

IRLab@iKAT24: Learned Sparse Retrieval with Multi-aspect LLM Query Generation for Conversational Search

Simon Lupart*
University of Amsterdam
Amsterdam, Netherlands
s.c.lupart@uva.nl

Zahra Abbasiantaeb*
University of Amsterdam
Amsterdam, Netherlands
z.abbasiantaeb@uva.nl

Mohammad Aliannejadi
University of Amsterdam
Amsterdam, Netherlands
m.aliannejadi@uva.nl

Abstract

The Interactive Knowledge Assistant Track (iKAT) 2024 focuses on advancing conversational assistants, able to adapt their interaction and responses from personalized user knowledge. The track incorporates a Personal Textual Knowledge Base (PTKB) alongside Conversational AI tasks, such as passage ranking and response generation. Query Rewrite being an effective approach for resolving conversational context, we explore Large Language Models (LLMs), as query rewriters. Specifically, our submitted runs explore multi-aspect query generation using the MQ4CS framework, which we further enhance with Learned Sparse Retrieval via the SPLADE architecture, coupled with robust cross-encoder models. We also propose an alternative to the previous interleaving strategy, aggregating multiple aspects during the reranking phase. Our findings indicate that multi-aspect query generation is effective in enhancing performance when integrated with advanced retrieval and reranking models. Our results also lead the way for better personalization in Conversational Search, relying on LLMs to integrate personalization within query rewrite, and outperforming human rewrite performance.

Keywords

conversational search, personalization, retrieval augmented generation, query understanding, neural sparse retrieval

1 Introduction

With the rapid advancement of conversational AI systems, conversational search (CS) has seen an important development recently, along with a need for systems that can engage in meaningful, context-aware interactions. The TREC interactive Knowledge Assistant Track (iKAT) 2024 specifically addresses these needs for a second year now [3], by developing conversational assistants capable of discussing a wide range of topics with users. These assistants are not only designed to engage in dialogue but also to support

users in making informed decisions by providing contextually relevant, personalized responses and dynamically ranked information within each conversational exchange.

The TREC iKAT achieves this by leveraging two main sources of knowledge: a broad knowledge collection and a Personal Text Knowledge Base (PTKB), which contains user-specific information. This PTKB captures individual preferences, constraints and prior experience of the user, all in a textual format. TREC iKAT builds upon the TREC Conversational Assistance Track (CAST) [5, 6, 21], extending its core task to incorporate personalization. By integrating user-centric data, the track makes the retrieval and response generation processes significantly more challenging, as systems must adapt to each user's unique context and preferences. To address this complexity, TREC iKAT includes several subtasks: passage ranking, PTKB classification, and response generation. To help understand the behavior of the different models, the track also proposes different submission formats: (1) an automatic run, where all tasks are processed together by the system, (2) a manual run, which includes human-generated rewrites of each conversational turn to enhance context modeling for retrieval, and (3) a generation-only run, which concentrates on response generation using an already ranked list of passages provided by the organizers.

Recent studies in conversational search have demonstrated that Large Language Models (LLMs) are effective in tackling such tasks [1, 3, 18–20], disambiguating the user's latest utterance in the conversation. Building on this, our submitted runs explore advanced query rewriting techniques to resolve conversational ambiguities. Our work leverages the recent MQ4CS framework [1], which we further extend. Additionally, we aim to push the state of the art in conversational search by deploying enhanced retrieval and reranking models, strengthening the accuracy and relevance of system-generated responses across personalized conversations.

Through our runs, we show how the MQ4CS framework generalizes to the new iKAT 2024 benchmark. We also propose a variation of the original MQ4CS framework and show the additional improvement of combining it with state-of-the-art retrieval and reranking models.

2 Tasks

As last year, iKAT 2024 provides user utterance, conversation history, and Personal Text Knowledge Base (PTKB) as input for each conversational turn. The conversation history includes the user's previous user utterances and gold responses from the system. TREC iKAT 2024 benchmark includes 103 turns across 13 topics, each paired with a unique user persona (i.e., PTKB) of 16.8 descriptive

*Equal contribution.

