

# ADACQR: Enhancing Query Reformulation for Conversational Search via Sparse and Dense Retrieval Alignment

Yilong Lai\*, Jialong Wu\*, Congzhi Zhang, Haowen Sun, Deyu Zhou†  
School of Computer Science and Engineering, Key Laboratory of Computer Network and Information Integration, Ministry of Education, Southeast University, China  
{yilong.lai, jialongwu, zhangcongzhi, haowensun, d.zhou}@seu.edu.cn

## Abstract

Conversational Query Reformulation (CQR) has significantly advanced in addressing the challenges of conversational search, particularly those stemming from the latent user intent and the need for historical context. Recent works aimed to boost the performance of CQR through alignment. However, they are designed for one specific retrieval system, which potentially results in sub-optimal generalization. To overcome this limitation, we present a novel framework ADACQR. By aligning reformulation models with both *term*-based and *semantic*-based retrieval systems, ADACQR enhances the generalizability of information-seeking queries among diverse retrieval environments through a two-stage training strategy. Moreover, two effective approaches are proposed to obtain superior labels and diverse input candidates, boosting the efficiency and robustness of the framework. Experimental results on the TopiOCQA, QReCC and TREC CAsT datasets demonstrate that ADACQR outperforms the existing methods in a more efficient framework, offering both quantitative and qualitative improvements in conversational query reformulation.<sup>1</sup>

## 1 Introduction

Conversational search extends the traditional information retrieval paradigms by addressing complex information-seeking requirements through multi-turn interactions (Radlinski and Craswell, 2017; Qu et al., 2020; Gao et al., 2023). A fundamental challenge in conversational search is to discover the latent user intent within the current query and historical context, which complicates the application of off-the-shelf retrievers due to issues such as omissions, ambiguity, and coreference (Anantha et al., 2021; Adlakha et al., 2022).

\* Equal Contribution.

† Corresponding Author.

<sup>1</sup>  <https://github.com/init0xyz/AdaCQR>

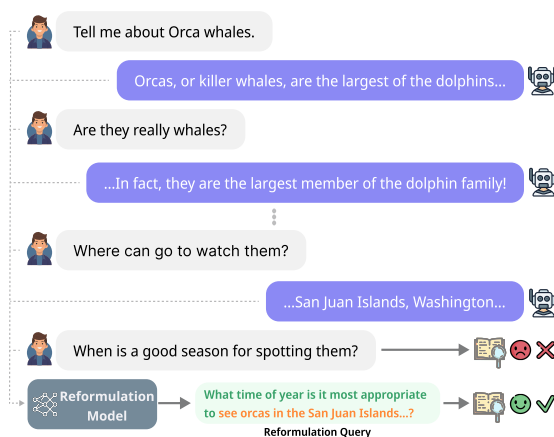


Figure 1: An example of CQR which takes the context and current query as input and generates a decontextualized query as output.

Existing methods to address this challenge can be broadly categorized into two types: dense retriever-based and query reformulation-based. For dense retriever-based approaches (Qu et al., 2020; Lin et al., 2021b; Kim and Kim, 2022; Mo et al., 2024d), long dialogue contexts can be effectively grasped by the dense retriever while incurring re-training costs and lacking the adaptability to sparse retrieval systems like BM25 (Robertson et al., 2009). Query reformulation-based approaches leverage a language model to decontextualize every user’s query into a stand-alone query, a process known as conversational query reformulation (CQR), as shown in Figure 1. Previous studies have demonstrated the effectiveness of CQR (Wu et al., 2022; Mo et al., 2023a; Ye et al., 2023).

As the training objectives are not aligned with task targets, *i.e.*, minimizing the cross-entropy loss for teacher forcing generation during training while expecting to maximize retrieval metric during inference, subsequent approaches have aimed to enhance the performance of CQR through alignment. For example, Jang et al. (2023) utilize Minimum

Bayes Risk (MBR) (Smith and Eisner, 2006) based on semantic similarity between the query and gold passage to achieve alignment. Yoon et al. (2024) create binarized comparisons based on retriever feedback and optimize the reformulation model via direct preference optimization. They also tackle the reliance on sub-optimal and costly human-annotated reformulation labels by using Large Language Models (LLMs) to generate labels via iterative prompting or multi-perspective prompting.

However, the methods mentioned above are designed for one specific retrieval system. For an information-seeking query to be generalized well across both sparse and dense retrieval systems, it must have: (1) precise *term* overlap (e.g., the presence of key entities in the query) and (2) high *semantic* similarity between the document and the query (Luan et al., 2021). Both characteristics of the information-seeking query play a crucial role in query reformulation; jointly leveraging them is beneficial. In addition, previous works leverage reinforcement learning to achieve alignment, but they exhibited stability issues and relied on an explicit reference model (Wu et al., 2022; Jang et al., 2023).

Therefore, in this paper, we introduce ADACQR, a novel framework that effectively aligns the training objective with the task target. In specific, ADACQR aligns the reformulation model and the retrievers from both *term* and *semantic* perspectives to achieve strong generalization abilities in sparse and dense retrieval. Furthermore, to address the issues of high complexity and instability inherent in MBR, we employ a two-stage training strategy to achieve alignment, where the reformulation model serves both as a generation model using cross-entropy loss for teacher forcing generation and a reference-free evaluation model using contrastive loss.

The framework works as follows: 1) A fusion metric is introduced to evaluate the generalization performance of the reformulated query across various retrieval systems. 2) By leveraging uncertainty estimation derived from the fusion metric, combined with insights from contrastive learning, we implicitly guide the large language model (LLM) to generate superior reformulation labels, which are then utilized to initialize the reformulation model during the first training stage. 3) In the second training stage, Diverse Beam Search (Vijayakumar et al., 2016) is employed to effectively gather multiple candidate reformulation queries. The reformula-

tion model is then aligned using a contrastive loss from both term and semantic perspectives, incorporating the candidates and their relative orders based on the fusion metric.

ADACQR achieves excellent performance on widely used conversation search datasets, including TopiOCQA (Adlakha et al., 2022), QReCC (Anantha et al., 2021) and TREC CAst (Dalton et al., 2020, 2021, 2022). Notably, to maintain the smoothness of the overall system, we selected a lightweight model as our backbone, which achieved performance comparable to those approaches fine-tuned on the LLaMA-7B backbone or aggregated multiple candidate queries from proprietary LLM. Experimental results demonstrate the quantitative and qualitative improvements of our proposed framework.

The contributions of this work are as follows:

- We propose ADACQR to align reformulation models from both *term* and *semantic* perspectives.
- By leveraging the proposed fusion metric, we can effectively acquire superior labels for generation and collect diverse ordered candidate queries for reference-free evaluation.
- Extensive experiments on several benchmark datasets conclusively demonstrate our proposed ADACQR significantly outperforms existing methods, establishing its superiority in performance.

## 2 Related Work

### 2.1 Conversational Search

Conversational search improves traditional information retrieval by using iterative, multi-turn interactions to address the complex information needs of users (Gao et al., 2023; Mo et al., 2024b). A key challenge is understanding the implicit intent of the user, requiring attention to both the current query and its historical context. Two main approaches to this problem are conversational dense retrieval (CDR) and conversational query reformulation (CQR).

CDR (Qu et al., 2020; Yu et al., 2021; Lin et al., 2021b) aims to improve the representation of the current query along with its historical context by training dense retrievers. Recent advancements in CDR have focused on mitigating the influence of irrelevant historical contexts (Kim and Kim, 2022; Mo et al., 2023b, 2024c,d) and enhancing interpretability (Mao et al., 2023c; Cheng et al., 2024).

However, this approach incurs additional training costs and lacks the adaptability to sparse retrieval systems like BM25 (Robertson et al., 2009).

Conversely, CQR (Elgohary et al., 2019) concentrates on decontextualizing the query of user into a stand-alone query suitable for use with off-the-shelf retrievers. Numerous prior studies have demonstrated the effectiveness of CQR by utilizing human annotations in supervised methods (Lin et al., 2020; Yu et al., 2020; Vakulenko et al., 2021) and integrating query expansion models (Mo et al., 2023a, 2024a). However, human-annotated labels are costly and reported to be sub-optimal (Lin et al., 2021b; Wu et al., 2022). In the era of LLMs, several studies have utilized LLMs to generate query reformulations directly (Ye et al., 2023; Mao et al., 2023b) and obtain reformulation labels for distillation (Jang et al., 2023; Yoon et al., 2024). This paper focuses on conversational query reformulation, proposing a novel framework ADACQR to align with *term*-based and *semantic*-based retrieval systems. To overcome the limitations of human annotation, we also developed two effective methods for obtaining superior labels and diverse input candidates.

## 2.2 Aligning LMs using Feedback

Aligning language models with feedback involves adjusting their behaviour and outputs based on evaluation feedback (Wang et al., 2023), employing various reward learning methodologies to provide accurate supervised signals (Schulman et al., 2017; Rafailov et al., 2023).

Recent studies have enhanced conversational query reformulation by aligning language models with retriever feedback (Jang et al., 2023; Yoon et al., 2024). Jang et al. (2023) achieve the alignment through minimizing Bayes Risk based on semantic similarity between the query and the gold passage. Yoon et al. (2024) leverages LLMs to generate numerous reformulations via multi-perspective prompting, creating binarized comparisons based on retriever feedback and optimizing the reformulation model using DPO (Rafailov et al., 2023). However, previous methods struggle with the high cost of generating reformulations with LLMs (Yoon et al., 2024), or the instability of MBR (Jang et al., 2023; Finkelstein and Freitag, 2023). In contrast, our framework utilizes a contrastive loss (Liu et al., 2022) to achieve alignment with retrievers. To the best of our knowledge, we are the first to employ the language model as a

reference-free evaluation model to align retrievers, thereby enhancing stability and reducing complexity.

