Findings of the Association for Computational Linguistics: EMNLP 2023, Singapore, December 6-10, 2023*, pages 10572–10601. Association for Computational Linguistics.
- Kelong Mao, Chenlong Deng, Haonan Chen, Fengran Mo, Zheng Liu, Tetsuya Sakai, and Zhicheng Dou. 2024. [Chatretriever: Adapting large language models for generalized and robust conversational dense retrieval](#). *CoRR*, abs/2404.13556.
- Kelong Mao, Zhicheng Dou, Fengran Mo, Jiewen Hou, Haonan Chen, and Hongjin Qian. 2023a. [Large language models know your contextual search intent: A prompting framework for conversational search](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023, Singapore, December 6-10, 2023*, pages 1211–1225. Association for Computational Linguistics.
- Kelong Mao, Zhicheng Dou, and Hongjin Qian. 2022a. [Curriculum contrastive context denoising for few-shot conversational dense retrieval](#). In *SIGIR '22: The 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, Madrid, Spain, July 11 - 15, 2022*, pages 176–186. ACM.
- Kelong Mao, Zhicheng Dou, Hongjin Qian, Fengran Mo, Xiaohua Cheng, and Zhao Cao. 2022b. [Convtans: Transforming web search sessions for conversational dense retrieval](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022*, pages 2935–2946. Association for Computational Linguistics.
- Kelong Mao, Hongjin Qian, Fengran Mo, Zhicheng Dou, Bang Liu, Xiaohua Cheng, and Zhao Cao. 2023b. [Learning denoised and interpretable session representation for conversational search](#). In *Proceedings of the ACM Web Conference 2023, WWW 2023, Austin, TX, USA, 30 April 2023 - 4 May 2023*, pages 3193–3202. ACM.
- Fengran Mo, Kelong Mao, Ziliang Zhao, Hongjin Qian, Haonan Chen, Yiruo Cheng, Xiaoxi Li, Yutao Zhu, Zhicheng Dou, and Jian-Yun Nie. 2024a. [A survey of conversational search](#). *arXiv preprint arXiv:2410.15576*.
- Fengran Mo, Jian-Yun Nie, Kaiyu Huang, Kelong Mao, Yutao Zhu, Peng Li, and Yang Liu. 2023. [Learning to relate to previous turns in conversational search](#). In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, KDD 2023, Long Beach, CA, USA, August 6-10, 2023*, pages 1722–1732. ACM.
- Fengran Mo, Chen Qu, Kelong Mao, Tianyu Zhu, Zhan Su, Kaiyu Huang, and Jian-Yun Nie. 2024b. [History-aware conversational dense retrieval](#). In *Findings of*