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

Conference'17, July 2017, Washington, DC, USA

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM

<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

statements on average. Passage ranking is done on the ClueWeb-iKAT collection, with each conversational turn evaluated independently. The iKAT subtasks are as follows:

- (1) **Passage Ranking:** aims to retrieve and rank the relevant passages from the given collection in response to each user utterance. To understand the user query per each turn, the system must do reasoning over relevant statements from PTKB and conversation history.
- (2) **Response Generation:** The goal of this task is to provide a natural answer to the user utterance, following the flow of the conversation, and based on the retrieved passages. The responses must be grounded from the retrieved passages from the collection. Hence, the responses must be generated using at least one passage, referred to as “provenance”, from the source collection.
- (3) **PTKB Classification:** the PTKB consists of several statements. The goal of this task is to classify the statements within the PTKB as relevant/irrelevant for each conversational turn. The output of this task is a binary label for each statement from the PTKB.

Besides, the organizers proposed different types of runs: first the automatic runs (including the three subtasks all at once), then the manual runs (where user utterances are disambiguated from the conversation history and already include relevant PTKB statements), and finally the response generation-only subtask. We focus on the automatic runs, as the most challenging run, and also on the manual runs.

3 Methodology

In this section we explain the method used for the different tasks including (1) Passage ranking, (2) Response generation and (3) PTKB classification.

3.1 Passage Ranking

Our proposed models for passage ranking is based on the MQ4CS [1] framework. We further make some changes to the rank fusion step of MQ4CS framework in our submissions. We will briefly explain the MQ4CS framework and our modification to the rank fusion step in the following.

Multiple Aspect Query Generation. We strongly rely on LLMs to disambiguate the conversational context and persona of the user. In particular, we use the MQ4CS framework proposed by [1], generating multiple queries each covering different aspects of the information need. MQ4CS proposes breaking the user information need into multiple queries rather than forming a single query rewrite which includes the complete information need. This allows for better coverage of the collections, as a query expansion, but also enables a decomposition of the query into several sub-pieces of simpler information needs. Note that the queries would be generated for each user utterance by doing reasoning over both the context and personal information about the user provided in PTKB.

Rank Fusion. In MQ4CS, a retrieval and reranking pipeline is applied to each of the generated queries, producing several ranked lists that will be interleaved into a final ranking. As mentioned by the authors, this interleaving is sub-optimal. Hence, we propose to only use the multiple generated queries for retrieval and then re-rank all retrieved passages using a single query rewrite. To be more




Figure 1: MQ4CS-QR, as a variation of the original MQ4CS framework, that uses multiple query generation for Retrieval, but an independent single rewrite for Reranking.

specific, instead of interleaving the passages retrieved by multiple generated queries, we re-rank them using a single query rewrite as described in Figure 1. We follow the zero-shot query-rewriting baseline based on the GPT-4 model proposed by [1] to generate a single query rewrite. This would benefit from a large recall, as the different queries cover different aspects, and using these queries ensures having a diverse set of passages from different sources of information, but still have high precision, as relying on a single query rewrite to re-rank them.

Retrieval & Re-ranking. We extend the original work on MQ4CS that uses term-based retrieval with learned sparse retrieval. Our runs use SPLADE, which adapts the MLM pre-trained layer of BERT to predict term importance and token expansion within the WordPiece vocabulary [7–9, 15].

For reranking, we also propose to use more effective re-rankers compared to the one used in the original MQ4CS paper. Those re-rankers were trained with negatives from SPLADE, making them more effective in detecting noisy passages that could have been retrieved from our first-stage model. We also rely on the ensembling of rankers, using a min-max normalization on the produced ranked list.

3.2 Response Generation

For response generation, we rely on the simple yet effective Retrieval Augmented Generation (RAG) pipeline using LLMs on top of retrieved passages [16]. In this pipeline, the model is given the user utterance, context of the conversation, PTKB of the user, and the top passages retrieved from our passage ranking. Each submission follows this pipeline, with a different top ranking. We use a similar prompt to Abbasiantaeb et al., including all PTKB within the prompt. This approach was effective for the last year of iKAT.

3.3 PTKB classification

In both previous tasks (Ranking and Response Generation), all PTKB are included in the model’s input. We thus entirely rely on LLMs to filter and extract relevant information from the user in different contexts. However, to better understand and explicit the model reasoning, we also prompted the model to classify relevant statements from the PTKB list. In particular, we follow the zero-shot approach proposed by [2] for iKAT 2023, where an LLM is prompted to return the relevant PTKB given the user utterance and context of the conversation, with an adaptation for classification.