## 3 Method

### 3.1 Task Formulation

The conversational search task discussed in this paper involves finding the passage most relevant to the intent of the user from a large collection of passages, given the current query of the user and historical context. To achieve this goal, the CQR task is to utilize a language model  $G_\theta$  to condense the current query  $q_k$  and historical context  $H_{k-1} = \{q_i, r_i\}_{i=1}^{k-1}$  into a stand-alone query  $\hat{Q}_k$ , where  $q_i$  and  $r_i$  denote the query and system answer of the  $i$ -th turn conversation, with  $k$  indicating the current turn. This decontextualized query  $\hat{Q}_k$  is subsequently input into an *off-the-shelf* retrieval system  $\mathcal{R}$ , which returns a ranked list of the top- $p$  relevant passages.

For sake of convenience, we formalize the CQR task as a session  $\mathcal{P} = \{q, H\}$ , where  $q$  represents the user’s current query and  $H$  denotes the historical context. The objective is to generate a reformulated query  $\hat{Q}$ , as discussed in the following sections.

### 3.2 Overall Framework

Figure 2 depicts the overall framework of ADACQR. The framework begins to introduce a fusion metric to evaluate the generalization performance of the reformulation queries across sparse and dense retrieval systems (§3.3). Leveraging the uncertainty estimation based on the fusion metric and the insight from contrastive learning, we select representative examples and implicitly guide LLM to generate superior labels  $Q^*$  to meet the needs of retrievers (§3.4). Subsequently, we employ a two-stage training strategy to align the reformulation model with the retrievers. In the first stage, we train the reformulation model with a cross-entropy loss  $\mathcal{L}_g$  using the superior  $Q^*$  to acquire the basic ability to generate reformulation queries (§3.5.1). Afterwards, the model is used to create a diverse set  $\mathcal{S}$  comprising candidates queries  $C_{(1)}, \dots, C_{(n)}$  (§3.5.2), which are then evaluated across sparse and dense retrieval, considering from term and semantic perspectives, and ranked based on a proposed fusion metric that synthesizes these evaluations. In the second stage, leveraging the relative order of the candidates, we apply a contrastive loss  $\mathcal{L}_c$  (§3.5.3)




Figure 2: The framework of the proposed ADACQR. A two-stage training is employed, where Stage 1 involves minimizing generation loss  $\mathcal{L}_g$ , followed by Stage 2 employing contrastive loss  $\mathcal{L}_c$ . The evaluation score is a distribution vector defined in Eq. (6).

to achieve alignment between the reformulation model and the retrievers, where the reformulation model acts as an evaluation model.

### 3.3 Fusion Metric for Sparse and Dense Retrieval

A good information-seeking query must have precise term overlap and high semantic similarity between the document and the query to generalize well across sparse and dense retrieval (Luan et al., 2021).

In sparse retrieval, the query is tokenized into terms and matches passage based on *term overlap*. In contrast, dense retrieval converts the query into an embedding by the encoder and searches passages based on *semantic similarity*. To measure the generalization ability of the reformulation queries, we input them into sparse and dense retrieval systems and assess their performance based on the ranking of the corresponding gold passages, as illustrated in the central part of Figure 2.

Inspired by reciprocal rank fusion (Cormack et al., 2009), we propose a fusion metric to merge the score of sparse retrieval based on term overlap and dense retrieval based on semantic similarity into an optimized score:

$$M(\hat{Q}, d) = \frac{1}{r_s(\hat{Q}, d)} + \frac{1}{r_d(\hat{Q}, d)} \quad (1)$$

where  $\hat{Q}$  is a reformulation query,  $d$  is the gold passage.  $r_s(q, d)$  and  $r_d(q, d)$  represent the rank of the gold passage  $d$  within the sparse and dense retrieval results for query  $q$ , respectively. The ranking  $r_s$  and  $r_d$  starts from 1, indicating the highest-ranked passage.

Leveraging Eq (1), the generalization performance of the reformulation query  $\hat{Q}$  could be evaluated, where a larger  $M(\hat{Q}, d)$  indicates better generalization performance for reformulation query  $\hat{Q}$  on sparse and dense retrieval systems.

### 3.4 Superior Reformulation Annotation

We leverage LLMs to generate high-quality reformulation labels by conveying the characteristics of effective query reformulation for retrieval. Given the challenge of explicitly defining optimal reformulation, we propose a prompting strategy that selects representative examples and implicitly guides LLMs to generate reformulation labels based on contrastive learning.

It begins with a vanilla model  $G_\pi$  with basic query reformulation ability. We employ  $G_\pi$  to generate a reformulation candidates set  $\mathcal{S}_\pi$  with diverse beam search for each reformulation session. Motivated by previous works showing that the annotated most uncertain examples can significantly enhance the effectiveness of in-context learning (Diao et al., 2024; Yue et al., 2024), the **variance**  $\text{Var}(\mathcal{S}_\pi)$ , is utilized as a measure for estimating uncertainty in reformulation tasks, where scores for  $\mathcal{S}_\pi$  are computed using the fusion metric described in Eq (1). A higher variance suggests greater instability in the performance of the retriever, indicating a more challenging reformulation problem. We then selected the top- $m$  representative reformulation problems exhibiting the highest variances on the validation set.

For the representative example annotation, inspired by contrastive learning (Paranjape et al.,

2021; He et al., 2022), we identified the best and worst reformulation candidates from set  $\mathcal{S}_\pi$  based on Eq (1), denoted as  $C_{\text{best}}$  and  $C_{\text{worst}}$  to implicitly guide the LLMs in generating labels aligned with the needs of the retrieval system. We then concatenate the  $m$  representative demonstrations, each consists of  $(q, H, C_{\text{best}}, C_{\text{worst}})$ , along with task instruction  $\mathcal{I}$ . Finally, we employed the LLM to obtain the superior reformulation labels  $Q^*$  through in-context learning (Brown et al., 2020; Dong et al., 2022; Xiang et al., 2024). The details of the annotation are presented in Appendix D.

### 3.5 Align LMs with Retrievers

After getting superior reformulation labels using a defined fusion metric, we can align LMs with retrievers through two-stage training. The reformulation model serves as a standard generation model at the training stage 1. (§3.5.1) Then we develop a method to generate multiple candidate queries using this trained model. (§3.5.2) By learning the relative order of these candidates, we implicitly guide the language model to generate queries that meet the requirements of the retrievers. Lastly, in training stage 2, the reformulation model serves both as a generation model using cross-entropy loss and a reference-free evaluation model using contrastive loss to achieve alignment. (§3.5.3)

#### 3.5.1 Training Stage 1 for Initialization

In the first training stage, we train a language model using the superior reformulation labels to endow it with the basic capability of query reformulation. To encourage more diverse generation results, a label smooth cross-entropy loss is used:

$$\mathcal{L}_1 = \mathcal{L}_g = \sum_{j=1}^l \sum_x p_s(x | \mathcal{P}, Q_{<j}^*) \log p_{G_\theta}(x | \mathcal{P}, Q_{<j}^*; \theta) \quad (2)$$

where  $\mathcal{P}$  is the reformulation session including current query  $q$  and historical context  $H$ ,  $Q_{<j}^*$  is the first  $j$  tokens of the reformulation label  $Q^*$ .  $p_s$  is a label smooth distribution, defined as follows:

$$p_s(x | \mathcal{P}, Q_{<j}^*) = \begin{cases} 1 - \beta & x = Q_j^* \\ \frac{\beta}{N-1} & x \neq Q_j^* \end{cases} \quad (3)$$

where  $\beta$  is the probability mass parameter, and  $N$  is the size of the dictionary. Now we have a trained language model  $G_\theta$  using cross-entropy loss, which can be used for candidate generation and serves as a reference-free evaluation model during the training at stage 2.

#### 3.5.2 Candidates Generation for Alignment

To efficiently generate a variety of candidates, we utilized Diverse Beam Search (Vijayakumar et al., 2016), an extension of the beam search strategy designed to generate a more diverse set of beam sequences for selection. Formally, given trained language model  $G_\theta$  and reformulation problem  $\mathcal{P}$ , we generate candidates set  $\mathcal{S} = \{C_{(1)}, \dots, C_{(n)}\}$  with diverse beam search, where  $C_{(i)}$  is the candidate of reformulation query,  $n$  is the number of candidates.

To align the retrievers from both term-based and semantic-based perspectives with the language model, we define the relative rank order as implicitly supervised signals, utilizing the metric proposed in Eq.(1), which simultaneously considers both types of retrievers, as follows:

$$C_{(i)} \succ C_{(j)} \iff M(C_{(i)}, d) > M(C_{(j)}, d) \quad (4)$$

where  $d$  is the gold passage of reformulation problem  $\mathcal{P}$ .

For reformulation problem  $\mathcal{P}$ , we now have candidates set  $\mathcal{S} = \{C_1, \dots, C_n\}$  and their relative rank order  $C_1 \succ C_2 \succ \dots \succ C_n$ , where  $C_i$  represents the  $i$ -th candidate in the sorted order.