- the Association for Computational Linguistics, ACL 2024, Bangkok, Thailand and virtual meeting, August 11-16, 2024*, pages 13366–13378. Association for Computational Linguistics.
- Fengran Mo, Bole Yi, Kelong Mao, Chen Qu, Kaiyu Huang, and Jian-Yun Nie. 2024c. [Convsgd: Session data generation for conversational search](#). In *Companion Proceedings of the ACM on Web Conference 2024, WWW 2024, Singapore, Singapore, May 13-17, 2024*, pages 1634–1642. ACM.
- OpenAI. 2022. [Openai: Introducing chatgpt](#).
- Paul Owoicho, Jeff Dalton, Mohammad Aliannejadi, Leif Azzopardi, Johanne R. Trippas, and Svitlana Vakulenko. 2022. [TREC cast 2022: Going beyond user ask and system retrieve with initiative and response generation](#). In *Proceedings of the Thirty-First Text REtrieval Conference, TREC 2022, online, November 15-19, 2022*, volume 500-338 of *NIST Special Publication*. National Institute of Standards and Technology (NIST).
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics, July 6-12, 2002, Philadelphia, PA, USA*, pages 311–318. ACL.
- Hongjing Qian, Yutao Zhu, Zhicheng Dou, Haoqi Gu, Xinyu Zhang, Zheng Liu, Ruofei Lai, Zhao Cao, Jian-Yun Nie, and Ji-Rong Wen. 2023. [Webbrain: Learning to generate factually correct articles for queries by grounding on large web corpus](#). *CoRR*, abs/2304.04358.
- Chen Qu, Liu Yang, Cen Chen, Minghui Qiu, W. Bruce Croft, and Mohit Iyyer. 2020. [Open-retrieval conversational question answering](#). In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval, SIGIR 2020, Virtual Event, China, July 25-30, 2020*, pages 539–548. ACM.
- Nir Ratner, Yoav Levine, Yonatan Belinkov, Ori Ram, Inbal Magar, Omri Abend, Ehud Karpas, Amnon Shashua, Kevin Leyton-Brown, and Yoav Shoham. 2023. [Parallel context windows for large language models](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2023, Toronto, Canada, July 9-14, 2023*, pages 6383–6402. Association for Computational Linguistics.
- Siva Reddy, Danqi Chen, and Christopher D Manning. 2019. [Coqa: A conversational question answering challenge](#). *Transactions of the Association for Computational Linguistics*, 7:249–266.
- Jiejun Tan, Zhicheng Dou, Wen Wang, Mang Wang, Weipeng Chen, and Ji-Rong Wen. 2024a. [Htmrlrag: HTML is better than plain text for modeling retrieved knowledge in RAG systems](#). *CoRR*, abs/2411.02959.
- Jiejun Tan, Zhicheng Dou, Yutao Zhu, Peidong Guo, Kun Fang, and Ji-Rong Wen. 2024b. [Small models, big insights: Leveraging slim proxy models to decide when and what to retrieve for llms](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2024, Bangkok, Thailand, August 11-16, 2024*, pages 4420–4436. Association for Computational Linguistics.
- ByteDance Doubao Team. 2023. [Doubao](#).
- Qwen Team. 2024. [Qwen2.5: A party of foundation models](#).
- Svitlana Vakulenko, Shayne Longpre, Zhucheng Tu, and Raviteja Anantha. 2021. [Question rewriting for conversational question answering](#). In *WSDM '21, The Fourteenth ACM International Conference on Web Search and Data Mining, Virtual Event, Israel, March 8-12, 2021*, pages 355–363. ACM.
- Nikos Voskarides, Dan Li, Pengjie Ren, Evangelos Kanoulas, and Maarten de Rijke. 2020. [Query resolution for conversational search with limited supervision](#). In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval, SIGIR 2020, Virtual Event, China, July 25-30, 2020*, pages 921–930. ACM.
- Liang Wang, Nan Yang, and Furu Wei. 2023a. [Query2doc: Query expansion with large language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023*, pages 9414–9423. Association for Computational Linguistics.
- Shuting Wang, Xin Yu, Mang Wang, Weipeng Chen, Yutao Zhu, and Zhicheng Dou. 2024. [Richrag: Crafting rich responses for multi-faceted queries in retrieval-augmented generation](#). *CoRR*, abs/2406.12566.
- Zhiruo Wang, Jun Araki, Zhengbao Jiang, Md. Rizwan Parvez, and Graham Neubig. 2023b. [Learning to filter context for retrieval-augmented generation](#). *CoRR*, abs/2311.08377.
- Zequ Wu, Yi Luan, Hannah Rashkin, David Reiter, Hannaneh Hajishirzi, Mari Ostendorf, and Gaurav Singh Tomar. 2022. [CONQRR: conversational query rewriting for retrieval with reinforcement learning](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022*, pages 10000–10014. Association for Computational Linguistics.
- Zequ Wu, Ryu Parish, Hao Cheng, Sewon Min, Prithviraj Ammanabrolu, Mari Ostendorf, and Hannaneh Hajishirzi. 2023. [Inscit: Information-seeking conversations with mixed-initiative interactions](#). *Transactions of the Association for Computational Linguistics*, 11:453–468.