As an alternative approach, one could only use those filtered PTKB for the context modeling and response generation tasks,

Table 1: Experimental results on the official evaluation set. Bold and underlined indicate the best and second-best results across both Automatic and Manual runs.

Runs	LLM	Retrieval	Reranker	nDCG@5	nDCG	MRR	Recall@100	P@20	mAP
(1) MQ4CS-QR-deberta	GPT-4	SPLADE	DebertaV3	<u>0.5213</u>	0.6098	0.8479	0.4366	<u>0.5927</u>	0.3486
(2) MQ4CS-QR-ensemble	GPT-4	SPLADE	Ensemble	0.5155	0.6192	0.8644	<u>0.4657</u>	0.6184	0.3698
(3) GPT4QR-deberta	GPT-4	SPLADE	DebertaV3	0.5071	0.5420	0.8315	0.4135	0.5684	0.3166
(4) GPT4QR-bm25-QD1	GPT-4	BM25	MiniLM	0.4768	0.4373	0.8476	0.3421	0.5199	0.2315
(5) HumanQR-deberta	Human	SPLADE	DebertaV3	0.4767	0.5496	0.7808	0.4263	0.5383	0.3072
(6) HumanQR-ensemble	Human	SPLADE	Ensemble	0.5384	0.5819	<u>0.8483</u>	0.4777	0.5922	<u>0.3533</u>

breaking down the personalization task, but we believe that LLMs are effective enough to combine multiple tasks.

4 Experimental Setup

4.1 Submitted Runs

The final submitted runs can be decomposed into automatic and manual runs. We did not experiment further with the answer generation-only subtask.

Automatic runs:

- (1) *MQ4CS-QR-deberta*: This run uses the MQ4CS framework, generating multiple queries ($\phi=5$, GPT4) for our first-stage SPLADE retrieval. For reranking, we generate an additional single query rewrite from GPT-4 and rerank it with a DebertaV3 cross-encoder. Both PTKB classification and response generation models use GPT-4 as LLM. All prompts are zero-shot, following MQ4CS [1].
- (2) *MQ4CS-QR-ensemble*: Similar to the first run, this submission relies on MQ4CS. The entire pipeline is similar except for reranking which uses an ensembling of 5 cross-encoder models (DebertaV2, DebertaV3, Roberta, Alert, Electra). Response generation and PTKB classification use the same method as in the previous run, with GPT-4.
- (3) *GPT4QR-deberta*: For this run, we follow the common practice of using a single query rewrite. This rewrite is the same single rewrite used in both previous runs for reranking, generated from GPT-4 in zero-shot. Retrieval and reranking use SPLADE and DebertaV3. Response generation and PTKB classification use the same method as in the previous runs, with GPT-4.
- (4) *GPT4QR-bm25-QD1*: This run uses a single query, but with the prompting strategy of MQ4CS, with $\phi=1$, GPT-4. It further uses BM25 and a miniLM cross-encoder. Response generation and PTKB classification remain unchanged. The objective of this run was to compare it with the baseline runs from the organizers, to see the impact of the prompting strategies with a single query.

Manual Runs:

- (5) *HumanQR-deberta*: This manual run uses the single human rewrite for both retrieval and reranking. It relies on SPLADE and DebertaV3. Response generation and PTKB classification are added afterward from GPT-4.
- (6) *HumanQR-ensemble*: similar to the first manual run, this submission uses the manual human rewrite, but with an

ensemble of 5 cross-encoders for reranking. Response generation and PTKB classification are also added afterward from GPT-4.

Note that all our submissions (both automatic and manual) use the same PTKB classification results.

4.2 Hyperparameters

We provide in this section the details and sources of the models we used. Our retrieval mostly relies on SPLADE, using the CoCondenser SelfDistil SPLADE++ checkpoint from HuggingFace¹. We also rely on BM25 for one of the runs, with default hyperparameters. For re-ranking, the list of rerankers is **debertav3 large**, **debertav2 xlarge** [10, 11], **roberta large** [17], **albert v2-xxlarge** [13] and **electra large** [4]. All those models were fine-tuned by Lassance and Clinchant for re-ranking. The list of rerankers can be found on HuggingFace². We also rely on the MS MARCO MiniLM cross-encoder³. For ensembling, we follow the traditional setup with the ranx library. All reranking rerank the top 1000 of retrieved passages.

