

CHIQ: Contextual History Enhancement for Improving Query Rewriting in Conversational Search

Fengran Mo¹, Abbas Ghaddar², Kelong Mao³,
Mehdi Rezagholizadeh², Boxing Chen², Qun Liu², Jian-Yun Nie¹

¹DIRO, Université de Montréal, Québec, Canada

²Huawei Noah's Ark Lab, Montreal Research Center, Canada

³Renmin University of China

fengran.mo@umontreal.ca, nie@iro.umontreal.ca

Abstract

In this paper, we study how open-source large language models (LLMs) can be effectively deployed for improving query rewriting in conversational search, especially for ambiguous queries. We introduce CHIQ, a two-step method that leverages the capabilities of LLMs to resolve ambiguities in the conversation history before query rewriting. This approach contrasts with prior studies that predominantly use closed-source LLMs to directly generate search queries from conversation history. We demonstrate on five well-established benchmarks that CHIQ leads to state-of-the-art results across most settings, showing highly competitive performances with systems leveraging closed-source LLMs. Our study provides a first step towards leveraging open-source LLMs in conversational search, as a competitive alternative to the prevailing reliance on commercial LLMs for query rewriting. Our code is publicly available at <https://github.com/fengranMark/CHIQ>.

1 Introduction

Conversational search enables users to interact with the system in a multi-turn fashion to satisfy their complex information needs (Gao et al., 2022; Zamani et al., 2023). One of the crucial steps is to compose adequate search queries for each context-dependent utterance. Recent advancements in the task-solving capabilities of Large Language Models (LLMs) (Ouyang et al., 2022; Chen et al., 2024b; Wang et al., 2024a; Huang et al., 2024) have motivated researchers to integrate these models into existing conversational search systems.

Most recent studies (Mao et al., 2023a; Ye et al., 2023) leverage LLMs to directly generate search queries based on the context of the conversation history. Although seemingly straightforward, this technique is shown to achieve higher effectiveness in query rewriting than fine-tuning a smaller language model, such as T5 (Raffel et al., 2020; Chung

et al., 2022). However, these performance gains are primarily achieved through the use of commercial, closed-source LLMs (OpenAI, 2023). This is primarily because closed-source LLMs can better perform complex reasoning tasks (Gudibande et al., 2023; Kaddour et al., 2023) compared to open-source models.

One of the main challenges in conversational search resides in the ambiguous nature of the conversation history. Figure 1 illustrates an example where solving co-reference relation in u_4 and elaborating the response in r_3 can help generate an adequate search query. Intuitively, performing these tasks requires basic NLP task-solving capabilities, which even small-scale open-source LLMs (e.g., 7B) possess (Touvron et al., 2023; Jiang et al., 2023). The key challenge is to unlock the full capabilities of open-source LLMs for conversational search. This requires carefully preparing the conversation history to enhance its quality, rather than directly using it to generate the search query.

In this paper, we propose **CHIQ**, a method that aims to enhance the quality of contextual history for improving query rewriting. As illustrated in Figure 1, we leverage the NLP capabilities of LLMs (e.g. solving coreference relation or expanding the context) to make the conversational history less ambiguous, consequently enhancing the relevance of the generated search query. We investigate various methods for integrating refined conversational history into existing frameworks, including ad-hoc query rewriting, generating pseudo supervision signals for fine-tuning query rewriting models, and the fusion of both approaches.

We conduct extensive experiments using the open-source LLM, LLaMA-2-7B (Touvron et al., 2023), across five well-established conversational search benchmarks under both dense and sparse retrieval settings. The experimental results indicate that enhancing the conversational history using our method achieves state-of-the-art performance

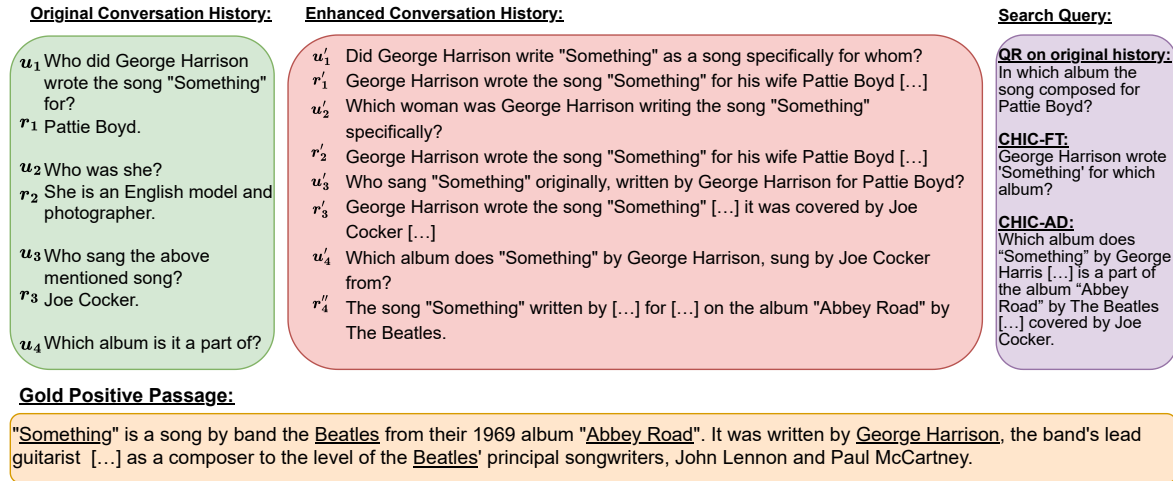


Figure 1: An illustrative example of a conversational history (left box) and the gold positive passage relevant to the last user turn. The enhanced history obtained using our method described in § 3.2 is in the middle box. The right box shows the three search queries generated by LLaMA-2-7B conditioned on the original history, and our CHIQ-FT and CHIQ-AD methods described in § 3.4 and § 3.3, respectively. Underlined terms in the gold passages are those that appear in the query generated by our approaches, which is conditioned on the enhanced history and did not appear in the query generated by the method that uses the original history.

across most settings, often surpassing systems powered by closed-source LLMs. Our analysis reveals that although closed-source LLMs benefit from enhancing the history, the gap with open-source models is narrower when using the enhanced history with different facets compared to the original. Our contributions are summarized as follows:

- We propose a two-step method for query rewriting that relies on open-source language models: enhancing the conversation history and then generating the search query.
- We introduce three approaches for generating the search query on top of the enhanced conversation history: ad-hoc query rewriting (CHIQ-AD), fine-tuning a small LM for the task (CHIQ-FT), and a fusion of both approaches (CHIQ-Fusion).
- Experiments conducted on five conversational search benchmarks demonstrate that CHIQ, using open-source LLMs, achieves state-of-the-art performance across most settings, often surpassing systems that rely on closed-source LLMs.

2 Related Work

Different from traditional ad-hoc retrieval, which assumes users submit a stand-alone query, conversational search provides a conversational interface

so that users can elaborate more complex search requirements, and interactively perform search. The main challenge lies in accurately understanding the user's real search intent, which may be embedded within a longer, noisy, and more complex conversational context history. There are two well-established approaches in the literature for conversational search: Conversational Dense Retrieval (CDR) (Qu et al., 2020; Yu et al., 2021; Mao et al., 2024; Mo et al., 2024a) and Conversational Query Rewriting (CQR) (Elgohary et al., 2019).

The CDR systems aim to fine-tune an end-to-end conversational dense retriever that can directly model the entire conversational history to return relevant documents (Kim and Kim, 2022; Mo et al., 2024b). Conversely, the CQR systems focus on formulating an adequate search query based on the conversational history. This query can then serve as the input to an existing, well-established retriever-ranker framework. We base our solution on CQR, leveraging its ability to integrate with existing ad-hoc search models, which has demonstrated significant practical value (Dalton et al., 2022).