- Fangyuan Xu, Weijia Shi, and Eunsol Choi. 2023. [RECOMP: improving retrieval-augmented lms with compression and selective augmentation](#). *CoRR*, abs/2310.04408.
- Shicheng Xu, Liang Pang, Jun Xu, Huawei Shen, and Xueqi Cheng. 2024. [List-aware reranking-truncation joint model for search and retrieval-augmented generation](#). In *Proceedings of the ACM on Web Conference 2024, WWW 2024, Singapore, May 13-17, 2024*, pages 1330–1340. ACM.
- An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, et al. 2024. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*.
- Haoyan Yang, Zhitao Li, Yong Zhang, Jianzong Wang, Ning Cheng, Ming Li, and Jing Xiao. 2023. [PRCA: fitting black-box large language models for retrieval question answering via pluggable reward-driven contextual adapter](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023*, pages 5364–5375. Association for Computational Linguistics.
- Linhao Ye, Zhikai Lei, Jianghao Yin, Qin Chen, Jie Zhou, and Liang He. 2024. [Boosting conversational question answering with fine-grained retrieval-augmentation and self-check](#). In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2024, Washington DC, USA, July 14-18, 2024*, pages 2301–2305. ACM.
- Shi Yu, Jiahua Liu, Jingqin Yang, Chenyan Xiong, Paul N. Bennett, Jianfeng Gao, and Zhiyuan Liu. 2020. [Few-shot generative conversational query rewriting](#). In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval, SIGIR 2020, Virtual Event, China, July 25-30, 2020*, pages 1933–1936. ACM.
- Shi Yu, Zhenghao Liu, Chenyan Xiong, Tao Feng, and Zhiyuan Liu. 2021. [Few-shot conversational dense retrieval](#). In *SIGIR '21: The 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, Virtual Event, Canada, July 11-15, 2021*, pages 829–838. ACM.
- Yaowei Zheng, Richong Zhang, Junhao Zhang, Yanhan Ye, Zheyang Luo, Zhangchi Feng, and Yongqiang Ma. 2024. [Llamafactory: Unified efficient fine-tuning of 100+ language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 3: System Demonstrations)*, Bangkok, Thailand. Association for Computational Linguistics.
- Yutao Zhu, Zhaoheng Huang, Zhicheng Dou, and Jirong Wen. 2024. [One token can help! learning scalable and pluggable virtual tokens for retrieval-augmented large language models](#). *CoRR*, abs/2405.19670.

Appendix

A Prompts of the Contextualization of Questions

When contextualizing questions, two steps need prompt. Firstly, we transform the node into the complete question. Secondly, we convert the complete questions into conversational questions. Table 6 illustrates the prompt for generating a complete question. Table 7 demonstrates the prompt for creating conversational questions. Following LLM4CS (Mao et al., 2023a), the prompt consists of three components: Instruction, Demonstration, and Input. The red section is designated for LDS prompting, the blue section for SIDS and STRW prompting, the green section for DTRW prompting, and the orange section for LDS, SIDS, and STRW prompting.

B Prompts of Generating Responses with Citation Labeling

Table 6 provides the prompt template for generating a response with citation labeling. The red part is for Raw Context and the Last Response strategy prompting. The blue part is for Rewrite Strategy prompting. The green part is for LLM Summarization Strategy prompting. The orange part is for Raw Context, the Last Response Strategy, and the LLM Summarization Strategy prompting.

C Prompts of LLM Summarization Strategy

Table 8 provides a general illustration of the prompt of generating a summary of conversation history.

D More Detailed Experimental Setting

D.1 Conversational Search Baselines

The conversational search baseline models are chosen for their prevalence and effectiveness in the field. We focus on two primary approaches: conversational dense retrieval (CDR) and conversational query rewriting (CQR). For CDR, we adopt KD-ANCE and Conv-ANCE, where ANCE is a base ad-hoc retriever. Following Yu et al. (2021), KD-ANCE uses an ad hoc query encoder as the teacher model, training the student session encoder to imitate the embeddings derived from the golden queries. Meanwhile, according to the methodology outlined by (Karpukhin et al., 2020; Lin et al., 2021a), Conv-ANCE is designed to implement

the classical ranking loss function. This function strives to minimize the distance between the session and its relevant passages while maximizing the separation from irrelevant ones. Dense retrieval is conducted using Faiss (Johnson et al., 2019). For CQR, we choose LLM4CS (Mao et al., 2023a), employing the proprietary commercial model GPT-3.5 to generate rewrites. Additionally, we choose an open-source LLM to generate rewrites as well, allowing for a comparative analysis.