4.3 Metrics

We report traditional metrics in IR, that we received from TREC. This includes Mean Reciprocal Rank (MRR), Normalized Discounted Cumulative Gain (nDCG) [12] and Precision, as well as recall and mean Average Precision. Those metrics are the ones used last year for iKAT [3]. We however could not use a ≥ 2 threshold for MRR and recall, as we did not have access to the qrels yet (only access to our results). For similar reason, we could also not evaluate the response generation, but we will update after we obtain the results from the organizers.

5 Results

In the section below, we present our retrieval results and main observations from the submitted runs.

5.1 Retrieval Performance

Multiple Query Generation. Table 1 contains the main retrieval results of our six submitted runs, manual and automatic. To first compare the performance when using multiple queries to represent the conversation, we look at the submitted runs (1) and (3), both reranked with DebertaV3. Results with the multiple queries show an important improvement in recall-oriented metrics, in particular

¹naver/splade-cocondenser-selfdistil

²naver/trecdl22-crossencoder-{debertav2, debertav3, electra, albert, roberta}

³cross-encoder/ms-marco-MiniLM-L-6-v2

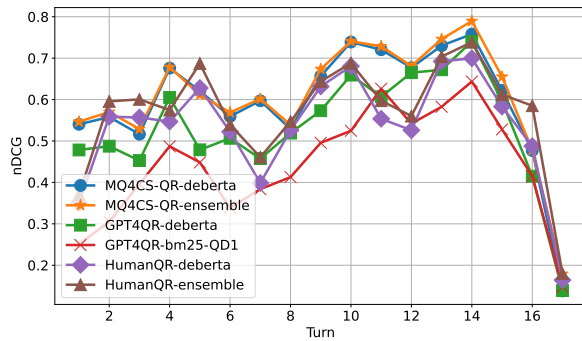


Figure 2: Performance per Depth.

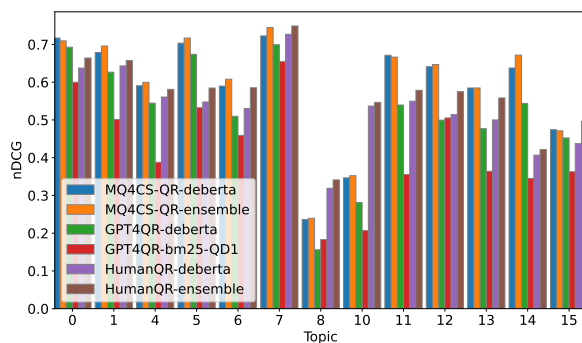


Figure 3: Performance per Topic.

a 2.3-point increase in Recall@100 and 3.2 points on mAP. This demonstrates the benefit of the MQ4CS framework combined with SPLADE. Also after the reranking, still comparing (1) and (3), we see that the improvement in recall for the first stage also leads to better precision after reranking. In particular, we see a 1.5-point gain on nDCG@5 and a 6.8-point increase on nDCG. This is possible because the set of candidate passages for re-ranking is of better quality.

Human vs. Automatic. Now looking at the manual runs, we see that the gap between manual and automatic runs is getting closer, with automatic runs outperforming the manual runs on nDCG, MRR, P@20, and mAP metrics. In particular when comparing (3) and (5) that both follow the same pipeline with a single rewrite either from Human or GPT4, we see how much GPT models now beat humans on the rewriting task. On the ensemble reranking approach ((2) and (6)), we also see that the automatic run outperforms on several metrics the manual one.

Ensembling of Rerankers. Ensembling uses several cross-encoders and averages their ranked list, to improve robustness and effectiveness. Looking at the Human rewrite, we see that ensembling leads to better performance compared to using only DebertaV3. We observe a similar trend on automatic runs combined with MQ4CS ((1) and (2)), with an average 2 points increase on all metrics, except for the nDCG@5 where the ensembling is 0.005 points below the single re-ranker.