Earlier approaches to CQR attempted to select useful tokens from the conversation context (Kumar and Callan, 2020; Voskarides et al., 2020; Fang et al., 2022) or to train generative rewriter models with conversational sessions to mimic the human-written query (Yu et al., 2020; Lin et al., 2020; Vakulenko et al., 2021). To optimize query rewrites

ing, some studies have adopted reinforcement learning (Wu et al., 2022; Chen et al., 2022), or used the ranking signals with the rewriting model training (Qian and Dou, 2022; Mo et al., 2023a; Mao et al., 2023b). In addition, there have been endeavors to improve the conversion history quality through context denoising (Lin et al., 2021b; Mao et al., 2022a; Krasakis et al., 2022; Mo et al., 2023b; Mao et al., 2023c), and data augmentation (Dai et al., 2022; Mao et al., 2022b; Mo et al., 2024c). Unlike them, we enhance query rewriting by leveraging the NLP capabilities of open-source LLMs to reduce the ambiguity of conversational history. There have been multiple endeavors to integrate LLMs to solve traditional ad-hoc search sub-tasks (Zhu et al., 2023), such as query expansion (Wang et al., 2023; Gao et al., 2023), dense retrieval (Ma et al., 2023; Wang et al., 2024b), and re-ranking (Sun et al., 2023). About conversational search, Jin et al. (2023), Jang et al. (2023), and Chen et al. (2024a) attempt to improve CDR with unsupervised fine-tuning. Mao et al. (2023a), Mo et al. (2024d), and Ye et al. (2023) explore how LLMs can understand users’ contextualized search intents via CQR. Unlike direct rewriting the query using LLMs, we investigate various approaches for integrating refined conversational history into CQR frameworks.

3 Methodology

3.1 Task Formulation

Let $\mathcal{H} = \{u_k, r_k\}_{k=1}^n$ represent the user-system conversational history, where u_k and r_k are the user question and the system response at the k -th turn. Given a new user question u_{n+1} , the goal of a conversational search system is to return a set of passages \mathcal{P}_{n+1} that are relevant to \mathcal{H} and u_{n+1} , which would eventually help generate the model response r_{n+1} . To solve the main challenge of uncovering the real search intents hidden in the user’s context-dependent query, a conversational query rewriting (CQR) module has been commonly employed as an intermediate step to obtain a rewritten query q_{n+1} , which in turn is used as input to an off-the-shelf retriever. Recently, LLMs have become the default option for obtaining q_{n+1} as follows:

$$q_{n+1} \leftarrow \mathcal{LLM}(\mathcal{I}^{CQR} \oplus \mathcal{H} \oplus u_{n+1}) \quad (1)$$

where \oplus denotes concatenation and \mathcal{I}^{CQR} is a manually-engineered instruction prompt describing

the CQR task. The choice of $\mathcal{LLM}(\cdot)$ has predominantly favored closed-source commercial models, mainly CHATGPT, some of which require multiple iterations to get the optimal query (Mao et al., 2023a; Ye et al., 2023). In this work, we leverage the basic NLP capabilities of open-source LLMs to generate \mathcal{H}' , a clearer and less noisy version of \mathcal{H} . This refined version can then serve as a substitute for \mathcal{H} in Eq. 1, potentially improving the quality of q'_{n+1} using an open source LLM.

3.2 History Enhancement

In this section, we propose five approaches to tackle the ambiguity problems inherent in conversational history \mathcal{H} and map each of them to a fundamental NLP task ability. Then, we explain how we design prompts for an LLM to make part or the entire history clearer. The exact prompts we used, along with illustrative examples for each case, are presented in Appendix A.

3.2.1 Question Disambiguation

Users often expect a human-to-human level of interaction with modern conversational systems. They often use acronyms, ambiguous words, or coreference substitutes when asking questions, expecting native understanding and default reasoning capabilities from these systems. For a search system, the search intent is often unclear and ambiguous. Therefore, we propose a prompt to an LLM, denoted as \mathcal{I}^{QD} , that takes the conversational history \mathcal{H} and the subsequent user question u_{n+1} as input, to generate u'_{n+1} , a self-contained and unambiguous version of u_{n+1} which can substitute it in Eq. 1.

3.2.2 Response Expansion

In conversation sessions, it is common for model responses to be short and concise, especially for the factoid query. While brevity is often convenient for the user to acquire needed information, it makes the response less informative for a search system, which requires abundant rewrite/expansion resources. To handle this issue, we design a prompt \mathcal{I}^{RE} which instructs an LLM to enrich the content of the last model response. The goal is to make it self-contained by leveraging the preceding conversational history. The enhanced history \mathcal{H}' is obtained by replacing the original response by r'_n .

3.2.3 Pseudo Response

Given that LLMs have been demonstrated to encapsulate human knowledge, one could employ them

to speculate on potential responses directly. The intuition is that, even if the response includes some noise, it may still contain relevant terms, particularly when the LLM is prompted to produce a self-contained answer. Therefore, we design a prompt \mathcal{I}^{PR} that takes the conversational history \mathcal{H} and u_{n+1} as inputs to generate a pseudo-response r'_{n+1} . The latter can be used to expand the input of Eq. 1 to improve the quality of the query generation.

3.2.4 Topic Switch

It is natural in a conversation that different turns may focus on different aspects. Some of them are relevant to the current turn, while others may not. This is especially the case when conversations are long. In such cases, using the full history is highly likely to distract the CQR module, leading to poor query generation. Therefore, we design a prompt \mathcal{I}^{TS} that instructs the LLM to determine whether a topic switch happens between u_{n+1} and \mathcal{H} . If a switch is identified, the enhanced history would only include the last turn to maintain the transition as $\mathcal{H}' = \{u'_n, r'_n\}$. The other turns in \mathcal{H} are deemed to be irrelevant for generating q'_{n+1} and ignored.

3.2.5 History Summary

As the conversation goes on, the historical context becomes longer, which contains more irrelevant/noisy parts. A summary of history context is expected to contain only the most useful information of the original long context and better serve for query expansion. Thus, we propose the prompt \mathcal{I}^{HS} , which takes a history context (original or even enhanced) and generates a summary of the conversation \mathcal{H}' .

3.3 Ad-hoc Query Rewriting

A straightforward method to obtain the enhanced rewritten query q'_{n+1} is to independently utilize the outputs of the five methods described previously. However, by complementing each other, the outputs of these methods can collectively contribute to a more enhanced conversational history, thereby significantly improving the retrieval performance by generating a better query. We intuitively define multiple combinatory configurations for updating the input in Eq. 1, which are denoted by different symbols. The addition of **+QD** or **+RE** indicates that we replace u_{n+1} and r_n with $u_{n+1} \oplus u'_{n+1}$ (§ 3.2.1) and r'_n (§ 3.2.2) in \mathcal{H}' , respectively. **+PR** signifies that r'_{n+1} (§ 3.2.3) is concatenated to the input of Eq. 1; Lastly, **+TS** indicates that \mathcal{H}' should

omit the previous turns except the last one if a topic switch is detected. Lastly, **+HS** means that \mathcal{H} is overwritten by the entire \mathcal{H}' obtained in § 3.2.5. In our default configuration, we first check for topic-switch (**TS**). If the result is affirmative, we only use the **QD+RE+PR** configuration on the top of the truncated history in § 3.2.4. Otherwise, we apply **QD+RE+PR** on top of the original history, followed by **+HS** which consists of obtaining the history summary on top of the enhanced history. We report the results of models using this configuration in the remaining sections.

3.4 Search-Oriented Fine-tuning

In existing studies, fine-tuning conversational query generators based on a small-scale language model, such as T5-base, have proven to be both effective and efficient (Lin et al., 2020). These models consist of using human-rewritten query (Wu et al., 2022) or LLM-generated query (Jang et al., 2023) to serve as supervision signals and take \mathcal{H} and u_{n+1} as input. However, they do not take the ranking signals into account during the training and the supervision signals might be sub-optimal (Lin et al., 2021b; Mo et al., 2023a).

Considering that oracle search queries are typically unavailable and costly to annotate, we propose extending the existing approach to generate pseudo-supervision signals for query generation by leveraging the outputs produced in § 3.2. More precisely, we propose three modifications to Eq. 1 to obtain a search-oriented q'_{n+1} as follow:

$$Q'_{n+1} \leftarrow \mathcal{LLM}(\hat{\mathcal{I}}^{CQR} \oplus \mathcal{H}' \oplus u_{n+1} \oplus p_{n+1}^*) \quad (2)$$

where we replace \mathcal{H} with the enhanced history \mathcal{H}' and update the instruction \mathcal{I}^{CQR} to $\hat{\mathcal{I}}^{CQR}$ to condition the query generation on the gold passage p_{n+1}^* and prompt the LLM to generate multiple pseudo-queries in the same forward pass.

$$q'_{n+1} \leftarrow \arg \max_{q'} \mathcal{S}(Q'_{n+1}, p_{n+1}^*), q' \in Q'_{n+1} \quad (3)$$

Then, we select q'_{n+1} , the one with the highest retrieval score \mathcal{S} determined by an off-the-shelf retriever and relevance judgment from the set of pseudo-queries Q'_{n+1} as Eq 3, which is used as the supervision signal to fine-tune a rewriting model $\mathcal{M}(\mathcal{H}' \oplus u_{n+1}) = q'_{n+1}$ by maximum likelihood estimation. It is important to mention that the process described in this section is conducted offline

and performed once, for the purpose of generating pseudo-labeled queries to fine-tune a search-oriented query rewriter. During inference, \mathcal{H} and u_{n+1} serve as inputs for the fine-tuned model to generate the query q_{n+1} , and no calls are made to the LLM, so that the latency is not much affected.