D.2 Generation with Citation Labeling

We compare the raw context baselines with another three conversation compression strategies introduced in Section 4. We choose Qwen2.5-7B-Instruct (Yang et al., 2024; Team, 2024), Mistral-7B-Instruct (Jiang et al., 2023), and Llama-3.1-8B-Instruct (Dubey et al., 2024) as the generator. For the scaling analysis, we use the Qwen2.5-Instruct series, specifically the 0.5B, 1.5B, 3B, and 7B models, as our generators for detailed examination. During the training process, we utilize the LLaMA-Factory (Zheng et al., 2024) framework, running on two A800 GPUs. The training parameters are set as follows: we employ a learning rate of $1.0e-5$. The batch size is maintained at 1, and the maximum token length for training instances is set to 4096. Because of the lack of training data of the LLM Summarization category, we use the checkpoint of Raw Context.

During the inference process, we leverage the vLLM (Kwon et al., 2023) framework to accelerate inference. The maximum input length is set to 32,000, top_p is set to 0.9, and temperature is set to 1.

D.3 More detailed Scaling Analysis

Table 5 provides detailed results of the generation quality and citation labeling accuracy.

E Dataset Format

Table 10 provides an example of CORAL. Our dataset CORAL has information-seeking questions, free-form responses with citation labeling, golden rewrites, and corresponding golden retrieval passage ID.

Category	Model	Generation		Citation Labeling	
		BLEU-1	ROUGE-L	Citation Recall	Citation Precision
Raw Context	Qwen2.5-0.5B	16.4	11.8	0.1	0.2
	Qwen2.5-1.5B	20.8	12.9	0.3	1.2
	Qwen2.5-3B	21.4	12.9	0.4	1.8
	Qwen2.5-7B	22.2	13.1	3.1	18.1
	Qwen2.5-0.5B-SFT	13.0	15.5	3.5	10.2
	Qwen2.5-1.5B-SFT	20.9	17.5	3.6	10.0
	Qwen2.5-3B-SFT	25.8	17.6	8.1	20.7
	Qwen2.5-7B-SFT	18.3	18.5	6.6	16.8
Last Response	Qwen2.5-0.5B	15.6	11.7	0.2	0.5
	Qwen2.5-1.5B	19.3	12.5	0.8	2.9
	Qwen2.5-3B	21.1	12.6	0.6	3.4
	Qwen2.5-7B	20.9	12.8	3.5	20.8
	Qwen2.5-0.5B-SFT	19.8	15.5	6.5	18.0
	Qwen2.5-1.5B-SFT	19.4	15.8	6.7	17.7
	Qwen2.5-3B-SFT	21.8	15.8	7.4	17.1
	Qwen2.5-7B-SFT	22.1	16.5	10.4	24.8
Rewrite	Qwen2.5-0.5B	17.3	12.0	0.4	0.8
	Qwen2.5-1.5B	19.9	12.5	0.4	1.2
	Qwen2.5-3B	20.8	12.5	4.0	14.9
	Qwen2.5-7B	21.1	12.8	2.4	9.5
	Qwen2.5-0.5B-SFT	21.4	16.1	6.5	16.5
	Qwen2.5-1.5B-SFT	21.7	16.6	5.9	14.9
	Qwen2.5-3B-SFT	23.3	16.8	7.2	16.5
	Qwen2.5-7B-SFT	18.9	16.4	7.4	19.8
LLM Summarization	Qwen2.5-0.5B	13.2	11.6	0.2	0.4
	Qwen2.5-1.5B	15.0	12.1	0.9	3.0
	Qwen2.5-3B	20.2	12.7	2.5	10.6
	Qwen2.5-7B	21.0	12.7	2.9	13.0
	Qwen2.5-0.5B-SFT	21.4	14.6	6.5	17.4
	Qwen2.5-1.5B-SFT	23.0	15.5	6.8	16.9
	Qwen2.5-3B-SFT	17.6	15.6	14.6	36.0
	Qwen2.5-7B-SFT	23.5	16.8	14.1	31.1

Table 5: The complete scaling analysis of generation and citation labeling performance. The best performance is **bold**.