5.2 Performance per Depths and Topics

Depth. Figure 2 presents the results of our set of models with respect to the depth of the turns of the conversation. While in previous years iKAT 2023 and generally in CAsT 2019-2022, longer conversations lead to harder context modeling, we see here that the nDCG increases as we go further within turns. An explanation for this can be the PTKB and the personalization task, as new information can be revealed by the user within the conversation, or relevant PTKB have been made explicit either by gold answers from previous turns or clarifying questions.

Topics. In Figure 3, we plot the average nDCG on the different topics from iKAT. We see from this figure that some conversations are more challenging than others, in particular topic 8 on mountain bikes. Furthermore, we can identify topics 11 to 14, on which our best automatic runs MQ4CS-QR Deberta and ensemble (blue and orange bars) outperform the two manual runs (purple, brown).

6 Conclusion

This report aims to describe the technical details of our iKAT 2024 submission. Through our submission, we demonstrate how the MQ4CS framework generalizes to iKAT 2024, and show the additional gains from combining it with Learned Sparse Retrieval and more effective cross-encoders models. We also propose a variation of the MQ4CS framework, that we will compare in future works to the original framework.

References

- [1] Zahra Abbasianteab, Simon Lupart, and Mohammad Aliannejadi. 2024. Generating Multi-Aspect Queries for Conversational Search. *arXiv preprint arXiv:2403.19302* (2024).
- [2] Zahra Abbasianteab, Chuan Meng, David Rau, Antonis Krasakis, Hossein A Rahmani, and Mohammad Aliannejadi. 2023. LLM-based Retrieval and Generation Pipelines for TREC Interactive Knowledge Assistance Track (iKAT) 2023. (2023).
- [3] Mohammad Aliannejadi, Zahra Abbasianteab, Shubham Chatterjee, Jeffrey Dalton, and Leif Azzopardi. 2024. TREC iKAT 2023: A Test Collection for Evaluating Conversational and Interactive Knowledge Assistants. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval* (Washington DC, USA) (SIGIR '24). Association for Computing Machinery, New York, NY, USA, 819–829. <https://doi.org/10.1145/3626772.3657860>
- [4] K Clark. 2020. Electra: Pre-training text encoders as discriminators rather than generators. *arXiv preprint arXiv:2003.10555* (2020).
- [5] Jeffrey Dalton, Chenyan Xiong, and Jamie Callan. 2020. CAsT 2020: The Conversational Assistance Track Overview. In *Text Retrieval Conference*. <https://api.semanticscholar.org/CorpusID:214735659>
- [6] Jeffrey Dalton, Chenyan Xiong, Vaibhav Kumar, and Jamie Callan. 2020. CAsT-19: A Dataset for Conversational Information Seeking. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval* (Virtual Event, China) (SIGIR '20). Association for Computing Machinery, New York, NY, USA, 1985–1988. <https://doi.org/10.1145/3397271.3401206>
- [7] Thibault Formal, Carlos Lassance, Benjamin Piwowarski, and Stéphane Clinchant. 2021. SPLADE v2: Sparse lexical and expansion model for information retrieval. *arXiv preprint arXiv:2109.10086* (2021).
- [8] Thibault Formal, Carlos Lassance, Benjamin Piwowarski, and Stéphane Clinchant. 2022. From Distillation to Hard Negative Sampling: Making Sparse Neural IR Models More Effective. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval* (Madrid, Spain) (SIGIR '22). Association for Computing Machinery, New York, NY, USA, 2353–2359. <https://doi.org/10.1145/3477495.3531857>
- [9] 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* (Virtual Event, Canada) (SIGIR '21). Association for Computing Machinery, New York, NY, USA, 2288–2292. <https://doi.org/10.1145/3404835.3463098>
- [10] Pengcheng He, Jianfeng Gao, and Weizhu Chen. 2021. Debertav3: Improving deberta using electra-style pre-training with gradient-disentangled embedding sharing. *arXiv preprint arXiv:2111.09543* (2021).