#### 3.5.3 Training Stage 2 for Alignment

Now we have sorted candidates  $\mathcal{S}$  and trained model  $G_\theta$  to perform training at stage 2. Leveraging the candidates set  $\mathcal{S}$  and their relative rank order  $C_1 \succ C_2 \succ \dots \succ C_n$ , a contrastive loss (Liu et al., 2022) for alignment:

$$\mathcal{L}_c = \sum_{i=1}^n \sum_{j>i} \max(0, f(C_j) - f(C_i) + (j-i) \times \lambda) \quad (5)$$

where  $j$  and  $i$  are the rank order in the candidates, and  $\lambda$  is the margin parameter.  $f(C)$  represents the length-normalized estimated log-probability, where the language model serves as a reference-free evaluation model:

$$f(C) = \frac{1}{|C|^\alpha} \sum_{t=1}^l \log p_{G_\theta}(c_t | \mathcal{P}, C_{<t}) \quad (6)$$

where  $|C|$  and  $l$  is the length of candidate,  $c_t$  is the generated  $t$ -th token given reformulation problem and previous  $t - 1$  tokens, and  $\alpha$  is the length penalty parameter.

To ensure the stability of the training process, we employed a multi-task learning loss function,

Type	Query Reform.	TopiOCQA				QReCC			
		MRR	NDCG	R@10	R@100	MRR	NDCG	R@10	R@100
Sparse (BM25)	T5QR (T5-base)	11.3	9.8	22.1	44.7	33.4	30.2	53.8	86.1
	CONQRR (T5-base)	-	-	-	-	38.3	-	60.1	88.9
	EDIRCS (T5-base)	-	-	-	-	41.2	-	62.7	<u>90.2</u>
	IterCQR (T5-base)	16.5	14.9	29.3	54.1	46.7	44.1	64.4	85.5
	LLM-Aided (ChatGPT)	-	-	-	-	49.4	46.5	67.1	88.2
	LLM4CS-REW(ChatGPT)	<u>16.8</u>	<u>15.0</u>	<u>31.3</u>	<u>56.7</u>	35.2	32.0	55.9	84.7
	ADACQR ( <i>Ours</i> , T5-base)	<b>17.8</b>	<b>15.8</b>	<b>34.1</b>	<b>62.1</b>	<b>52.4</b>	<b>49.9</b>	<b>70.9</b>	<b>91.0</b>
	ConvGQR (T5-base) <sup>♡</sup>	12.4	10.7	23.8	45.6	44.1	41.0	64.4	88.0
	RETPO (LLaMA2-7B) <sup>♡</sup>	<b>28.3</b>	<b>26.5</b>	48.3	<b>73.1</b>	50.0	47.3	69.5	89.5
	LLM4CS-RAR(ChatGPT) <sup>♡</sup>	<u>27.9</u>	<u>26.4</u>	<u>48.4</u>	71.1	<u>51.6</u>	<u>49.3</u>	<u>75.3</u>	<u>92.6</u>
ADACQR+Expansion <sup>♡</sup>	<b>28.3</b>	<b>26.5</b>	<b>48.9</b>	<u>71.2</u>	<b>55.1</b>	<b>52.5</b>	<b>76.5</b>	<b>93.7</b>	
Dense (ANCE)	T5QR (T5-base)	23.0	22.2	37.6	54.4	34.5	31.8	53.1	72.8
	CONQRR (T5-base) <sup>‡</sup>	-	-	-	-	41.8	-	65.1	<u>84.7</u>
	EDIRCS (T5-base)	-	-	-	-	42.1	-	65.6	<b>85.3</b>
	IterCQR (T5-base)	26.3	25.1	42.6	62.0	42.9	40.2	65.5	84.1
	LLM-Aided (ChatGPT)	-	-	-	-	<u>43.5</u>	<u>41.3</u>	<u>65.6</u>	82.3
	LLM4CS-REW (ChatGPT)	30.3	29.0	49.9	<u>67.7</u>	36.8	34.0	56.1	74.2
	ADACQR ( <i>Ours</i> , T5-base)	<b>32.8</b>	<b>31.5</b>	<b>54.6</b>	<b>73.0</b>	<b>45.1</b>	<b>42.4</b>	<b>66.3</b>	83.4
	ConvGQR (T5-base) <sup>♡</sup>	25.6	24.3	41.8	58.8	42.0	39.1	63.5	81.8
	RETPO (LLaMA2-7B) <sup>♡</sup>	30.0	28.9	49.6	68.7	44.0	41.1	66.7	<b>84.6</b>
	LLM4CS-RAR (ChatGPT) <sup>♡</sup>	35.4	34.4	<u>55.2</u>	<u>72.2</u>	44.7	41.8	67.2	84.0
ADACQR+Expansion <sup>♡</sup>	<b>38.5</b>	<b>37.6</b>	<b>58.4</b>	<b>75.0</b>	<b>45.8</b>	<b>42.9</b>	<b>67.3</b>	83.8	

Table 1: Evaluation results of various retrieval system types on the test sets of QReCC and TopiOCQA. The best results among all methods with similar settings are **bolded**, and the second-best results are underlined. ♡ denotes the method involved in using query expansion. ‡ denotes the baselines utilizing another dual encoder dense retrieval. +Expansion denotes the addition of query expansion, details in Appendix F.

where the language model served as both a generation model and an evaluation model:

$$\mathcal{L}_2 = \mathcal{L}_g + \gamma \mathcal{L}_c \quad (7)$$

where  $\gamma$  is the weight of the contrastive loss.

## 4 Experiments

**Datasets** We train and evaluate our model using two widely utilized conversational search datasets: QReCC (Anantha et al., 2021) and TopiOCQA (Adlakha et al., 2022). We also conduct zero-shot experiments on TREC CAsT 19-21 (Dalton et al., 2020, 2021, 2022). The details of these datasets are shown in Appendix B.2.

**Retrieval Systems** Following prior works in the CQR task (Wu et al., 2022; Mo et al., 2023a; Jang et al., 2023; Yoon et al., 2024), we evaluate ADACQR using sparse and dense retrieval systems. The sparse retrieval system used is BM25 (Robertson et al., 2009). For dense retrieval, we use ANCE (Xiong et al., 2020), trained on the MS MARCO (Nguyen et al., 2016) retrieval task.

**Baselines** To compare with previous CQR baselines, we define two variants of our framework:

- ADACQR generates the reformulation queries through the aligned T5-base model, as illustrated in Section 3.5.
- ADACQR+Expansion concats the reformulation queries generated by ADACQR and query expansions generated by vanilla LLaMA2-7B leveraging the pseudo answers and keywords expansion techniques as described in the Appendix F.

We categorize previous works into two groups: *query reformulation* and *query reformulation with expansion*<sup>♡</sup>. In the query reformulation section, we compare ADACQR with the fine-tuned T5-base models including: T5QR (Lin et al., 2020), CONQRR (Wu et al., 2022), EDIRCS (Mao et al., 2023a), IterCQR (Jang et al., 2023) and LLM-based prompting methods including LLM-Aided (Ye et al., 2023) and LLM4CS (Mao et al., 2023b) under the Rewriting Prompting(REW) setting. In the query reformulation with expansion section, we compare ADACQR+Expansion with T5-based model ConvGQR<sup>2</sup> (Mo et al., 2023a),

<sup>2</sup>ADACQR (T5-base) without expansion, is showing superior performance than ConvGQR<sup>♡</sup> whose expansions are generated by another T5-base model.

Query Reform.	CAST-19			CAST-20			CAST-21		
	MRR	R@10	R@100	MRR	R@10	R@100	MRR	R@10	R@100
T5QR(T5-base)	70.1	-	33.2	42.3	-	35.3	46.9	-	40.8
EDIRCS(T5-base)	70.9	-	35.3	43.8	-	37.5	-	-	-
LLM4CS-REW(ChatGPT)	65.6	10.5	33.4	<u>49.1</u>	<b>16.9</b>	<b>41.3</b>	<u>57.7</u>	<b>21.9</b>	<b>55.3</b>
ADACQR( <i>Ours</i> , T5-base)	<b>71.6</b>	<b>12.1</b>	<b>36.7</b>	<b>51.4</b>	<u>15.7</u>	<u>40.9</u>	<b>58.3</b>	<u>21.3</u>	<u>54.5</u>
CONVGQR(T5-base) <sup>♥</sup>	70.8	-	33.6	46.5	-	36.8	43.3	-	33.0
LLM4CS-RAR(ChatGPT) <sup>♥</sup>	70.6	12.6	37.7	<b>57.3</b>	<b>19.8</b>	<b>47.6</b>	<b>65.8</b>	<u>24.6</u>	<b>59.5</b>
ADACQR+Expansion <sup>♥</sup>	<b>74.5</b>	<b>13.8</b>	<b>39.2</b>	<u>56.6</u>	<u>19.2</u>	<u>45.6</u>	<u>64.2</u>	<b>25.0</b>	<u>58.7</u>

Table 2: Zero-shot experiment results on TREC CAST 19-21 datasets. The best results among all methods with similar settings are **bolded**, and the second-best results are underlined.

the fine-tuned LLaMA2-7B model RETPO (Yoon et al., 2024) and LLM-based prompting method LLM4CS (Mao et al., 2023b) with the Rewrite-and-Resonse(RAR) setting.