4 Experimental Setup

4.1 Datasets and Evaluation Metrics

To be comparable with state-of-the-art systems (Mo et al., 2023a; Mao et al., 2023a), we consider two standard benchmarks for conversational search: TopiOCQA (Adlakha et al., 2022), QReCC (Anantha et al., 2021). TopiOCQA focuses on the challenge of the topic switch under the conversational setting, while QReCC focuses on the query rewriting problem. We run experiments on the official train-test splits and report MRR, NDCG@3, and Recall@10 to evaluate the passage retrieval results as in previous works. In addition, we evaluate three CAsT datasets (Dalton et al., 2020, 2021, 2022) which are used solely as test sets, to further validate the zero-shot or transfer learning ability of our approach, e.g., when CQR models are trained on TopiOCQA and tested on CAsTs.

4.2 Baselines

We define three main configurations for the approaches that use our enhanced generated queries:

- **CHIQ-AD** directly use the queries generated by the ad-hoc method described in § 3.3 as input for an off-the-shelf retriever.
- **CHIQ-FT** search queries generated by a small LM (e.g., T5) fine-tuned for the CQR task following the approach described in § 3.4.
- **CHIQ-Fusion** We fuse the rank list retrieved by CHIQ-AD and CHIQ-FT using the *result-level* fusion technique (Lin et al., 2021b).¹

We compared our methods with a variety of systems that can mainly be classified into three categories. More precisely, we first compare against traditional systems that fine-tune small-scale CQR models (e.g., T5-base) including: QuReTeC (Voskarides et al., 2020) T5QR (Lin et al., 2020), CONQRR (Wu et al., 2022), ConvGQR (Mo et al., 2023a), EDIRCS (Mao et al., 2023b). Then,

¹It consists of aggregating multiple ranked lists retrieved by each query into a single list to produce \mathcal{P}_{n+1} .

we compare with the systems that fine-tune an LLM-based retriever, e.g., create the query and document representation by the ending token from a decoder-only model, including RepLLaMA (Ma et al., 2023), E5-Mistral (Wang et al., 2024b), and LLM-Embedder (Zhang et al., 2023), or fine-tune an LLM-based CQR model as RETPO (Yoon et al., 2024) and IterCQR (Jang et al., 2023). Besides, we include the systems that directly obtain the rewritten query by prompting LLMs such as LLM-Aided IQR (Ye et al., 2023), HyDE (Gao et al., 2023), Query2doc (Wang et al., 2023) and LLM4CS (Mao et al., 2023a). Although not directly comparable, we report results of systems that fine-tune an ad-hoc search retriever for conversational scenarios, including the one without LLMs ConvDR (Yu et al., 2021) and with LLMs InstructorR (Jin et al., 2023). A detailed description of each aforementioned baseline is presented in Appendix B.3.

4.3 Implementation Details

We conduct experiments with the instruct-tuning variants² of both LLaMA-2-7B (Touvron et al., 2023) and Mistral-2-7B (Jiang et al., 2023) as $\mathcal{LLM}(\cdot)$ in Eq. 1 and Eq. 2. We experiment with both BM25 (Robertson et al., 2009) sparse retriever and ANCE dense retriever (Xiong et al., 2020). In addition, we use FlanT5-base³ (Chung et al., 2022) and large models as the backbone when fine-tuning a CQR model on TopiOCQA and QReCC. The fine-tuning process consists of 10 epochs with a learning rate of 1e-5 and a batch size of 8 for both datasets. More implementation details can be found in Appendix B.2.

5 Results and Analysis

5.1 Main Results

Table 1 shows both the dense and sparse retrieval performances of systems with diverse properties on the TopiOCQA and QReCC. We report the results of our systems using LLaMA-2-7B as the backbone LLM to make the results comparable with previous work. First, we observe that using our enhanced conversation history significantly improves performance over vanilla baselines that use the original history, for both ad-hoc QR (LLM4CS) and fine-tuning a small QR model (T5QR). For dense

²Concretely, the version of LLMs we used are *meta-llama/Llama-2-7b-chat-hf* and *mistralai/Mistral-7B-Instruct-v0.2* on <https://huggingface.co/>, respectively.

³We report our main results using FlanT5-base to ensure the results are comparable with previous studies.

Type	System	System Properties						TopiOCQA			QReCC		
		DR	QR	CS	OS	FT	QF	MRR	N@3	R@10	MRR	N@3	R@10
Dense (ANCE)	ConvDR	✓	✗	✗	✗	✗	✗	27.2	26.4	43.5	38.5	35.7	58.2
	InstructorR	✓	✗	✗	✓	✗	✗	25.3	23.7	45.1	43.5	40.5	66.7
	QuReTeC	✗	✓	✗	✗	✓	✓	11.2	10.5	20.2	35.0	32.6	55.0
	T5QR	✗	✓	✗	✗	✓	✗	23.0	22.2	37.6	34.5	31.8	53.1
	CONQRR	✗	✓	✗	✗	✓	✗	-	-	-	41.8	-	65.1
	ConvGQR	✗	✓	✗	✗	✓	✓	25.6	24.3	41.8	42.0	39.1	63.5
	EDIRCS	✗	✓	✗	✓	✓	✗	-	-	-	42.1	-	65.6
	IterCQR	✗	✓	✓	✗	✓	✗	26.3	25.1	42.6	42.9	40.2	65.5
	RETPO [‡]	✗	✓	✓	✓	✓	✓	30.0	28.9	49.6	44.0	41.1	66.7
	LLM-Aided	✗	✓	✓	✗	✗	✗	-	-	-	43.9	41.3	65.6
	LLM4CS	✗	✓	✗	✓	✗	✓	27.7	26.7	43.3	44.8	42.1	66.4
	CHIQ-FT	✗	✓	✗	✓	✓	✗	30.0 [†]	28.9 [†]	51.0 [†]	36.9	34.0	57.6
	CHIQ-AD	✗	✓	✗	✓	✗	✗	<u>33.2[†]</u>	<u>32.2[†]</u>	<u>53.0[†]</u>	<u>47.0[†]</u>	44.6[†]	70.8[†]
	CHIQ-Fusion	✗	✓	✗	✓	✓	✓	38.0[†]	37.0[†]	61.6[†]	47.2[†]	<u>44.2[†]</u>	<u>70.7[†]</u>
Sparse (BM25)	QuReTeC	✗	✓	✗	✗	✓	✓	8.5	7.3	16.0	34.0	30.5	55.5
	T5QR	✗	✓	✗	✗	✓	✗	11.3	9.8	22.1	33.4	30.2	53.8
	CONQRR	✗	✓	✗	✗	✓	✗	-	-	-	38.3	-	60.1
	ConvGQR	✗	✓	✗	✗	✓	✓	12.4	10.7	23.8	45.6	44.1	41.0
	EDIRCS	✗	✓	✗	✓	✓	✗	-	-	-	41.2	-	62.7
	IterCQR	✗	✓	✓	✗	✓	✗	16.5	14.9	29.3	46.7	44.1	64.4
	LLM-Aided	✗	✓	✓	✗	✗	✗	-	-	-	48.9	46.3	66.4
	LLM4CS	✗	✓	✗	✓	✗	✓	18.9	17.7	33.7	47.8	45.0	69.1
	CHIQ-FT	✗	✓	✗	✓	✓	✗	17.0 [†]	15.4 [†]	32.3 [†]	37.8	35.0	57.1
	CHIQ-AD	✗	✓	✗	✓	✗	✗	<u>22.5[†]</u>	<u>20.5[†]</u>	<u>40.4[†]</u>	<u>53.1[†]</u>	<u>50.7[†]</u>	<u>77.2[†]</u>
	CHIQ-Fusion	✗	✓	✗	✓	✓	✓	25.6[†]	23.5[†]	44.7[†]	54.3[†]	51.9[†]	78.5[†]

Table 1: Performance of dense and sparse retrieval on TopiOCQA and QReCC with different systems. We list the attributes of the reported baseline systems, which include: **DR** based on conversational dense retrieval, **QR** perform query rewriting, **CS** leverage close-source LLMs (e.g., ChatGPT or GPT-4), **OS** leverage open-source LLMs (mainly LLaMA-2-7B), **FT** fine-tune a small LM (mainly T5-base) for QR, and **QF** fuse multiple queries for retrieval. RETPO[‡] involves high-cost supervised fine-tuning an LLM for QR. † denotes significant improvements with t-test at $p < 0.05$ over all compared baselines (except CONQRR, RETPO, and IterCQR). **Bold** and underline indicate the best and the second-best results within the categories of dense and sparse retrieval.