- [11] Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2020. Deberta: Decoding-enhanced bert with disentangled attention. *arXiv preprint arXiv:2006.03654* (2020).
- [12] Kalervo Järvelin and Jaana Kekäläinen. 2000. IR Evaluation Methods for Retrieving Highly Relevant Documents. In *Proceedings of the 23rd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval* (Athens, Greece) (SIGIR '00). ACM, New York, NY, USA, 41–48. <https://doi.org/10.1145/345508.345545>
- [13] Z Lan. 2019. Albert: A lite bert for self-supervised learning of language representations. *arXiv preprint arXiv:1909.11942* (2019).
- [14] Carlos Lassance and Stéphane Clinchant. 2023. Naver Labs Europe (SPLADE)@TREC Deep Learning 2022. *arXiv preprint arXiv:2302.12574* (2023).
- [15] Carlos Lassance, Hervé Déjean, Thibault Formal, and Stéphane Clinchant. 2024. SPLADE-v3: New baselines for SPLADE.
- [16] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. Retrieval-augmented generation for knowledge-intensive NLP tasks. In *Proceedings of the 34th International Conference on Neural Information Processing Systems* (Vancouver, BC, Canada) (NIPS '20). Curran Associates Inc., Red Hook, NY, USA, Article 793, 16 pages.
- [17] Yinhan Liu. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692* 364 (2019).
- [18] Simon Lupart, Mohammad Aliannejadi, and Evangelos Kanoulas. 2024. DiSCo Meets LLMs: A Unified Approach for Sparse Retrieval and Contextual Distillation in Conversational Search. *arXiv preprint arXiv:2410.14609* (2024).
- [19] Kelong Mao, Zhicheng Dou, Fengran Mo, Jiewen Hou, Haonan Chen, and Hongjin Qian. 2023. Large Language Models Know Your Contextual Search Intent: A Prompting Framework for Conversational Search. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, Houda Bouamor, Juan Pino, and Kalika Bali (Eds.). Association for Computational Linguistics, Singapore, 1211–1225. <https://doi.org/10.18653/v1/2023.findings-emnlp.86>
- [20] Fengran Mo, Abbas Ghaddar, Kelong Mao, Mehdi Rezagholizadeh, Boxing Chen, Qun Liu, and Jian-Yun Nie. 2024. CHIQ: Contextual History Enhancement for Improving Query Rewriting in Conversational Search. *arXiv preprint arXiv:2406.05013* (2024).
- [21] Paul Owoicho, Jeffrey Dalton, Mohammad Aliannejadi, Leif Azzopardi, Johanne R. Trippas, and Svitlana Vakulenko. 2022. TREC CAsT 2022: Going Beyond User Ask and System Retrieve with Initiative and Response Generation. In *Text Retrieval Conference*. <https://api.semanticscholar.org/CorpusID:261288646>

A Prompts

All our prompts re-use the one proposed within the MQ4CS [1] framework or previous work on iKAT 2023. We only changed the PTKB prompt, to adapt it to the classification task. Table 2 recall the multiple query generation prompt, Table 3 the RAG prompt for response generation, and Table 4 the prompt for PTKB classification.

Table 2: The prompt designed for multiple QD using GPT-4 as a zero-shot learner.

Multiple Query re-writing (QD).

Instruction: *I will give you a conversation between a user and a system. Imagine you want to find the answer to the last user question by searching on Google. You should generate the search queries that you need to search on Google. Please don't generate more than $\{\phi\}$ queries and write each query on one line.*

Background knowledge: $\{ptkb\}$

Context: $\{ctx\}$

User question: $\{user\ utterance\}$

Generated queries:

Table 3: The prompt designed for answer generation, using retrieved documents.

Retrieval Augmented Answer Generation (RAG).

Doc1: $\{doc_{(1)}\}$

Doc2: $\{doc_{(2)}\}$

Doc3: $\{doc_{(3)}\}$

Doc4: $\{doc_{(4)}\}$

Doc5: $\{doc_{(5)}\}$

I will give you a conversation between a user and a system. Also, I will give you some background information about the user. You should answer the last utterance of the user by providing a summary of the relevant parts of the given documents. Please remember that your answer shouldn't be more than 200 words.

Background information about the user: $\{ptkb\}$

Conversation: $\{ctx\}$

User query: $\{user\ utterance\}$

Table 4: The prompt designed for PTKB classification.

PTKB classification.

I will give you some background information about a user and a conversation between the user and a system. You should tell me which of the background information is relevant for answering the last question of the user.

Here is the background information about the user: $\{ptkb\}$

Please just copy the relevant background information to the last user utterance.