The details regarding **baselines, implementation and evaluation metrics** are provided in Appendix B.1, Appendix C and Appendix B.3, respectively.

#### 4.1 Main Results

To evaluate the efficacy of our framework, we conducted comprehensive experiments datasets with the ADACQR and baselines, presented in Table 1. We consider three kinds of backbones as baselines: the T5-based, the LLaMA2-7B-based, and the ChatGPT-based.

The results under the without expansion setting demonstrate that ADACQR significantly outperforms previous models utilizing T5-base as the backbone. Among methods with expansion, the exceptional performance of RETPO and LLM4CS-RAR, which use LLaMA2-7B and ChatGPT as the backbone, respectively, can be attributed to the inherently strong common-sense reasoning capabilities of backbone models. ADACQR with expansion achieves results comparable to RETPO and LLM4CS-RAR, with only a slight disadvantage in the R@100 metric, while empirically outperforming them in other settings. Specifically, ADACQR+Expansion shows superior performance in dense retrieval on the TopiOCQA, attaining the best MRR (38.5), NDCG (37.6), R@10 (58.4), and R@100 (75.0). The reported results show significant improvements with the  $t$ -test at  $p < 0.05$  overall compared to baselines in all the settings on QReCC and TopiOCQA (except ADACQR+Expansion with ANCE on TopiOCQA).

These results underscore the efficacy and generalizability of ADACQR in enhancing retrieval

Type	Ablation Variants	MRR	R@10
Sparse	ADACQR ( <i>Ours</i> )	<b>52.4</b>	<b>70.9</b>
	w/o. Contrastive Loss	43.3	62.8
	w/o. Fusion Metric	50.5	67.7
	w/o. Sparse Rank	50.9	69.7
	w/o. Dense Rank	51.6	70.5
Dense	ADACQR ( <i>Ours</i> )	<b>45.1</b>	<b>66.3</b>
	w/o. Contrastive Loss	38.5	58.9
	w/o. Fusion Metric	42.4	63.7
	w/o. Sparse Rank	43.5	64.2
	w/o. Dense Rank	42.9	63.0

Table 3: Ablation study for alignment and ranking of ADACQR on QReCC dataset.

performance across different retrieval systems.

#### 4.2 Zero-shot Results

To access the generalization performance of ADACQR, we conduct the experiment on TREC CAST datasets under a zero-shot setting, shown in Table 2. Among all non-expansion methods, ADACQR achieves the best performance across all metrics on the CAST-19 and the highest MRR scores on the CAST-20 and CAST-21. Among methods with expansion, ADACQR still performs best on the CAST-19 dataset and demonstrates comparable results to the LLM4CS-RAR approach on the CAST-20 and CAST-21. The strongest competitor in expansion setting, LLM4CS-RAR, leverages advanced proprietary LLM and candidates queries aggregation to enhance its expansion capability, so it is reasonable that ADACQR falls slightly short. Achieving comparable performance further demonstrates the strength of our method. These results provide solid evidence of the generalization capabilities of our approach.

#### 4.3 Ablation Study

To investigate the impact of each component on the performance of ADACQR, we conducted abla-

Type	Query Reform.	MRR	R@10
Sparse	Human Rewrite	39.8	62.7
	LLM Rewrite $Q^*$	45.4	65.5
	ADACQR		
	Trained on #HR	44.9	63.7
	Trained on $Q^*$	<b>52.4</b>	<b>70.9</b>
Dense	Human Rewrite	38.4	58.6
	LLM Rewrite $Q^*$	40.1	60.2
	ADACQR		
	Trained on #HR	41.0	60.5
	Trained on $Q^*$	<b>45.1</b>	<b>66.3</b>

Table 4: Ablation study for reformulation labels on QReCC dataset. The superior LLM reformulation  $Q^*$  is obtained by in-context learning (detailed in §3.4).

tion experiments focusing on the below modules in Table 3 and Table 4.

**Alignment** The training Stage of ADACQR incorporates a contrastive loss to align the retrievers. To assess the influence of contrastive loss, we executed a single-stage training process without alignment. The removal of contrastive loss results in the most notable decline in performance, with a 9.1% decrease in MRR metric.

**Ranking** We introduce a fusion metric to evaluate query performance across semantic and term perspectives. To determine the effect of the fusion metric, we substitute it with the metric used in previous work (Jang et al., 2023), which ranks candidates solely based on the *cosine similarity* between the candidate query and the gold passage. The performance degradation confirms the effectiveness of using signals from the ranking of both types of retrievers. To further investigate the effectiveness of considering both perspectives in the fusion metric, we separately remove sparse ranking  $r_s$  and dense ranking  $r_d$  within it for analysis. Removing any rank degrades performance for both retrievers, more significantly for the corresponding retriever. This confirms the rationale behind considering both perspectives simultaneously.

**Reformulation Labels** The performance of directly using human label and  $Q^*$  as the queries is reported in Table 4. It is worth noting that superior labels  $Q^*$  outperform human labels and have strong performance both in sparse and dense retrievals, which validates the effectiveness of the proposed fusion metric and the annotation method. In the Training Stage 1 of ADACQR, we use  $Q^*$  as the golden labels. We also experiment by replacing it with human labels, and the resulting performance

drop further validates the effectiveness of  $Q^*$ . Additionally, the final results of ADACQR after both Training Stage 1 and Stage 2 show significant improvement over the original queries (whether  $Q^*$  or human labels), demonstrating the effectiveness of the Stage 2 alignment.

#### 4.4 Performance against CDR Methods

Conversational Dense Retrieval (CDR) is an orthogonal method for conversational query reformulation in conversational search, which trains dense retrievers to improve the representation of the current query and historical context. Although not directly comparable, we still present a performance comparison of the ADACQR and CDR methods on the QReCC, TopiOCQA, and CAsT datasets, as shown in Table 5.

We compare our approach against four baseline models: Conv-ANCE (Xiong et al., 2020), InstructoR-ANCE (Jin et al., 2023), Conv-SPLADE (Formal et al., 2021), and LeCoRE (Mao et al., 2023c). Conv-ANCE trains a dense retriever by enhancing session representation using a conventional contrastive ranking loss. InstructoR-ANCE leverages large language models (LLMs) to predict the relevance score between sessions and passages, followed by retriever training. Conv-SPLADE fine-tunes a strong lexical-based retriever on conversational search data, utilizing ranking loss. LeCoRE extends the SPLADE model with multi-level denoising techniques to enhance lexical session representation. ADACQR achieves the best average performance across four datasets, demonstrating superiority in both effectiveness and generalizability. Our approach employs the *off-the-shelf* ANCE retriever, and our rewrite method is **orthogonal** to the aforementioned retrievers, leaving the exploration of combining these methods for achieving higher performance in future research. Moreover, while CDR methods focus on dense representations, there are scenarios where the sparse representation of BM25 demonstrates a retrieval advantage. Our approach ADACQR takes both types of retrievers into account and enhances performance through a rewrite strategy.

## 5 Analysis

### 5.1 Analysis of the Aligned Query

To evaluate the effectiveness of the aligned reformulation queries, we analyse the reformulation queries across the first 5 epochs during Stage 2



Method	TopiOCQA	QReCC	CAsT-20	CAsT-21	Avg.
Conv-ANCE (Xiong et al., 2020)	22.9	47.1	42.2	52.3	41.1
InstructoR-ANCE (Jin et al., 2023)	25.3	43.5	43.7	53.0	41.4
Conv-SPLADE (Formal et al., 2021)	30.7	50.0	36.9	47.9	41.4
LeCoRE (Mao et al., 2023c)	32.0	51.1	37.7	50.8	42.9
ADACQR( <i>Ours</i> )	32.8	45.1	51.4	58.3	46.9
ADACQR+Expansion( <i>Ours</i> )	38.5	45.8	56.6	64.2	<b>51.3</b>

Table 5: MRR performance comparison of the ADACQR and CDR methods.




Figure 3: Analysis of the aligned reformulation query across epochs in Stage 2 training, focusing on the term overlap (DICE coefficient) and semantic similarity with the gold passage (cosine similarity).

training in Figure 3. We conduct analyses focusing on the average term overlap and semantic similarity between the queries and the gold passages. The DICE Coefficient (Dice, 1945) is utilized to assess term overlap, while cosine similarity is employed to measure semantic similarity. This analysis indicates that both term overlap and semantic similarity between the reformulated queries and the gold passages exhibit an increasing trend with each epoch in Stage 2, demonstrating the effectiveness of our method in considering both perspectives.

## 6 Conclusion

In this paper, to achieve alignment between the reformulation model and both term and semantic retrieval systems, ADACQR is proposed to enhance the generalizability of information-seeking queries through a two-stage training strategy. By leveraging a fusion metric that assesses the generalization performance across different retrieval systems, we can effectively obtain superior labels for the generation and gather a diverse set of ordered candidate queries for reference-free evaluation. Extensive experiments on five datasets demonstrate the superiority of ADACQR, achieving performance comparable to methods using the fine-tuned LLaMA2-7B and proprietary LLM.

## Limitations

Although ADACQR demonstrates remarkable performance in experimental evaluations, it also has several limitations.

During the ADACQR training process, we leverage ChatGPT for superior reformulation label annotation, and our annotation prompt requires training a basic model, which incurs additional costs and training expenses. Furthermore, due to budget constraints, we do not use more powerful LLMs, such as GPT-4, to obtain reformulation labels, although it is obvious that employing a more powerful LLM would yield better reformulation labels.