System	\mathcal{LLM}	CASt-19			CASt-20			CASt-21		
		MRR	N@3	R@10	MRR	N@3	R@10	MRR	N@3	R@10
RepLLaMA	LLaMA-2-7B	62.4	31.6	10.6	26.8	18.3	10.4	47.4	32.7	19.6
E5-Mistral	Mistral-7B	62.2	31.3	9.5	22.0	15.4	8.4	48.2	32.5	20.5
LLM-Embedder	LLaMA-2-7B	63.3	36.6	11.4	25.2	15.4	8.7	46.8	31.2	17.3
HyDE	ChatGPT-3.5	55.6	39.2	10.0	44.8	29.3	16.9	-	-	-
Query2doc	ChatGPT-3.5	58.8	42.4	11.6	48.6	32.5	17.3	-	-	-
InstructorR	ChatGPT-3.5	61.2	46.6	10.4	43.7	29.6	8.3	46.7	32.5	18.4
	LLaMA-2-7B	68.4	45.9	11.2	52.3	37.2	<u>17.9</u>	57.0	41.5	20.2
LLM4CS	Mistral-2-7B	67.6	44.5	10.9	48.3	33.5	17.0	53.0	35.3	19.6
	ChatGPT-3.5	70.4	46.8	11.7	58.6	41.5	19.3	66.1	46.9	<u>24.4</u>
CHIQ-FT	LLaMA-2-7B	68.5	45.1	<u>11.9[†]</u>	46.3	31.6	15.9	53.9	36.0	20.4
CHIQ-AD	LLaMA-2-7B	<u>70.8[†]</u>	47.6 [†]	<u>11.9[†]</u>	51.0	34.4	<u>17.9</u>	57.7	42.0	22.6
CHIQ-Fusion	LLaMA-2-7B	73.3[†]	50.5[†]	12.9[†]	<u>54.0</u>	<u>38.0</u>	19.3	<u>62.9</u>	<u>46.5</u>	25.2[†]

Table 2: Zero-shot retrieval performances of the systems involved with LLMs under the dense retrieval (ANCE). † denotes significant improvements with t-test at $p < 0.05$ over all compared baselines. **Bold** and underline indicate the best and the second-best results, respectively.

retrieval, CHIQ-AD outperforms LLM4CS by 5.5% and 2.2% MRR on TopiOCQA and QReCC respectively, while CHIQ-FT reports a gain of 7.0% and 1.9% over T5QR on the same datasets. Similar gains are also observed using the sparse retriever, indicating the strong effectiveness of our methods.

Second, we notice that vanilla QR systems on top of an enhanced history can outperform systems that utilize additional training techniques and sophisticated modules. For instance, CHIQ-AD outperforms both ConvDR and InstructorR, which need relevance judgments to fine-tune a conversational dense retriever on the raw input; ReTPO, which fine-tunes an LLM for QR and in addition leverages GPT-4 for data augmentation. While CHIQ-FT outperforms its direct competitors, primarily ConvGQR and IterCQR that refine the supervision signals on TopiOCQA, it underperforms on QReCC mainly because previous fine-tuned QR models rely on QReCC’s human-rewritten queries. The contrasting observations between the two datasets suggest that enhancing the history is crucial for performance when no QR-supervised annotations exist.

Third, we observe that by systematically fusing the outputs of our approaches, CHIQ-Fusion outperforms the models that use each component separately, achieving the best performance across most settings. The gains are more significant on the topic-mixed and more challenging TopiOCQA compared to QReCC, with 4.8% and 0.2% MRR score improvements on each dataset. Interestingly, this occurs even though CHIQ-FT significantly underperforms compared to CHIQ-AD, suggesting that CHIQ-FT still generates query content that is complementary and not captured by CHIQ-AD.

5.2 Zero-shot Results

We compare the dense retrieval performances of different systems that leverage LLMs under a zero-shot manner on three CAsT datasets in Table 2. We observe a consistent pattern as previous results in Table 1 when comparing the performances within our approaches. More precisely, although CHIQ-FT performs slightly worse compared to CHIQ-AD, fusing their outputs systematically leads to better performances across all three datasets. Besides, we can see that CHIQ-AD outperforms most systems either utilizing open-source or close-source LLMs and yields results competitive with the state-of-the-art system LLM4CS, which requires multiple calling for each query

turn. Specifically, CHIQ-AD surpasses LLM4CS with LLaMA-2-7B, on CAsT-19 and CAsT-21. In addition, our top-performing approach, CHIQ-Fusion, outperforms all compared systems, except the LLM4CS with close-source ChatGPT-3.5 on CAsT-20 and CAsT-21, indicating the superior effectiveness of our approaches. We also find that ad-hoc fine-tuned LLM-based retrievers (RepLLaMA, E5-Mistral, and LLM-Embedder) underperform the systems with LLM-based query generation (HyDE and Query2doc) and the InstructorR with conversational fine-tuning adaption. These systems also underperform CHIQ-FT, which only fine-tunes a small LM on TopiOCQA with enhanced supervision signals. The observation indicates the importance of improving the generalization capabilities of the models to handle complex and diverse conversational scenarios.

5.3 Open vs. Close Source LLMs

In addition to conducting experiments with open-source LLMs, we also deploy the closed-source LLM ChatGPT-3.5 to isolate the effects of history enhancement. Table 3 shows the dense retrieval performances on three CAsT test sets when query rewriting (QR) is performed by CHIQ-AD approach on both the original history and the enhanced one. We observe that ChatGPT-3.5 benefits from QR on enhanced history, with NDCG@3 score improvements of 5.4%, 4.3%, and 0.9% across CAsT-19, CAsT-20, and CAsT-21, respectively. Such results indicate that despite the superior reasoning abilities of closed-source LLM, enhancing the conversational history is deemed important for handling complex queries within conversational scenarios.

\mathcal{LLM}	CAsT-19		CAsT-20		CAsT-21	
	MRR	N@3	MRR	N@3	MRR	N@3
Original History						
LLaMA	67.4	42.5	40.9	27.9	52.7	37.2
Mistral	67.9	42.0	44.2	30.4	59.5	41.6
ChatGPT	69.3	40.8	53.0	36.2	60.3	41.9
Our Enhanced History						
LLaMA	70.8	47.6	51.0	34.4	57.7	42.0
Mistral	71.4	47.2	49.2	34.4	67.0	47.2
ChatGPT	71.7	46.4	55.7	40.5	62.2	42.8

Table 3: Dense retrieval results for systems using various LLMs as backbones, where QR is performed either directly on top of the original conversation history or on our enhanced history using the CHIQ-AD method.

Also, we find that conducting QR on enhanced conversational history helps to narrow the performance gap between open-source and closed-source

LLMs. For instance, the gap of MRR score between LLaMA and ChatGPT-3.5 on the original history is 1.9%, 12.1%, and 7.6% across three CASt test sets, respectively. In contrast, when utilizing enhanced history, the gaps reduced significantly to 0.9%, 4.7%, and 4.5%, indicating that our designed approach can adequately leverage the capacity of open-source LLMs for conversational search and be competitive with close-source ones.

5.4 Search-Oriented Fine-tuning Ablation

We analyze the potential choices for generating search queries in Eq. 2 and Eq. 3 as supervision signals for CHIQ-FT models. Table 4 presents the dense retrieval performances via the queries generated by CHIQ-FT models, which are fine-tuned using manually rewritten queries or the variants of the approach outlined in §3.4. The ablations are based on the results of either without using enhanced history, without generating multiple queries, or not conditioning on the gold passage. We observe that using the queries generated by LLMs as supervision signals outperform the one using manual annotation, which is consistent with previous studies that have identified human-written queries as sub-optimal (Wu et al., 2022; Mo et al., 2023a).