Instruction	<p>Given the keyword chain of the question and its response, generate the original question. If the response is not informative enough to help you reconstruct the question, please rely on the provided keyword chain to generate the question. The keyword chain consists of terms where each term is a more specific or detailed subset of the previous one, with the last term being the most specific or important. The question you generate should focus on the last keyword in the chain and include it explicitly.</p>
Input	<p>Given the following keyword chain and response: Keyword Chain: 72nd Primetime Emmy Awards, ceremony information, category and rule changes Response: Several rule changes were announced in December 2019. first, episodes that were scheduled to air after the eligibility period closed... (Now, you should give me the original question given the keyword chain and its response. The output format should always be: Question: \$Question. Note that you should always try to generate it. Never ask for clarification or say you don't understand it in the generated question. Go ahead!)</p>
Model Output	<p>Question: What were the category and rule changes for the 72nd Primetime Emmy Awards ceremony?</p>

Table 6: An illustration of the prompt for question generation. The prompt consists of two parts: Instruction and Input.

Instruction

Given a topic and corresponding question and response pairs. **The questions are arranged in a logical, progressively deeper sequence, where each subsequent question delves deeper into the topic based on the previous one.** The questions are organized in a logical sequence where they are interconnected and may delve deeper into earlier topics, rather than following a direct, linear progression. **The questions initially follow a logical progression but may shift to another topic as needed, reflecting a dynamic conversational flow rather than a strict linear order.** I would like you to convert the original question into a conversational form. The goal is to rewrite it without any grammatical errors while preserving its original intent as closely as possible. It is necessary to consider the omission and reference to the previous question and response in the generated conversational question.

Demonstration

I will give you one example multi-turn dialog, where each turn contains an original question, a conversational question, a response, and the corresponding analysis.

Example:**Topic:** depression**Conversations:****Turn #1:****Original Question1:** What are the mechanisms of depression?**Conversational Question1:** What are the mechanisms of depression?**Response1:** The major neurotransmitters are acetylcholine, norepinephrine, dopamine, and serotonin. Many experts believe that an imbalance among the different neurotransmitters is the cause of depression.**Analysis:** The initial question addresses the mechanisms of depression, which leads to an explanation of neurotransmitters and their role in depression.**Turn #2:****Original Question2:** What is the role of serotonin in depression?**Conversational Question2:** What is the role of serotonin?**Response2:** When it comes to biological factors for depression, it is commonly accepted that serotonin plays a crucial role. In the brain, this neurotransmitter acts as a relay, transporting messages from one area to another. Given this critical task, an imbalance in serotonin levels may influence our emotions.**Analysis:** Since "depression" was already mentioned in the previous turn, it's redundant to repeat it in the subsequent question. Therefore, by omitting "in depression", the question flows more smoothly and maintains coherence in the conversation.**Turn #3:****Original Question3:** How can I balance serotonin?**Conversational Question3:** So how can I balance it?**Response3:** You can create your own dopamine rewards by setting a goal, and then take small steps toward it every day. This can be starting a new exercise program, learning French, or challenging yourself to drive home from work a different way every day, preferably without the use of your GPS.**Analysis:** The use of "it" in the question effectively refers back to "serotonin" from the question in the previous turn. This helps maintain coherence in the conversation and avoids repetitive expressions. And we use "so" as a connecting word to make the conversation look smoother.**Turn #4:****Original Question4:** What foods boost dopamine?

Conversational Question4: What foods boost it?

Response4: Bananas, especially ripe bananas, are an exceptional food for regulating dopamine because they have a high concentration of tyrosine. Other foods that increase dopamine through the conversion of phenylalanine to tyrosine include almonds, apples, watermelons, cherries, yogurt, beans, eggs, and meats.

Analysis: Continuing the discussion on depression, the question asks about foods that boost dopamine, with "it" referring back to "dopamine" from the previous response.

...

Now I will give you a topic and the corresponding question and response pairs:

Now I will give you the question and response pairs:

Topic: Solve Education

Id: 1

Question: What is Solve Education and what are its main objectives?

Response: solve education is a global technology non-profit organization founded in 2015 by peng tsin ong and Janine teo...

Id: 2

Question: What is the history of the organization Solve Education! since its inception?

Response: solve education was founded in December 2015 after peng tsin ong attended the Africa Innovation Summit in Cape Verde in 2014 where he led a plenary session on financing innovation...

Id: 3

Question: What is the purpose of the Dawn of Civilization game developed by Solve Education, and how does it utilize its services to benefit education in low-income countries?