Although no further costs are introduced during reformulation model inference, aligning AdaCQR with retrievers introduces additional training time. Furthermore, generating the ordered candidate set for alignment demands extra retrieval time and increased storage capacity.

## Acknowledgments

The authors would like to thank the anonymous reviewers for their insightful comments. This work is funded by the National Natural Science Foundation of China (Grant No.62176053). This work is supported by the Big Data Computing Center of Southeast University.

## References

- Vaibhav Adlakha, Shehzaad Dhuliawala, Kaheer Suleman, Harm de Vries, and Siva Reddy. 2022. [TopiOCQA: 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, Online. Association for Computational Linguistics.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Yiruo Cheng, Kelong Mao, and Zhicheng Dou. 2024. [Interpreting conversational dense retrieval by rewriting-enhanced inversion of session embedding](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2879–2893, Bangkok, Thailand. Association for Computational Linguistics.
- Gordon V Cormack, Charles LA Clarke, and Stefan Buettcher. 2009. Reciprocal rank fusion outperforms condorcet and individual rank learning methods. In *Proceedings of the 32nd international ACM SIGIR conference on Research and development in information retrieval*, pages 758–759.
- Jeffrey Dalton, Chenyan Xiong, and Jamie Callan. 2020. Cast 2019: The conversational assistance track overview. In *In Proceedings of TREC*.
- Jeffrey Dalton, Chenyan Xiong, and Jamie Callan. 2021. Cast 2020: The conversational assistance track overview. In *In Proceedings of TREC*.
- Jeffrey Dalton, Chenyan Xiong, and Jamie Callan. 2022. Trec cast 2021: The conversational assistance track overview. In *In Proceedings of TREC*.
- Shizhe Diao, Pengcheng Wang, Yong Lin, Rui Pan, Xiang Liu, and Tong Zhang. 2024. [Active prompting with chain-of-thought for large language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1330–1350, Bangkok, Thailand. Association for Computational Linguistics.
- Lee R Dice. 1945. Measures of the amount of ecologic association between species. *Ecology*, 26(3):297–302.
- Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Zhiyong Wu, Baobao Chang, Xu Sun, Jingjing Xu, and Zhifang Sui. 2022. A survey on in-context learning. *arXiv preprint arXiv:2301.00234*.
- 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, Hong Kong, China. Association for Computational Linguistics.
- Mara Finkelstein and Markus Freitag. 2023. Mbr and qe finetuning: Training-time distillation of the best and most expensive decoding methods. In *The Twelfth International Conference on Learning Representations*.
- Thibault Formal, Benjamin Piwowarski, and Stéphane Clinchant. 2021. [Splade: Sparse lexical and expansion model for first stage ranking](#). In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '21*, page 2288–2292, New York, NY, USA. Association for Computing Machinery.
- Jianfeng Gao, Chenyan Xiong, Paul Bennett, and Nick Craswell. 2023. *Neural approaches to conversational information retrieval*, volume 44. Springer Nature.
- Xuanli He, Islam Nassar, Jamie Kiros, Gholamreza Haffari, and Mohammad Norouzi. 2022. [Generate, annotate, and learn: NLP with synthetic text](#). *Transactions of the Association for Computational Linguistics*, 10:826–842.
- Rolf Jagerman, Honglei Zhuang, Zhen Qin, Xuanhui Wang, and Michael Bendersky. 2023. Query expansion by prompting large language models. *arXiv preprint arXiv:2305.03653*.
- 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*.
- 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, Singapore. Association for Computational Linguistics.
- Jeff Johnson, Matthijs Douze, and Hervé Jégou. 2019. Billion-scale similarity search with GPUs. *IEEE Transactions on Big Data*, 7(3):535–547.
- Sungdong Kim and Gangwook 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, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles*.
- 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 Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR 2021)*, 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, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- 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*.
- Yixin Liu, Pengfei Liu, Dragomir Radev, and Graham Neubig. 2022. [BRIO: Bringing order to abstractive summarization](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2890–2903, Dublin, Ireland. Association for Computational Linguistics.
- Ilya Loshchilov and Frank Hutter. 2018. Decoupled weight decay regularization. In *International Conference on Learning Representations*.
- Yi Luan, Jacob Eisenstein, Kristina Toutanova, and Michael Collins. 2021. [Sparse, dense, and attentional representations for text retrieval](#). *Transactions of the Association for Computational Linguistics*, 9:329–345.
- Kelong Mao, Zhicheng Dou, Bang Liu, Hongjin Qian, Fengran Mo, Xiangli Wu, Xiaohua Cheng, and Zhao Cao. 2023a. [Search-oriented conversational query editing](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 4160–4172, Toronto, Canada. Association for Computational Linguistics.
- Kelong Mao, Zhicheng Dou, Fengran Mo, Jiewen Hou, Haonan Chen, and Hongjin Qian. 2023b. [Large language models know your contextual search intent: A prompting framework for conversational search](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 1211–1225, Singapore. Association for Computational Linguistics.
- 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, WWW '23*, page 3193–3202, New York, NY, USA. Association for Computing Machinery.
- Fengran Mo, Abbas Ghaddar, Kelong Mao, Mehdi Rezagholizadeh, Boxing Chen, Qun Liu, and Jian-Yun Nie. 2024a. [CHIQ: Contextual history enhancement for improving query rewriting in conversational search](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 2253–2268, Miami, Florida, USA. Association for Computational Linguistics.
- Fengran Mo, Kelong Mao, Ziliang Zhao, Hongjin Qian, Haonan Chen, Yiruo Cheng, Xiaoxi Li, Yutao Zhu, Zhicheng Dou, and Jian-Yun Nie. 2024b. [A survey of conversational search](#). *Preprint*, arXiv:2410.15576.
- Fengran Mo, Kelong Mao, Yutao Zhu, Yihong Wu, Kaiyu Huang, and Jian-Yun Nie. 2023a. [ConvGQR: Generative query reformulation for conversational search](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4998–5012, Toronto, Canada. Association for Computational Linguistics.
- 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 *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, KDD '23*, page 1722–1732, New York, NY, USA. Association for Computing Machinery.
- Fengran Mo, Chen Qu, Kelong Mao, Yihong Wu, Zhan Su, Kaiyu Huang, and Jian-Yun Nie. 2024c. [Aligning query representation with rewritten query and relevance judgments in conversational search](#). In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management, CIKM '24*, page 1700–1710, New York, NY, USA. Association for Computing Machinery.
- Fengran Mo, Chen Qu, Kelong Mao, Tianyu Zhu, Zhan Su, Kaiyu Huang, and Jian-Yun Nie. 2024d. [History-aware conversational dense retrieval](#). In *Findings of the Association for Computational Linguistics ACL 2024*, pages 13366–13378, Bangkok, Thailand and virtual meeting. Association for Computational Linguistics.
- Tri Nguyen, Mir Rosenberg, Xia Song, Jianfeng Gao, Saurabh Tiwary, Rangan Majumder, and Li Deng. 2016. Ms marco: A human-generated machine reading comprehension dataset.
- OpenAI. 2022. Introducing chatgpt. <https://openai.com/blog/chatgpt>. Accessed: 2024-02-06.
- Bhargavi Paranjape, Julian Michael, Marjan Ghazvininejad, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2021. [Prompting contrastive explanations for commonsense reasoning tasks](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 4179–4192, Online. Association for Computational Linguistics.
- 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.
- Filip Radlinski and Nick Craswell. 2017. [A theoretical framework for conversational search](#). In *Proceedings of the 2017 Conference on Conference Human Information Interaction and Retrieval, CHIIR '17*, page 117–126, New York, NY, USA. Association for Computing Machinery.

- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. 2023. [Direct preference optimization: Your language model is secretly a reward model](#). In *Advances in Neural Information Processing Systems*, volume 36, pages 53728–53741. Curran Associates, Inc.
- 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.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- David A. Smith and Jason Eisner. 2006. [Minimum risk annealing for training log-linear models](#). In *Proceedings of the COLING/ACL 2006 Main Conference Poster Sessions*, pages 787–794, Sydney, Australia. Association for Computational Linguistics.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv: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.
- Ashwin K Vijayakumar, Michael Cogswell, Ramprasath R Selvaraju, Qing Sun, Stefan Lee, David Crandall, and Dhruv Batra. 2016. Diverse beam search: Decoding diverse solutions from neural sequence models. *arXiv preprint arXiv:1610.02424*.
- Yufei Wang, Wanjun Zhong, Liangyou Li, Fei Mi, Xingshan Zeng, Wenyong Huang, Lifeng Shang, Xin Jiang, and Qun Liu. 2023. Aligning large language models with human: A survey. *arXiv preprint arXiv:2307.12966*.
- Zequ Wu, Yi Luan, Hannah Rashkin, David Reitter, 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*, pages 10000–10014, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Yanzheng Xiang, Hanqi Yan, Lin Gui, and Yulan He. 2024. [Addressing order sensitivity of in-context demonstration examples in causal language models](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 6467–6481, Bangkok, Thailand. Association for Computational Linguistics.
- 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, Singapore. Association for Computational Linguistics.
- 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, SIGIR ’21*, page 829–838, New York, NY, USA. Association for Computing Machinery.
- Murong Yue, Jie Zhao, Min Zhang, Liang Du, and Ziyu Yao. 2024. [Large language model cascades with mixture of thought representations for cost-efficient reasoning](#). In *The Twelfth International Conference on Learning Representations*.