Signal	TopiOCQA		QReCC	
	MRR	N@3	MRR	N@3
Manual	-	-	34.7	31.9
CHI-Q-FT	30.0	28.9	36.9	34.0
w.o. \mathcal{H}'	27.6	26.7	35.4	33.8
w.o. Q'_{n+1}	24.2	23.4	33.4	31.7
w.o. p_{n+1}^*	18.2	17.2	26.8	23.9

Table 4: Dense retrieval performances of fine-tuned QR models, which utilize different supervised signals. These include human-written queries and variants from CHIQ-FT, either without enhanced history, without multiple queries, or not conditioned on the gold passage.

Additionally, we observe that all of our proposed adaptive modifications significantly enhance the final performance of the system, especially applying them all to generate the pseudo supervision signals. Such results indicate that improving the quality of supervision signals is crucial for QR model fine-tuning and justifying the effectiveness of our approaches for search-oriented fine-tuning.

5.5 History Enhancement Ablation

We study the contribution of each of our proposed prompts in § 3.2 for query enhancement by conducting ablation studies on \mathcal{H}' when one of the

prompts is not used. Table 5 presents the ablation performances of dense and sparse retrieval on TopiOCQA and QReCC. The History Summary (**HS**) is ablated alongside (**TS**), as **HS** is not activated if a new topic is detected.

Type	Ablation	TopiOCQA		QReCC	
		MRR	N@3	MRR	N@3
Dense (ANCE)	CHI-Q-AD	33.2	32.2	47.0	44.6
	w.o. QD	32.5	31.4	44.6	41.9
	w.o. RE	28.3	27.0	46.6	44.0
	w.o. PR	26.4	25.2	43.5	40.8
	w.o. TS	20.0	18.7	46.9	44.4
Sparse (BM25)	CHI-Q-AD	22.5	20.5	53.1	50.7
	w.o. QD	22.1	20.1	47.3	44.6
	w.o. RE	19.0	16.8	50.1	47.3
	w.o. TS	17.8	16.5	51.7	48.8
	w.o. PR	16.9	15.3	46.9	44.5

Table 5: Dense and sparse retrieval results of ablating CHIQ-AD by not using one history enhancement prompt at each line on TopiOCQA and QReCC datasets.

We observe that all our proposed enhancements to the history context contribute positively to the performance of the CHIQ-AD method, although some enhancements are more effective than others. On one hand, detecting topic switching is particularly crucial on TopiOCQA, leading to performance improvements of 13.2% and 5.2% MRR scores in dense and sparse retrieval, respectively. This is mainly due to the multi-topic focus design of the dataset within the same conversation. On the other hand, we notice that while question disambiguation (**QD**) improves performance, it is less critical compared to predicting a pseudo response (**PR**) or enhancing the quality of the last system response (**RE**). In addition, we notice that all our proposed enhancements contribute similarly to the generated search queries across both dense and sparse retrieval settings.

5.6 Case Analysis

We manually analyze the content of the enhanced history to better understand the mechanisms and limitations of our approach. This analysis shows the complementary roles each enhancement prompt plays in improving the quality of the original history. **QD** and **RE** primarily assist in resolving coreferences and clarifying acronyms to full names, **TS** helps remove irrelevant content, **PR** speculates on relevant terms that may occur in the response, and **HS** not only converts the conversation into plain text but also ensures that key terms from the conversation are preserved. While prompts such as

PR and **RE** generation generally aid in retrieval, they may also introduce noisy terms due to the wrong fact generated by LLMs that hurt the ranking results. Finally, we also notice that the queries generated by CHIQ-AD and CHIQ-FT are of different styles. The first focuses on expanding more relevant terms to increase the matching scores, while the latter queries are more concise with higher efficiency for retrieval. Nevertheless, aggregating the output rank lists from both approaches helps refine the final results by ranking the relevant passages higher. The concrete examples of these case analyses are presented in Appendix D.

6 Conclusion

In this paper, we propose CHIQ, an approach that leverages the basic NLP capabilities of LLMs to enhance the quality of contextual history for improving the query rewriting performance in terms of conversational search. Despite its simplicity, our approach achieves superior performance across various datasets and settings, using open-source LLMs compared to closed-source alternatives. This study shows that instead of simply ask an LLM to generate a search query, it is critical to design strategies to generate different facets of enhancement in view of finding the target information.

Limitations

Potential limitations of this work include not experimenting with larger open-source LLMs, such as the 56B Mixtral (Mixtral AI team, 2023) or 70B LLaM a, as well as other recent models like Gemma (Team et al., 2024). Additionally, the study did not incorporate more closed-source models such as GPT-4 (OpenAI, 2023) or Claude (AnthropicAI, 2023) to further study the impact of history enhancement. This is mainly due to limitations in computation (open source) and financial (close source) resources. Despite the straightforward and significant gains, some design choices could be further analyzed to potentially boost the performance even more. For instance, adding a backoff filtering strategy could detect when the LLM is producing noisy outputs, or exploring approaches that interpolate between the use of human and pseudo-queries when its higher quality as training signals for CHIQ-FT. Besides, we have considered 5 directions of enhancement in this paper. More strategies can be incorporated so that other useful enhancements can be integrated.

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.
- 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*, pages 520–534.
- AnthropicAI. 2023. Introducing claude.
- Haonan Chen, Zhicheng Dou, Kelong Mao, Jiongnan Liu, and Ziliang Zhao. 2024a. Generalizing conversational dense retrieval via llm-cognition data augmentation. *arXiv preprint arXiv:2402.07092*.
- Pei Chen, Shuai Zhang, and Boran Han. 2024b. Comm: Collaborative multi-agent, multi-reasoning-path prompting for complex problem solving. In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 1720–1738.
- Zhiyu Chen, Jie Zhao, Anjie Fang, Besnik Fetahu, Rokhlenko Oleg, and Shervin Malmasi. 2022. Reinforced question rewriting for conversational question answering.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. 2022. Scaling instruction-finetuned language models. *arXiv preprint arXiv:2210.11416*.
- 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*, pages 4558–4586. PMLR.
- Jeffrey Dalton, Chenyan Xiong, and Jamie Callan. 2020. Trec cast 2019: The conversational assistance track overview. *arXiv preprint arXiv:2003.13624*.
- Jeffrey Dalton, Chenyan Xiong, and Jamie Callan. 2021. Cast 2020: The conversational assistance track overview. Technical report.
- Jeffrey Dalton, Chenyan Xiong, and Jamie Callan. 2022. Trec cast 2021: The conversational assistance track overview. In *In Proceedings of TREC*.
- Ahmed Elgohary, Denis Peskov, and Jordan Boyd-Graber. 2019. Can you unpack that? learning to rewrite questions-in-context. 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)*, pages 5918–5924.