Response: Dawn of Civilization is a city-building game which is based on various subjects and itself contains mini-games...

Id: 4

Question: What is Ed the Learning Bot and how does it contribute to the services provided by Solve Education?

Response: It is an AI-based chatbot used by solve education! on telegram. it helps underserved young people to learn about any subject with the help of gamified pictures and text...

Id: 5

Question: What is the purpose of the Content+ service on the Solve Education platform, and how do educational content providers interact with it?

Response: it is an open platform which facilitates educational content providers to add their content in order to meet the needs of beneficiaries...

Id: 6

Question: What is Learnalytics and how does it function within Solve Education services?

Response: It is a tool developed by solving education! which helps measure the performance. additionally, learning portal analytics, provides context, the learning progress, and the performance of the students as they learn...

...

Input

	<p>(Please consider the question and response from the previous text when generating the current conversational question, but there is no need to generate the response. The output should be in the following format: Turn #$\\$turn_id$ Original Question: $\\$original\ question$ Conversational Question: $\\$conversational\ question$ Reason: $\\$reason$)</p>
Model Output	<p>Turn #1 Original Question: What is Solve Education! and what are its main objectives? Conversational Question: What is Solve Education and what are its main objectives? Reason: This is the first question in the series, so it introduces the topic and organization directly, laying the foundation for the conversation. No reference to prior questions is needed.</p>
	<p>Turn #2 Original Question: What is the history of the organization Solve Education since its inception? Conversational Question: Could you share more about the history of Solve Education since it started? Reason: The conversational question transitions smoothly from the introduction of the organization by asking for historical details. The phrase "since it started" connects directly back to the organization introduced in Turn #1.</p>
	<p>Turn #3 Original Question: What is the purpose of the Dawn of Civilization game developed by Solve Education, and how does it utilize its services to benefit education in low-income countries? Conversational Question: Can you tell me about the Dawn of Civilization game and how it benefits education in low-income areas? Reason: This question builds on the previous introduction of Solve Education's initiatives. The phrase "the Dawn of Civilization game" is specific and follows logically from learning about the organization's history and broader objectives.</p>
	<p>Turn #4 Original Question: What is Ed the Learning Bot and how does it contribute to the services provided by Solve Education? Conversational Question: What role does Ed the Learning Bot play in Solve Education's services? Reason: The question uses "Ed the Learning Bot" directly and asks about its role, making it clear and focused on this specific service. It follows naturally after discussing another tool (Dawn of Civilization game) used by the organization.</p>
	<p>...</p>

Table 7: An illustration of the prompt for generating conversational questions consists of three parts: Instruction, Demonstration, and Input. The red section is designated for LDS prompting, the blue section for SIDS and STRW prompting, the green section for DTRW prompting, and the orange section for LDS, SIDS, and STRW prompting.

Instruction	Given the current question and the previous conversation history, summarize the conversation history. The summary should contain relevant information to help the agent provide a more informative response to the current question.
Demonstration	<p>I will give you several example dialogs, where each example contains the current question, the conversation history, the summarization, and the reason of generating such summarization.</p> <p>Example #1: Current Question: What about their impact on battery longevity? Question1: What are some of the key advancements in electric vehicle technology lately? Response1: There are improvements in battery technology for longer ranges, faster charging methods, and the integration of solar panels to help extend range. Question2: Do any particular models feature these solar panels? Response2: Yes, a number of new sedans and SUVs have solar roofs which can significantly increase daily driving range. Question3: What about the charging stations? Are they getting better too? Response3: Indeed, ultra-fast charging stations are now available that can boost a battery to 80% in as little as 20 minutes. Summarization: The dialogue has highlighted advancements in electric vehicle technology, especially in battery efficiency, solar panel integration, and faster charging solutions. The conversation has particularly focused on the introduction of ultra-fast charging stations and briefly touched upon their efficiency in reducing charging times. Reason: The summary efficiently condenses prior discussions to provide context specifically for the user’s current question about the impact of ultra-fast charging on battery longevity. ... </p>
Input	<p>Given the following question and its conversation history:</p> <p>Current Question: What are some other applications of this method, particularly in solving constrained optimization problems?</p> <p>Conversation History: Question1: Can you explain what the multiplicative weight update method is and how it’s used in decision making and prediction? Response1: The multiplicative weights update method is an algorithmic technique most commonly used for decision making and prediction. . . (Now, you should give me the summarization given the current question and its conversation history. The output format should always be: “Summarization: \$Summarization”. Note that you should always try to generate it. Never ask for clarification or say you can’t generate the summarization. Go ahead!)</p>
Model Output	<p>Summarization: The conversation has detailed the multiplicative weight update method, describing its primary use in decision making and prediction, particularly through the example of prediction from expert advice. It also touched on its wide applications in fields like machine learning, optimization, theoretical computer science, and game theory, which relates to its potential in solving constrained optimization problems as queried in the current question.</p>