## A Discussion

### A.1 Advantage of ADACQR in Alignment

Previous works like CONQRR (Wu et al., 2022) and IterCQR (Jang et al., 2023) explored the concept of alignment between the reformulation model and the retriever. CONQRR leverages reinforcement learning to maximize the reward of the sampled rewrite  $q_s$  and the baseline rewrite  $q$ . IterCQR uses Minimum Bayes Risk (MBR) and Top-1 Selection to achieve the alignment between the reformulation model and dense retriever.

Compared with previous works, the main advantage of ADACQR is that it **no longer requires an explicit reward model during training**; the language model serves simultaneously as both the generation model and the evaluation model (the evaluation model of CONQRR is BM25 retrieval system, and the evaluation model of IterCQR is a dense passage encoder). By decoupling the model training and candidates ranking (*i.e.*, the candidates ranking is an offline process), our method achieves improved training efficiency.

Apart from this, ADACQR achieves alignment based on contrastive learning, which is **easier to implement, converges faster, and is more stable during training** compared to the reinforcement learning (RL) of CONQRR and the Minimum Bayes Risk (MBR) of IterCQR (Liu et al., 2022). Both CONQRR and IterCQR employ special techniques like min-max normalization to ensure model convergence, which increases the training complexity (Wu et al., 2022; Jang et al., 2023). Additionally, IterCQR requires two stages, MBR and Top-1 Selection, to achieve alignment in 7-10 epochs (Jang et al., 2023), whereas ADACQR achieves alignment in just 3-5 epochs.

### A.2 Effectiveness of Prompt Setting

To evaluate the effectiveness of the prompt design method proposed in Section 3.4, we applied our prompt design method for reformulation label annotation on the QReCC test set.

We compared the results with the 0-shot approach (*i.e.*, using only the Instruction and Annotated Sample parts from Table 11) and the 3-shot-random approach (*i.e.*, randomly selecting 3 examples from the validation set). The results are shown in Table 6.

Based on these results, our prompt setting significantly improves performance in both sparse and dense retrieval compared to the 0-shot and 3-shot-

Type	Prompt Setting	QReCC	
		MRR	R@10
Sparse	0-shot	36.3	54.9
	3-shot ( <b>Random</b> )	39.1	58.0
	3-shot ( <b>Representative</b> )	<b>45.4</b>	<b>65.5</b>
Dense	0-shot	34.5	52.6
	3-shot ( <b>Random</b> )	37.2	56.0
	3-shot ( <b>Representative</b> )	<b>40.1</b>	<b>60.2</b>

Table 6: The annotation results generated by ChatGPT under different prompt settings on the QReCC test set. **Random** denotes examples randomly chosen from the validation set, while **Representative** refers to select examples as described in Section 3.4.

Coefficient( $\gamma$ )	QReCC			
	MRR	NDCG	R@10	R@100
0	43.3	41.0	62.8	88.5
0.1	45.3	42.7	65.2	90.2
1	48.8	46.1	68.7	<b>91.2</b>
10	50.2	47.7	68.8	89.0
100	<b>52.4</b>	<b>49.9</b>	<b>70.9</b>	91.0
1000	49.4	46.7	68.6	90.7
$+\infty$	44.5	41.8	65.5	90.9

Table 7: ADACQR performance with different  $\gamma$  coefficients weighting of the contrastive loss in Eq. (7).  $+\infty$  indicates only using the contrastive loss. 0 indicates only using the cross-entropy loss. BM25 is used as the retriever for experiments.

random methods, showing the effectiveness of our prompt setting.

### A.3 Effect of the Multi-Task Loss

The multi-task loss defined in Eq. (7) is designed to align with retrievers by incorporating both cross-entropy loss and contrastive loss. We conducted experiments with various  $\gamma$  coefficients, as shown in Table 7. The results indicate that increasing  $\gamma$  improves the performance of ADACQR within a certain range, highlighting the crucial role of contrastive loss for alignment. However, the importance of cross-entropy loss is also evident: when  $\gamma$  is excessively high or cross-entropy loss is omitted, the performance declines. Therefore, it concludes that including cross-entropy loss is essential to prevent excessive model variation, illustrating its necessity in the design of this multi-task loss.

### A.4 Robustness to Topic Shifts in Conversation

In the conversational search task, the frequent topic changes during the dialogue pose challenges for CQR. To evaluate the robustness of ADACQR

Model	Topic-Concentrated		Topic-Shifted	
	MRR	R@10	MRR	R@10
T5QR	35.2	54.4	25.2	45.1
CONQRR	41.9	63.1	25.2	45.9
IterCQR	54.4	72.4	24.9	49.7
Human Rewrite	44.0	66.7	31.8	56.7
ADACQR	<b>66.0</b>	<b>82.4</b>	<b>34.1</b>	<b>58.3</b>

Table 8: Performance of ADACQR on topic-concentrated and topic-shifted samples on QReCC, MRR and R@10 are reported. The result is reported on the BM25 Retrieval System.

in handling topic shifts, we divided the QReCC dataset into two parts: Topic-Concentrated and Topic-Shifted. Following previous work (Jang et al., 2023), we determine whether a topic shift has occurred in the current conversation by checking if the gold passage ID associated with the current query appears in the gold passage IDs corresponding to the previous context. The results presented in Table 8 indicate that ADACQR substantially outperforms previous models in both parts of conversations. Additionally, ADACQR exceeds human rewrites in topic-shifted dialogues, showing the robustness of our approach in query reformulation when addressing topic shifting.

## B Experimental Details

### B.1 Baseline Details

In our study, we compare ADACQR with the following representative baselines in the CQR task:

- **T5QR** (Lin et al., 2020) is a vanilla baseline that utilizes the T5-base (Raffel et al., 2020) model to perform CQR tasks.
- **CONQRR** (Wu et al., 2022) aligns the T5-base reformulation model with retrievers through direct optimization using reinforcement learning.
- **ConvGQR** (Mo et al., 2023a) enhances retrieval performance by employing two fine-tuned T5-base models, one for query reformulation and the other for query expansion.
- **EDIRCS** (Mao et al., 2023a) leverages the incorporation of non-autoregressive text-selecting techniques and autoregressive tokens generation to generate reformulation queries effectively on a fine-tuned T5-base model.
- **LLM-Aided** (Ye et al., 2023) employs ChatGPT (OpenAI, 2022) to conduct query reformulation via a “rewrite-then-edit” prompting strategy.
- **IterCQR** (Jang et al., 2023) achieves the alignment of the T5-base reformulation model and

	QReCC			TopiOCQA		
	Train	Valid	Test	Train	Valid	Test
# Dialogues	10822	769	2775	3509	720	205
# Turns	62701	800	16451	44650	800	2514
# Turns with Gold	28796	800	8209	44650	800	2514

Table 9: The statistics of QReCC and TopiOCQA datasets.

dense retriever by minimizing Bayes Risk based on the semantic similarity between the query and the gold passage.

- **RETPO** (Yoon et al., 2024) utilizes large language models to generate multiple reformulations through multi-perspective prompting, creates binarized comparisons based on retriever feedback, and optimizes LLaMA2-7B (Touvron et al., 2023) using direct preference optimization (DPO) (Rafailov et al., 2023).
- **LLM4CS** (Mao et al., 2023b) is the state-of-the-art conversational query reformulation baseline, utilizing LLM prompting techniques to generate multiple reformulated queries and aggregate their embeddings. We adopt rewriting prompting (REW) and rewriting-and-response (RAR) as baselines for query reformulation and query reformulation with expansion, respectively. Our implementation of REW and RAR follows standard settings (*i.e.*, generating five queries and aggregating them) but excludes chain-of-thought (CoT) content due to the extensive human annotation required, which is not part of our baselines and impractical in real-world scenarios. It is important to note that aggregation techniques are applicable only in dense retrieval; therefore, for BM25 retrieval, we implement RAR by concatenating the query and response with the top-1 generation result.

### B.2 Datasets Details

The QReCC (Anantha et al., 2021) dataset comprises 14K conversations with 80K question-answer pairs, and we aim to retrieve the gold passage from a collection containing 54M passages. Conversely, the TopiOCQA (Adlakhia et al., 2022) dataset includes 3.9K topic-switching conversations with 51K question-answer pairs, where the passage collection is sourced from Wikipedia and contains about 20M passages. Notably, a few examples from the QReCC and TopiOCQA training sets were randomly partitioned to create respective validation sets. The details of QReCC and TopiOCQA are described in Table 9.

TREC CAsT 2019, 2020, and 2021 (Dalton et al., 2020, 2021, 2022) are datasets known for their complexity and challenges in conversational search with a zero-shot setting. The statistics of CAsT 19-21 are shown in Table 10.

	CAsT-19	CAsT-20	CAsT-21
# Dialogues	50	25	26
# Turns	479	208	239
# Collections	38M	38M	40M

Table 10: The statistics of CAsT 2019, 2020, and 2021 datasets.

### B.3 Evaluation Metrics

We evaluate ADACQR’s retrieval performance using several widely used metrics, such as Mean Reciprocal Rank (MRR), Normalized Discounted Cumulative Gain (NDCG), Recall@10, and Recall@100. MRR is a ranking quality metric that considers the position of the first relevant passage among the ranked passages. NDCG@3 evaluates the retrieval results by considering the relevance and the rank of the top three results. Recall@K measures whether the gold passage is present within the top-K results.