- Hung-Chieh Fang, Kuo-Han Hung, Chen-Wei Huang, and Yun-Nung Chen. 2022. Open-domain conversational question answering with historical answers. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2022*, pages 319–326.
- Jianfeng Gao, Chenyan Xiong, Paul Bennett, and Nick Craswell. 2022. Neural approaches to conversational information retrieval. *arXiv preprint arXiv:2201.05176*.
- Luyu Gao, Xueguang Ma, Jimmy Lin, and Jamie Callan. 2023. Precise zero-shot dense retrieval without relevance labels. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1762–1777.
- Arnav Gudibande, Eric Wallace, Charlie Snell, Xinyang Geng, Hao Liu, Pieter Abbeel, Sergey Levine, and Dawn Song. 2023. The false promise of imitating proprietary llms. *arXiv preprint arXiv:2305.15717*.
- Kaiyu Huang, Fengran Mo, Hongliang Li, You Li, Yuanchi Zhang, Weijian Yi, Yulong Mao, Jinchun Liu, Yuzhuang Xu, Jinan Xu, et al. 2024. A survey on large language models with multilingualism: Recent advances and new frontiers. *arXiv preprint arXiv:2405.10936*.
- Yunah Jang, Kang-il Lee, Hyunkyung Bae, Seungpil Won, Hwanhee Lee, and Kyomin Jung. 2023. Itercqr: Iterative conversational query reformulation without human supervision. *arXiv preprint arXiv:2311.09820*.
- 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.
- 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*, pages 6649–6675.
- Jeff Johnson, Matthijs Douze, and Herv   J  gou. 2019. Billion-scale similarity search with gpus. *IEEE Transactions on Big Data*, 7(3):535–547.
- Jean Kaddour, Joshua Harris, Maximilian Mozes, Herbie Bradley, Roberta Raileanu, and Robert McHardy. 2023. Challenges and applications of large language models. *arXiv preprint arXiv:2307.10169*.
- Sungdong Kim and Gangwoo Kim. 2022. Saving dense retriever from shortcut dependency in conversational search. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 10278–10287. Association for Computational Linguistics.
- Antonios Minas Krasakis, Andrew Yates, and Evangelos Kanoulas. 2022. Zero-shot query contextualization for conversational search. *arXiv preprint arXiv:2204.10613*.
- 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*, pages 3971–3980.
- Jimmy Lin, Xueguang Ma, Sheng-Chieh Lin, Jheng-Hong Yang, Ronak Pradeep, and Rodrigo Nogueira. 2021a. Pyserini: A python toolkit for reproducible information retrieval research with sparse and dense representations. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2356–2362.
- Sheng-Chieh Lin, Jheng-Hong Yang, and Jimmy Lin. 2021b. Contextualized query embeddings for conversational search. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1004–1015.
- Sheng-Chieh Lin, Jheng-Hong Yang, Rodrigo Nogueira, Ming-Feng Tsai, Chuan-Ju Wang, and Jimmy Lin. 2020. Conversational question reformulation via sequence-to-sequence architectures and pretrained language models. *arXiv preprint arXiv:2004.01909*.
- Xueguang Ma, Liang Wang, Nan Yang, Furu Wei, and Jimmy Lin. 2023. Fine-tuning llama for multi-stage text retrieval. *arXiv preprint arXiv:2310.08319*.
- 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. *arXiv preprint arXiv:2404.13556*.
- Kelong Mao, Zhicheng Dou, Haonan Chen, Fengran Mo, and Hongjin Qian. 2023a. Large language models know your contextual search intent: A prompting framework for conversational search.
- Kelong Mao, Zhicheng Dou, Bang Liu, Hongjin Qian, Fengran Mo, Xiangli Wu, Xiaohua Cheng, and Zhao Cao. 2023b. Search-oriented conversational query editing. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 4160–4172.
- Kelong Mao, Zhicheng Dou, and Hongjin Qian. 2022a. Curriculum contrastive context denoising for few-shot conversational dense retrieval. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 176–186.
- Kelong Mao, Zhicheng Dou, Hongjin Qian, Fengran Mo, Xiaohua Cheng, and Zhao Cao. 2022b. Contrans: Transforming web search sessions for conversational dense retrieval. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 2935–2946.

- Kelong Mao, Hongjin Qian, Fengran Mo, Zhicheng Dou, Bang Liu, Xiaohua Cheng, and Zhao Cao. 2023c. Learning denoised and interpretable session representation for conversational search. In *Proceedings of the ACM Web Conference 2023*, pages 3193–3202.
- Mixtral AI team. 2023. [Mixtral of experts a high quality sparse mixture of experts](#).
- Fengran Mo, Kelong Mao, Yutao Zhu, Yihong Wu, Kaiyu Huang, and Jian-Yun Nie. 2023a. Convgqr: Generative query reformulation for conversational search. *arXiv preprint arXiv:2305.15645*.
- Fengran Mo, Jian-Yun Nie, Kaiyu Huang, Kelong Mao, Yutao Zhu, Peng Li, and Yang Liu. 2023b. Learning to relate to previous turns in conversational search. In *29th ACM SIGKDD Conference On Knowledge Discover and Data Mining (SIGKDD)*.
- Fengran Mo, Chen Qu, Kelong Mao, Yihong Wu, Zhan Su, Kaiyu Huang, and Jian-Yun Nie. 2024a. Aligning query representation with rewritten query and relevance judgments in conversational search. *arXiv preprint arXiv:2407.20189*.
- Fengran Mo, Chen Qu, Kelong Mao, Tianyu Zhu, Zhan Su, Kaiyu Huang, and Jian-Yun Nie. 2024b. History-aware conversational dense retrieval. *arXiv preprint arXiv:2401.16659*.
- Fengran Mo, Bole Yi, Kelong Mao, Chen Qu, Kaiyu Huang, and Jian-Yun Nie. 2024c. Convsdg: Session data generation for conversational search. In *Companion Proceedings of the ACM on Web Conference 2024*, pages 1634–1642.
- Fengran Mo, Longxiang Zhao, Kaiyu Huang, Yue Dong, Degen Huang, and Jian-Yun Nie. 2024d. How to leverage personal textual knowledge for personalized conversational information retrieval. *arXiv preprint arXiv:2407.16192*.
- OpenAI. <https://platform.openai.com/docs/models/gpt-3-5-turbo>.
- OpenAI. 2023. [Gpt-4 technical report](#). *ArXiv*, abs/2303.08774.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744.
- Hongjin Qian and Zhicheng Dou. 2022. Explicit query rewriting for conversational dense retrieval. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 4725–4737.
- 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*, pages 539–548.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21(140):1–67.
- Stephen Robertson, Hugo Zaragoza, et al. 2009. The probabilistic relevance framework: Bm25 and beyond. *Foundations and Trends® in Information Retrieval*, 3(4):333–389.
- Weiwei Sun, Lingyong Yan, Xinyu Ma, Shuaiqiang Wang, Pengjie Ren, Zhumin Chen, Dawei Yin, and Zhaochun Ren. 2023. Is chatgpt good at search? investigating large language models as re-ranking agents. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 14918–14937.
- Gemma Team, Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak, Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, et al. 2024. Gemma: Open models based on gemini research and technology. *arXiv preprint arXiv:2403.08295*.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton-Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurélien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. [Llama 2: Open foundation and fine-tuned chat models](#). *CoRR*, abs/2307.09288.
- Svitlana Vakulenko, Shayne Longpre, Zhucheng Tu, and Raviteja Anantha. 2021. Question rewriting for conversational question answering. In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining*, pages 355–363.
- Christophe Van Gysel and Maarten de Rijke. 2018. Pytrec_eval: An extremely fast python interface to trec_eval. In *SIGIR*. 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*, pages 921–930.
- Jiayin Wang, Fengran Mo, Weizhi Ma, Peijie Sun, Min Zhang, and Jian-Yun Nie. 2024a. A user-centric benchmark for evaluating large language models. *arXiv preprint arXiv:2404.13940*.
- Liang Wang, Nan Yang, Xiaolong Huang, Linjun Yang, Rangan Majumder, and Furu Wei. 2024b. [Improving text embeddings with large language models](#). *CoRR*, abs/2401.00368.
- Liang Wang, Nan Yang, and Furu Wei. 2023. Query2doc: Query expansion with large language models. *arXiv preprint arXiv:2303.07678*.
- Zequ Wu, Yi Luan, Hannah Rashkin, David Reitter, and Gaurav Singh Tomar. 2022. Conqrr: Conversational query rewriting for retrieval with reinforcement learning.
- Lee Xiong, Chenyan Xiong, Ye Li, Kwok-Fung Tang, Jialin Liu, Paul N Bennett, Junaid Ahmed, and Arnold Overwijk. 2020. Approximate nearest neighbor negative contrastive learning for dense text retrieval. In *International Conference on Learning Representations*.
- Fanghua Ye, Meng Fang, Shenghui Li, and Emine Yilmaz. 2023. Enhancing conversational search: Large language model-aided informative query rewriting. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 5985–6006.
- Chanwoong Yoon, Gangwoo Kim, Byeongguk Jeon, Sungdong Kim, Yohan Jo, and Jaewoo Kang. 2024. Ask optimal questions: Aligning large language models with retriever’s preference in conversational search. *arXiv preprint arXiv:2402.11827*.
- Shi Yu, Jiahua Liu, Jingqin Yang, Chenyan Xiong, Paul 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*, pages 1933–1936.
- Shi Yu, Zhenghao Liu, Chenyan Xiong, Tao Feng, and Zhiyuan Liu. 2021. Few-shot conversational dense retrieval. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 829–838.
- Hamed Zamani, Johanne R Trippas, Jeff Dalton, Filip Radlinski, et al. 2023. Conversational information seeking. *Foundations and Trends® in Information Retrieval*, 17(3-4):244–456.
- Peitian Zhang, Shitao Xiao, Zheng Liu, Zhicheng Dou, and Jian-Yun Nie. 2023. [Retrieve anything to augment large language models](#). *CoRR*, abs/2310.07554.
- Yutao Zhu, Huaying Yuan, Shuting Wang, Jiongnan Liu, Wenhan Liu, Chenlong Deng, Zhicheng Dou, and Ji-Rong Wen. 2023. Large language models for information retrieval: A survey. *arXiv preprint arXiv:2308.07107*.