Table 8: An illustration of the prompt for generating a conversation summary. The prompt consists of three components: Instruction, Demonstration, and Input.

Instruction	Based on the provided passages, please answer the current question while considering the context of the conversations history. If the current question lacks complete information or contains omissions, use the conversation history to enhance understanding and provide a response. please answer the question. please answer the current question considering the summarization of the conversation history please. If the current question lacks complete information or contains omissions, use the summarization of conversation history to enhance understanding. Support your answer with citations in brackets. If the provided passages are not informative, please generate the answer based on your own knowledge. Don't say you can't answer the question.
Input	Current Question: ... Question: ... Conversation History: ... Summarization of Conversation History: ... Passages: ...

Table 9: The prompt template of generating the response. The prompt consists of two parts, i.e., Instruction, and Input. **red part** is for Raw Context and Last Response strategy prompting. The **blue part** is for Rewrite Strategy prompting. The **green part** is for LLM Summarization Strategy prompting. The **orange part** is for Raw Context, Last Response Strategy, and LLM Summarization Strategy prompting.

Question1: What were the key details and outcomes of the 72nd Primetime Emmy Awards?

Response1: The 72nd Primetime Emmy Awards honored the best in American prime time television programming from June 1, 2019, until May 31, 2020, as chosen by the Academy of Television Arts & Sciences.[65215]. . .

Golden Retrieval Passages IDs: 65215

Golden Rewrite: What were the key details and outcomes of the 72nd Primetime Emmy Awards?

URL: https://en.wikipedia.org/wiki/72nd_Primetime_Emy_Awards

Question2: Could you tell me about the winners and nominees from this event?

Response2: The nominations for the 72nd Primetime Emmy Awards were announced on July 28, 2020, by host Leslie Jones and presenters Laverne Cox, Josh Gad, and Tatiana Maslany.[65226]. . .

Golden Retrieval Passage IDs: 65216, 65217, 65218, 65219, 65220, 65221, 65222, 65223, 65224, 65225, 65226, 65227, 65228, 65229, 65230, 65231, 65232

Golden Rewrite: Who were the winners and nominees at the 72nd Primetime Emmy Awards?

URL: https://en.wikipedia.org/wiki/72nd_Primetime_Emy_Awards

Question3: Who received the Governors Award and why were they chosen?

Response3: The Governors Award was presented to Tyler Perry and the Perry Foundation “in recognition of their unparalleled contributions to shaping the television medium” and for their “inclusion, engagement, employment and other philanthropic initiatives”. [65233][65234] The award was moved to the main telecast from its usual presentation at the Creative Arts Emmys.[65235]

Golden Retrieval Passage IDs: 65233, 65234, 65235

Golden Rewrite: Who received the Governors Award at the 72nd Primetime Emmy Awards and for what reasons?

URL: https://en.wikipedia.org/wiki/72nd_Primetime_Emy_Awards

...

...

...

Question10: Who performed during the in Memoriam Segment, and which notable figures were mentioned?

Response10: The annual in Memoriam Segment featured H.E.R. performing “Nothing Compares 2 U” on piano and electric guitar .[65284][65285]. . .

Golden Retrieval Passage IDs: 65284, 65285, 65286, 65287

Golden Rewrite: Who performed during the in Memoriam Segment at the 72nd Primetime Emmy Awards, and which notable figures were mentioned?

URL: https://en.wikipedia.org/wiki/72nd_Primetime_Emy_Awards

Table 10: A conversation from CORAL.