## C Implementation Details

All experiments are conducted on a server equipped with four Nvidia GeForce 3090 GPUs.

### C.1 ADACQR Details

For the implementation of ADACQR, we use Huggingface *transformers* library<sup>3</sup> and *Pytorch Lightning*<sup>4</sup> framework.

We use T5-base<sup>5</sup> (Raffel et al., 2020) as the backbone of ADACQR. After conducting a comprehensive grid search, we configured the number of candidates  $n = 32$ , the margin parameter  $\lambda = 0.1$ , the weight of the contrastive loss  $\gamma = 100$ , the length penalty parameter  $\alpha = 0.6$ , and the probability mass parameter in label smooth distribution  $\beta = 0.1$ . The model parameters are optimized by the AdamW optimizer (Loshchilov and Hutter, 2018).

ADACQR is trained for 10 epochs in Stage 1 with a learning rate set to  $2e-5$  and 8 epochs in

<sup>3</sup><https://github.com/huggingface/transformers>

<sup>4</sup><https://github.com/Lightning-AI/pytorch-lightning>

<sup>5</sup><https://huggingface.co/google-t5/t5-base>

Stage 2 with a learning rate adjusted to  $5e-6$ . Both stages incorporate linear learning rate schedulers with a warm-up ratio of 0.1. We employ grid search for hyperparameter selection.

The vanilla reformulation model  $G_\pi$  in Section 3.4 is trained on reformulation labels of the QReCC dataset acquired by zero-shot prompting with ChatGPT, and the prompt is shown in Appendix D. This model is trained in 10 epochs, and the learning rate is set to  $2e-5$  with a linear learning rate scheduler with a warm-up ratio of 0.1.

For candidate generation in Section 3.5.2, we used diverse beam search with a diverse penalty of 2.0. The minimum token length for generated candidates is set to 8, and the maximum token length is set to 64. For the generation of reformulation queries, we employed beam search with a beam size of 5, and the maximum token length is set to 64 for generated queries.

### C.2 Retrieval Systems Details

We implement the retrieval systems using Faiss (Johnson et al., 2019) and Pyserini (Lin et al., 2021a). For BM25, as in previous work (Mo et al., 2023a; Jang et al., 2023; Yoon et al., 2024), we set  $k_1 = 0.82$ ,  $b = 0.68$  in QReCC, and  $k_1 = 0.9$ ,  $b = 0.4$  in TopiOCQA. The  $k_1$  controls the non-linear term frequency normalization, and  $b$  is the scale of the inverse document frequency. For ANCE<sup>6</sup>, the maximum token length is set to 128 tokens for reformulation query and 384 tokens for passage.

For both sparse and dense retrieval systems, we retrieved the top 100 relevant passages for each query and obtained the result of evaluation metrics with *pytrec\_eval* (Van Gysel and de Rijke, 2018).

## D ChatGPT Annotation Details

We use gpt-3.5-turbo-0125 (OpenAI, 2022)<sup>7</sup> to obtain the initial and superior reformulation labels via zero-shot and few-shots prompting.

For initial reformulation labels of  $G_\pi$ , we use the “Instruction” and “Annotated Sample” parts shown in Table 11, *i.e.*, zero-shot.

For superior reformulation labels for ADACQR, we utilize the top-3 most challenging demonstrations (*i.e.*,  $m = 3$ ) for the QReCC dataset and the top-5 most challenging demonstrations (*i.e.*,

<sup>6</sup><https://huggingface.co/sentence-transformers/msmarco-roberta-base-ance-firstp>

<sup>7</sup><https://platform.openai.com/docs/models/gpt-3-5-turbo>

$m = 5$ ) for the TopiOCQA dataset, *i.e.*, few-shots. The final prompt is formed  $\mathcal{D} = \mathcal{I} \parallel T_1 \parallel \dots \parallel T_m$ , where  $\parallel$  donates concatenation. The detailed prompts to annotate the QReCC dataset and the TopiOCQA dataset are shown in Table 11 and Table 12, respectively.

To encourage a more deterministic output, we set the temperature to 0.1, and the seed is set to 42 for reproductivity. The total consumption to annotate QReCC and TopiOCQA datasets for initial and superior reformulation labels is about 151M tokens, which cost about 120\$.

## E Case Study

In this section, we present several examples of how ADACQR succeeded or failed on the QReCC and TopiOCQA datasets.

Table 15 demonstrates a case where ADACQR successfully retrieved the gold passage through query rewriting, whereas human rewrites failed, showing the superiority of ADACQR over human rewrites. After being written by ADACQR, the query is decontextualized, resulting in overlaps while concurrently offering more specific information. This enhanced specificity aids the retriever toward the most relevant passages effectively. Additionally, in Tables 16 and 17, we show examples of how the ADACQR and ADACQR with Expansion models successfully retrieved the gold passage. We also provide a failure case, as illustrated in Table 18.

## F Query Expansion Details

For query expansion, we leverage LLaMA2-7B-Chat<sup>8</sup> as the backbone for a fair comparison with prior work (Yoon et al., 2024). The query expansion process involves directly answering the given query (Mo et al., 2023a) and generating relevant keywords (Jagerman et al., 2023). Then the reformulation queries are concatenated with the generated answers and keywords for retrieval. The prompts employed for query expansion are presented in Table 13 and Table 14.

*vLLM* framework (Kwon et al., 2023) is used for inference, with the temperature parameter set to 0.5 and the maximum token limit set to 50 during the generation process.

---

<sup>8</sup><https://huggingface.co/meta-llama/Llama-2-7b-chat-hf>



## Prompt for QReCC Annotation

### Instruction

Given a question and its context, decontextualize the question by addressing coreference and omission issues. The resulting question should retain its original meaning and be as informative as possible, and should not duplicate any previously asked questions in the context.

### Demonstrations

Context: [Q: What was Ridley Scott's directing approach to directing? A: Russell Crowe commented about Ridley Scott's directing, I like being on Ridley's set because actors can perform and the focus is on the performers. Q: Were there others who commented about Scott's approach as a director and producer? A: Charlize Theron praised the Ridley Scott's willingness to listen to suggestions from the cast for improvements in the way their characters are portrayed on screen. Q: What was Ridley Scott's style? A: In Ridley Scott's visual style, he incorporates a detailed approach to production design and innovative, atmospheric lighting Q: How did that translate into his films? A: In his movies, Ridley Scott commonly uses slow pacing until the action sequences. Q: What popular movies did he take this approach and use this style? A: Examples of Ridley Scott's directing style include Alien and Blade Runner.]

Question: Is there anything else interesting about his style?

Good Rewrite: Is there anything else interesting about Ridley Scott's style besides his slow pacing until the action sequences?

Bad Rewrite: is there anything else interesting about Ridley Scott's directing style?

Context: [Q: What was the health issues did Bad Brains frontman H.R. have? A: On March 15, 2016, Bad Brains frontman H.R. was reportedly diagnosed with a rare type of headache called Short-lasting unilateral neuralgiform headache with conjunctival injection and tearing (SUNCT syndrome) Q: Was there anything to cure it? A: As diagnostic criteria have been indecisive and its pathophysiology remains unclear, no permanent cure is available for short-lasting unilateral neuralgiform headache with conjunctival injection and tearing (SUNCT syndrome) Q: Are there any other interesting aspects about this article? A: On November 3, 2015, Bad Brains announced on their Facebook page that Dr. Know (Gary Miller) was hospitalized and on life support, after many other musicians reported so.]

Question: What did they do in 2015?

Good Rewrite: What did Bad Brains do in 2015 after Dr. Know (Gary Miller) was hospitalized and on life support?

Bad Rewrite: What do the Bad Brains do in 2015?

<— Omit One demonstration —>

### Annotated Sample

Context: [{{current\_context}}]

Question: {{current\_query}}

Good Rewrite:

Table 11: The prompt used to obtain QReCC annotated labels.

### Instruction

Given a question and its context, decontextualize the question by addressing coreference and omission issues. The resulting question should retain its original meaning and be as informative as possible, and should not duplicate any previously asked questions in the context.

### Demonstrations

Context: [Q: what is the fallacy of the argumentum ad hominem A: That it is not always fallacious, and that in some instances, questions of personal conduct, character, motives, etc., are legitimate and relevant to the issue, as when it directly involves hypocrisy, or actions contradicting the subject's words. Q: what does that last phrase mentioned above mean? A: It is an argumentum(a quarrel; altercation) ad hominem, refers to several types of arguments, not all are fallacious. Q: where does this phrase come from? A: The ancient Greek. Q: are there any philosophers who have written about this? A: Yes, Greeks. Aristotle, Sextus Empiricus, John Locke, Charles Leonard Hamblin, Douglas N. Walton. Q: who is the first mentioned person? A: He was a Greek philosopher and polymath during the Classical period in Ancient Greece. Q: has he written any book? A: He has written on subjects including physics, biology, zoology, metaphysics, logic, ethics, aesthetics, poetry, theatre, music, rhetoric, psychology, linguistics, economics, politics, and government. Q: what did he theorize about dreaming? A: He explained that dreams do not involve actually sensing a stimulus. In dreams, sensation is still involved, but in an altered manner.He also explains that when a person stares at a moving stimulus such as the waves in a body of water, and then look away, the next thing they look at appears to have a wavelike motion. Q: who is the second philosopher mentioned earlier? A: Sextus Empiricus was a Pyrrhonist philosopher and a physician mostly involved in ancient Greek and Roman Pyrrhonism.]