A History Enhancement and Query Rewriting Prompts

In this section, we list the prompts that we have carefully designed to enhance different parts of the conversation history, as well as the prompt used for query rewriting and pseudo supervision signals for search-originated fine-tuning. For each prompt, we designed the instruction part through trial and error iterations until we confirmed that both models (LLaMA-2-7B and Mistral-v0.2-7B) could follow the instructions and generate outputs in the required format. We observed no benefits from using in-context examples for any model, as the outputs remained mostly stable, with minor to no changes in the model responses even after adding these examples.

A.1 Question Disambiguation

You are given a set of question-answers pairs and a new question that is ambiguous. Your goal is to rewrite the question so it becomes clear. Write the new question without any introduction.

A.2 Response Expansion

Given a series of question-and-answer pairs, along with a new question, your task is to give a one-sentence response to the new question.

A.3 Pseudo Response

You are given a question-and-answer pair, where the answer is not clear. Your goal is to write a long version of the answer based on its given context. The generated answer should be one sentence only and less than 20 words.

A.4 Topic Switch

Given a series of question-and-answer pairs, along with a new question, your task is to determine whether the new question continues the discussion on an existing topic or introduces a new topic. Please respond with either "new_topic" or "old_topic" as appropriate.

A.5 History Summary

You are given a context in the form of question-answer pairs. Your goal is to write a paragraph that summarizes the information in the context. The summary should be short with one sentence for each question answer pair.

A.6 Query Rewriting

Given a series of question-and-answer pairs as context, along with a new question, your task is to convert the new question into a search engine query that can be used to retrieve relevant documents. The output should be placed in a JSON dictionary as follows: {"query": ""}

A.7 Pseudo Supervision Signals

You are given a relevant passage, a series of question-and-answer pairs as context along with a new question, your task is to generate a set of search queries based on the relevancy between the new question and the relevant passage and also rely on the given context. The output format should be in a list with indexes e.g., 1. 2. 3.

B Experimental Setup

B.1 Datasets Details

Dataset	Split	#Conv.	#Turns(Qry.)	#Collection
TopiOCQA	Train	3,509	45,450	25M
	Test	205	2,514	
QReCC	Train	10,823	29,596	54M
	Test	2,775	8,124	
CAsT-19	Test	50	479	38M
CAsT-20	Test	25	208	
CAsT-21	Test	26	239	

Table 6: Statistics of conversational search datasets.

The statistics of each dataset are presented in Table 6. We discard the samples without gold passages. The manually rewritten query for each turn is provided in all datasets except TopiOCQA. The relevance judgments in the CAsT datasets are made by experts with multi-level annotations. The relevance judgment thresholds are set at 1, 2, and 2 for CAsT-19, CAsT-20, and CAsT-21, respectively.

B.2 Implementation Details

We implement the retrieval evaluation metrics from the `pytrec_eval` tool (Van Gysel and de Rijke, 2018). We leverage the Pyserini (Lin et al., 2021a) and Faiss (Johnson et al., 2019) libraries for implementing the BM25 and ANCE retrievers, respectively. Following previous works (Lin et al., 2021b; Mo et al., 2023a), we set BM25 parameters as follows: $k_1 = 0.9$, $b = 0.4$ on TopiOCQA and $k_1 = 0.82$, $b = 0.68$ on the QReCC. The lengths of the query, concatenated input, and passage are truncated to 32, 512, and 384 tokens, respectively.

In all experiments, we use sampling and set the temperature to 0.7 when generating with LLMs.

For the search-oriented fine-tuning, we use the NDCG@3 score as the standard metric to select the generated query as the pseudo supervision signals, while we set the maximum length for the generated query is set to 32, which is the same as (Lin et al., 2020; Wu et al., 2022; Mo et al., 2023a). For the rank-list fusion, we set the balance factor α in Lin et al. (2021b) as 1, which indicates the same importance of different retrieved results.

B.3 Baselines Details

We provide a more detailed introduction to the following baselines used for comparison:

QuReTeC (Voskarides et al., 2020): A traditional sequence tagger query rewriting approach fine-tuned with weakly supervision signals to determine whether a term in a historical context should be expanded to the current query.

T5QR (Lin et al., 2020): A query rewriting approach fine-tuned with manual annotations provided in QReCC as the supervised signals via the T5-base model.

CONQRR (Wu et al., 2022): A query rewriting approach fine-tuned with manual annotations provided in QReCC and the ranking signals using reinforcement learning via the T5-base model.

ConvGQR (Mo et al., 2023a): A unified framework that integrates query rewriting and query expansion mechanisms by two T5-base models and fine-tuned them with manual annotations and ground-truth response, respectively.

EDIRCS (Mao et al., 2023b): A query rewriting approach based on text editing technique with ranking signals fine-tuned on the T5-base model.

IterCQR (Jang et al., 2023): An iterative query rewriting method using the initial query generated by the ChatGPT-3.5 and refining a query turn multiple times according to the ranking signals feedback during the training stage.

LLM-Aided (Ye et al., 2023): An informative conversational query rewriting by directly prompting ChatGPT-3.5 as both query rewriters and rewrite editors twice to incorporate all the desirable properties for producing the final rewritten queries.

RETPO (Yoon et al., 2024): A retriever preference adapted query rewriting method that fine-tunes LLaMA-2-7B as a QR model with an external QR dataset generated by GPT-4.

ConvDR (Yu et al., 2021): A traditional conversational dense retrieval method that uses

knowledge distillation to learn the session embeddings with relevance judgments from the human-rewritten queries based on the ANCE model.

InstructorR (Jin et al., 2023): A LLM-based general conversational dense retriever tailored to various tasks and domains by fine-tuned with various task-specific instructions and relevance judgments based on FlanT5-XL model.

RepLLaMA (Ma et al., 2023): A large ad-hoc dense retriever fine-tuned on top of the LLaMA-7B model on the MSMARCO dataset.

E5-Mistral (Wang et al., 2024b): A large ad-hoc retriever fine-tuned on top of Mistral-7B model on the synthetic dataset generated by ChatGPT-3.5 and MSMARCO.

LLM-Embedder (Zhang et al., 2023): A unified retrieval model that can support diverse retrieval augmentation needs of LLMs, which is fine-tuned on various tasks and datasets such as MSMARCO, NQ, ToolLLM, QReCC, FLAN, Books3, and Multi-Session Chat.

HyDE (Gao et al., 2023): A zero-shot retrieval method, which adopts ChatGPT-3.5 to generate hypothetical documents for the query, then retrieves real documents with hypothetical documents.

Query2doc (Wang et al., 2023): A zero-shot query expansion approach, which expands the original query with the generated documents from ChatGPT-3.5.

LLM4CS (Mao et al., 2023a): A state-of-the-art LLM-based prompting method for conversational query rewriting. We implement it with full aggregation by calling LLMs five times for the query and response generation but without the chain-of-thought (CoT) content because of the efficient annotation consideration in practical scenarios.

C Results of Mistral-2-7B

Table 7 and Table 8 show the performances of our methods using Mistral-2-7B as the backbone LLM, replacing LLama-2-7B as previously shown in Table 1 and Table 2, respectively. These results exhibit trends similar to those observed with LLama-2-7B, as discussed in Sections 5.1 and 5.2. More specifically, CHIQ-AD consistently outperforms CHIQ-FT across all settings, and combining the results of both methods (CHIQ-Fusion) yields the best performance. It is worth noting that results with Mistral-2-7B are systematically lower than the ones with LLama-2-7B on most datasets (except on CAsT-21). Besides, by comparing the fusion per-

Type	Systems	TopiOCQA			QReCC		
		MRR	NDCG@3	Recall@10	MRR	NDCG@3	Recall@10
Dense	CHIQ-FT	26.2	25.4	45.3	31.1	28.5	50.6
	CHIQ-AD	28.9 [†]	28.3 [†]	46.8 [†]	46.7 [†]	44.2[†]	70.7[†]
	CHIQ-Fusion	36.3[†]	35.0[†]	59.6[†]	47.1[†]	44.1 [†]	70.3 [†]
Sparse	CHIQ-FT	15.2	14.3	30.5	32.9	30.0	51.4
	CHIQ-AD	19.2 [†]	17.3 [†]	35.6 [†]	51.7 [†]	48.8 [†]	76.2 [†]
	CHIQ-Fusion	21.4[†]	19.2[†]	39.4[†]	51.9[†]	49.0[†]	76.3[†]

Table 7: Performance of dense and sparse retrieval on TopiOCQA and QReCC datasets based on Mistral-2-7B model. The system properties and the settings are inherited from the Table 1.