Question: do his teachings/work have any similarities with buddhism?

Good Rewrite: do sextus empiricus' teachings/work have any similarities with buddhism?

Bad Rewrite: there are any similarities between the philosophers mentioned above and buddhism.

Context: [Q: who was the french leader the diplomats were trying to meet with A: French foreign minister Talleyrand Q: what was this affair about? A: Confrontation between the United States and Republican France that led to the Quasi-War. Q: what was this confrontation about? A: To negotiate a solution to problems that were threatening to break out into war. Q: can you name any one who attended the previous meetings? A: Charles Cotesworth Pinckney Q: who was he? A: He was an early American statesman of South Carolina, Revolutionary War veteran, and delegate to the Constitutional Convention. Q: where was he born? A: Charleston, South Carolina]

Question: what was his views regarding slaves?

Good Rewrite: what was charles cotesworth pinckney's views regarding slaves?

Bad Rewrite: whatever were carlos castellanos' views regarding slaves?

<— Omit Three Demonstrations —>

### Annotated Sample

Context: [{{current\_context}}]

Question: {{current\_query}}

Good Rewrite:

### Prompt for Query Expansion (Answer)

#### **Instruction**

Given a question, please answer the question in a sentence. The answer should be as informative as possible.

#### **Demonstrations**

Question: and by whom was the game the last of us established?

Answer: Andy Gavin and Jason Rubin. Naughty Dog, LLC (formerly JAM Software, Inc.) is an American first-party video game developer based in Santa Monica, California. Founded by Andy Gavin and Jason Rubin in 1984 as an independent developer.

Question: is chelsea a club?

Answer: Yes, chelsea is an English professional football club.

Question: is call me by your name a movie?

Answer: Yes, based on a book of the same name. Call Me by Your Name is a 2017 coming-of-age romantic drama film directed by Luca Guadagnino. Its screenplay, by James Ivory, who also co-produced, is based on the 2007 novel of the same name by Andr Aciman.

Question: where was alan menken born?

Answer: Ian Irwin Menken was born on July 22, 1949, at French Hospital in Manhattan, to Judith and Norman Menken.

Question: where was ulysses s. grant from?

Answer: Hiram Ulysses Grant was born in Point Pleasant, Ohio, on April 27, 1822, to Jesse Root Grant, a tanner and merchant, and Hannah Simpson Grant.

#### **Annotated Sample**

Question: `{{reformulation_query}}`

Answer:

Table 13: The prompt for query expansion by directly answering the question.

### Prompt for Query Expansion (Keywords)

#### **Instruction**

Write a few keywords for the given query.

#### **Annotated Sample**

Query: `{{reformulation_query}}`

Keywords:

Table 14: The prompt for query expansion by giving keywords.

---

**Conversation:**

Q1: What was the Securities Act of 1933?

A1: The Securities Act of 1933 has two basic objectives: To require that investors receive financial and other significant information concerning securities being offered for public sale; and. To prohibit deceit, misrepresentations, and other fraud in the sale of securities.

Q2: What is exempt from it?

A2: However, there are exempt securities, under Section 4 of the Securities Act of 1933. These securities are financial instruments that carry government backing and typically have a government or tax-exempt status

Q3: Why was it needed?

A3: The act took power away from the states and put it into the hands of the federal government. The act also created a uniform set of rules to protect investors against fraud.

Q4: What was the reason for creating the 1934 act?

A4: The SEA of 1934 was enacted by Franklin D. Roosevelt's administration as a response to the widely held belief that irresponsible financial practices were one of the chief causes of the 1929 stock market crash.

Q5: What is the largest securities exchange in the world?

A5: The New York Stock Exchange founded on May 17, 1792, is the world's biggest stock exchange in trader value and has a capitalization of \$19.223 Trillion USD.

**Original Query:** How and when when was it created? (**rank: Not Found**)

**Human Rewrite:** How and when was the largest securities exchange in the world created? (**rank: Not Found**)

**ADACQR(Ours):** How and when was the Securities Act of 1933 created? The New York Stock Exchange founded on May 17, 1792? (**rank: 2**)

---

Table 15: Case study on QReCC (id: 8\_6) when using BM25. The underline part shows the decontextualized information in the reformulation query. We do not show the gold passage of this case because it is too long.

---

**Conversation:**

Q1: when did the first episode of sesame street air?

A1: 10 November 1969.

Q2: is it a series?

A2: Yes, an educational children's television series.

Q3: is the series still running?

A3: Yes, it is on its second run.

Q4: when was the pilot episode done?

A4: July 1969.

Q5: is it a puppet show?

A5: Yes.

Q6: name a few characters from the series?

A6: Human characters included Susan and muppet characters included Elmo.

Q7: can you name more muppet characters?

A7: Big Bird and Oscar the Grouch.

Q8: how do the latter look like?

A8: He has a green body with no visible nose.

Q9: does the muppet perform any oscar related play?

A9: UNANSWERABLE

Q10: who performed the aforementioned bird muppet?

A10: It was performed by Caroll Spinney till his retirement.

Q11: who is he by profession?

A11: He was an American puppeteer, cartoonist, author and speaker.

**Original Query:** did he do comics too? (**rank: Not Found**)

**ADACQR(Ours):** did Caroll Spinney do Caroll comics? (**rank: 1**)

**Gold Passage:** Caroll Spinney Comics and cartoons While in the Air Force, Spinney wrote and illustrated "Harvey", a comic strip about military life. He also animated a series of black-and-white cartoons called "Crazy Crayon".

---

Table 16: Successful case study on TopiOCQA (id: 16\_12) when using BM25. The underline part shows the decontextualized information in the reformulation query.

---

**Conversation:**

Q1: does callie baby die in season 7 episode 18?

A1: No.

Q2: who plays the character mentioned above?

A2: Sara Ramirez.

Q3: apart from acting, does she have a career in any other profession?

A3: She is a singer and songwriter.

Q4: name some of her songs ?

A4: Silent Night.

Q5: what is the significance of the above song?

A5: It is a popular Christmas carol.

Q6: who has written it?

A6: Joseph Mohr

Q7: the above mentioned episode is from which series?

A7: "Grey's Anatomy"

Q8: name some characters of it.

A8: Meredith Grey, Alex Karev, Miranda Bailey and Richard Webber

Q9: what is the real name of the third character mentioned in the above list?

A9: Chandra Wilson

Q10: which movie did she debut in?

A10: "Philadelphia"

**Original Query:** what was it about? (**rank: Not Found**)

**ADACQR:** what was the movie "Philadelphia" about? (**rank: Not Found**)

**AdaCQR + Expansion:** what was the movie "Philadelphia" about? **Philadelphia is a 1993 American drama film directed by Jonathan Demme and starring Tom Hanks and Denzel Washington. The movie tells the story of Andrew Beckett, a gay lawyer who is fired from his job because of his sexual orientation, and his subsequent fight for justice and equality in the legal system. Philadelphia, movie, Tom Hanks, Denzel Washington, AIDS, discrimination, lawsuit. (rank: 1)**

**Gold Passage:** Philadelphia (film) Introduction Philadelphia is a 1993 American legal drama film written by Ron Nyswaner, directed by Jonathan Demme and starring Tom Hanks and Denzel Washington. It was one of the first mainstream Hollywood films to acknowledge HIV/AIDS, homosexuality, and homophobia. For his role as Andrew Beckett, Hanks won the Academy Award for Best Actor at the 66th Academy Awards, while the song "Streets of Philadelphia" by Bruce Springsteen won the Academy Award for Best Original Song. Nyswaner was also nominated for the Academy Award for Best Original Screenplay, but lost to Jane Campion for "The Piano".

---

Table 17: Successful case study with query expansion on TopiOCQA (id: 55\_11) when using BM25. The **part** and the **part** represent the answers and keywords generated by LLM, respectively. These components furnish additional information that assists the retriever in enhancing its performance.

---

**Conversation:**

Q1: How can you tell if someone is suffering from depression?

A1: You may be depressed if, for more than two weeks, you've felt sad, down or miserable most of the time, or have lost interest or pleasure in usual activities, and have also experienced several of the signs and symptoms across at least three of the categories below.

Q2: What are common types?

A2: The four most common types of depression are major depression, persistent depressive disorder (formerly known as dysthymia), bipolar disorder, and seasonal affective disorder.

Q3: What causes it?

A3: Rather, there are many possible causes of depression, including faulty mood regulation by the brain, genetic vulnerability, stressful life events,

Q4: What is the role of brain chemicals?

A4: A chemical imbalance in the brain is said to occur when there's either too much or too little of certain chemicals, called neurotransmitters, in the brain. Neurotransmitters are natural chemicals that help facilitate communication between your nerve cells.

Q5: Does a lack of sunlight cause it?

A5: If you're not careful, a lack of sunlight can actually lead to a form of clinical depression.

**Original Query:** How can you treat SAD? (**rank: 2**)

**Human Rewrite:** How can you treat SAD? (**rank: 2**)

**ADACQR(Ours):** How can SAD be treated if a lack of sunlight in the brain can actually lead to depression? (**rank: Not Found**)

---

Table 18: Failure case on QReCC (id: 57\_6) when using BM25.