Systems	CASt-19			CASt-20			CASt-21		
	MRR	NDCG@3	R@10	MRR	NDCG@3	R@10	MRR	NDCG@3	R@10
CHIQ-FT	58.3	35.4	9.0	37.1	24.7	12.0	44.4	29.1	16.5
CHIQ-AD	71.4 [†]	47.2 [†]	12.2[†]	49.2	34.4	17.5	67.0[†]	47.2 [†]	25.5 [†]
CHIQ-Fusion	71.5[†]	47.7[†]	11.9 [†]	51.1	35.5	18.4	66.1	48.9[†]	27.5[†]

Table 8: Performance of dense retrieval on three CASt datasets based on Mistral-2-7B model. The system properties and the settings are inherited from the Table 2.

formance between Mistral and LLaMa, we notice that in most cases when the gap between CHIQ-AD and CHIQ-FT is large, CHIQ-Fusion results are either slightly better or worse than CHIQ-AD. This is mainly because the poor quality of the rank list obtained by CHIQ-FT negatively impacts the one from CHIQ-AD. However, when the gap is smaller, we notice a significant gain for CHIQ-Fusion, suggesting that both variants are generating good and complementary rank lists. It will be interesting to investigate how we can better take advantage of CHIQ-AD and CHIQ-FT in an adaptive fusion. Nevertheless, performing QR on top of the enhanced history with our approach still outperforms most other settings and datasets.

D Case Analysis

Table 9 and Table 10 showcase two examples that support the case study analysis conducted in § 5.6. In the example of Table 9, **QD** and **RE QD** and **RE** contribute by adding the terms “hormone” into “cortisol, glucagon, adrenaline, cytokines orexin, and melatonin” to the enhanced history. In addition, **PR** enriches the emotional context by including “excitement, anxiety, or fear”, which co-occurs in the gold positive passage, thereby improving the scores. Since a topic switch (**TS**) is detected, earlier turns containing noisy terms, such as “Adenosine triphosphate (ATP)”, are dropped from the history summary (**HS**). Therefore, performing QR with our methods on the enhanced history results in top pas-

sages being ranked higher compared to those based on the original history (LLM-QR). Similar trends are observed in the second case shown in Table 10, where CHIQ-FT and CHIQ-AD outperform QR with the original history. This improvement may be attributed to the enhanced history increasing semantic similarity through references to titles of Tjader’s artworks from **PR** and **HS**, such as “Inauguration of the Pleasure Dome”. However, names of various collaborators like “Kenneth Anger and Stan Brakhage” may introduce noise in **PR**.

Original History	Our Enhanced History
<p>u_1: The primary high energy molecule in human metabolism is?</p> <p>r_1: Adenosine Triphosphate (ATP).</p> <p>u_2: What is catabolism?</p> <p>r_2: It is the set of metabolic processes that breaks down large molecules.</p> <p>u_3: Which <u>hormones</u> are related to it?</p> <p>r_3: Cortisol, Glucagon, <u>Adrenaline</u>, Cytokines, Orexin, and Melatonin.</p> <p>u_4: What is the emotional response due to the <u>third one</u>?</p>	<p>QD: Which hormone among cortisol, glucagon, adrenaline, cytokines orexin, and melatonin is associated with an emotional response? TS: Yes</p> <p>RE: These hormones, including cortisol, glucagon, adrenaline, cytokines, orexin, and melatonin, play various roles in regulating metabolic processes involving the breakdown of larger molecules to produce ATP during catabolism.</p> <p>PR: The emotional responses to hormones can vary greatly among individuals. Adrenaline (epinephrine) is often associated with excitement, anxiety, or fear.</p> <p>HS: Adenosine triphosphate (ATP) serves as the primary high energy molecule in human metabolism. Catabolism refers to the metabolic processes that break down large molecules into smaller ones. Cortisol, glucagon, adrenaline, cytokines, orexin, and melatonin are associated hormones involved in this process.</p>
Rewritten Query q'_4	Gold Positive Passage p_4^*
<p>LLM-QR: What is the emotional response associated with the third <u>hormone</u>? (Rank: 41)</p> <p>CHIQ-AD: QD + PR + RE (Rank: 4)</p> <p>CHIQ-FT: How does <u>adrenaline</u> impact mood? (Rank: 3)</p> <p>CHIQ-Fusion (Rank: 1)</p>	<p>Every emotional response has a behavioral component, an autonomic component, and a hormonal component. The hormonal component includes the release of adrenaline, an adrenomedullary response that occurs in response to stress and that is controlled by the sympathetic nervous system. The major emotion studied in relation to adrenaline is fear. In an experiment, subjects who were injected with adrenaline expressed more negative and fewer positive facial expressions to fear films compared to a control group.</p>

Table 9: The first case for analyzing the content between the original and our enhanced history, as well as the rewritten queries generated on top of them. The blue tokens and the orange tokens stand for the effective and noisy patterns for matching, respectively. The underline tokens denote the coreference relation between rewritten queries and the original context. The Rank indicates the ranking position of the gold positive passage w.r.t each query.

Original History	Our Enhanced History
<p>u_1: Are there any interesting aspects about Cal Tjader?</p> <p>r_1: Cal Tjader entered the United States Navy in 1943 served as a medical corpsman in the Pacific Theater until March 1946.</p> <p>u_2: Were any albums released?</p> <p>r_2: At San Francisco State he met Dave Brubeck, a young pianist fresh from a stint in the Army. Brubeck introduced Tjader to Paul Desmond.</p> <p>u_3: What were the jazz album released for?</p> <p>r_3: Cal Tjader continued the trio work in California with bassist Jack Weeks and pianists John Marabuto ... recording his first LP with them for Fantasy Records.</p> <p>u_4: What was a title of <u>one of the albums</u>?</p>	<p>QD: What was the outcome or reception of Kenneth Anger's collaboration with Stan Brakhage on their film project in the United States?</p> <p>PR: It is unclear if Kenneth Anger and Stan Brakhage's confiscated film was successful as it is presumed to have been destroyed. However, Anger's later film, Inauguration of the Pleasure Dome (1954), received critical acclaim for its themes. TS: No</p> <p>HS: In 1953, Kenneth Anger returned to the United States following his mother's death to handle her estate. There, he befriended Stan Brakhage and together they produced a controversial film, which was confiscated and likely destroyed due to its obscene content. Despite this setback, Anger went on to create his groundbreaking 38-minute surrealist work, "Inauguration of the Pleasure Dome," in 1954, showcasing Crowleyan and Thelemite themes.</p>
Rewritten Query q'_4	Gold Positive Passage p_4^*
<p>LLM-QR: Which album did Cal Tjader record with Jack Weeks and John Marabuto in California?(Rank:16)</p> <p>CHIQ-AD: QD + PR + HS (Rank: 7)</p> <p>CHIQ-FT: Which specific album title did <u>Cal Tjader</u> record with Jack Weeks on bass and either John with Marabuto or Vince Guaraldi on piano? (Rank: 10)</p> <p>CHIQ-Fusion: (Rank: 5)</p>	<p>The most obvious deviation from Tjader's Latin jazz sound was Several Shades of Jade and the follow-up Breeze From the East. Both albums attempted to combine jazz and Asian music, much as Tjader and others had done with Afro-Cuban. The result was ... Other experiments were not so easily dismissed. Tjader teamed up New Yorker Eddie Palmieri in 1966 to produce El Sonido Nuevo A companion LP was recorded for Palmieri's ...</p>

Table 10: The second case for analyzing the content between the original and our enhanced history, as well as the rewritten queries generated on top of them. The indication is consistent with Table 9.