

ConvMix: A Mixed-Criteria Data Augmentation Framework for Conversational Dense Retrieval

Fengran Mo^{1*}, Jinghan Zhang^{2*}, Yuchen Hui¹, Jia Ao Sun¹, Zhichao Xu³, Zhan Su^{1†}, Jian-Yun Nie^{1†}

¹University of Montreal, Canada

²Clemson University, USA

³University of Utah, USA

fengran.mo@umontreal.ca; jinghaz@clemson.edu; zhan.su@umontreal.ca; nie@iro.umontreal.ca

Abstract

Conversational search aims to satisfy users' complex information needs via multiple-turn interactions. The key challenge lies in revealing real users' search intent from the context-dependent queries. Previous studies achieve conversational search by fine-tuning a conversational dense retriever with relevance judgments between pairs of context-dependent queries and documents. However, this training paradigm encounters data scarcity issues. To this end, we propose **ConvMix**, a mixed-criteria framework to augment conversational dense retrieval, which covers more aspects than existing data augmentation frameworks. We design a two-sided relevance judgment augmentation schema in a scalable manner via the aid of large language models. Besides, we integrate the framework with quality control mechanisms to obtain semantically diverse samples and near-distribution supervisions to combine various annotated data. Experimental results on five widely used benchmarks show that the conversational dense retriever trained by our ConvMix framework outperforms previous baseline methods, which demonstrates our superior effectiveness.

Introduction

Conversational search aims to satisfy users' complex information needs via multiple-turn interaction in natural languages. The central hurdle lies in accurately understanding users' genuine search intent in the conversational session, given that their queries are context-dependent and prone to linguistic issues like omission, coreference, and ambiguity (Gao et al. 2022; Mo et al. 2025b).

When tackling conversational search, an intuitive approach is to use the conversational query rewriting (CQR) technique that breaks down the task by employing a query rewrite model to convert the query of the current turn into a de-contextualized one and then conducting an ad-hoc search retriever with the rewritten query (Elgohary, Peskov, and Boyd-Graber 2019). This approach allows for the use of existing retrievers for the search process. However, the CQR paradigm is challenging to directly optimize the rewriting towards search, due to the training signals do not directly impact the downstream retrieval tasks (Wu et al. 2022; Mo

et al. 2023a; Mao et al. 2023b; Ye et al. 2023; Mao et al. 2023a; Yoon et al. 2024). Another paradigm is to conduct conversational dense retrieval (CDR) that aims to train a conversational dense retriever to grasp the search intent by implicitly learning the latent representations of encoded queries and documents, which enable the model to naturally learn the relevant patterns between context-dependent queries and documents based on the relevance judgment signals (Yu et al. 2021; Kim and Kim 2022; Mao et al. 2023c; Jin et al. 2023; Mo et al. 2024c; Lupart, Aliannejadi, and Kanoulas 2025). Although these approaches achieve recognized improvements in conversational search, they encounter the data scarcity issue (Mao et al. 2022; Dai et al. 2022). This is because the conversational search engines are not widely deployed, and thus the available user query logs to generate relevance judgments for dense retriever training are not as abundant as the ad-hoc search scenarios (Bajaj et al. 2016).

Constructing conversational search datasets manually is a costly and difficult endeavor, so an intuitive alternative to enrich the training data is to automatically augment the existing data. Some prior studies (Mao et al. 2022; Dai et al. 2022) attempt to generate additional annotated data, but the generation process remains complex and relies on the assumption that a substantial amount of relevant in-domain data is available to build data augmentation frameworks, which is infeasible for scalable data augmentation. Recent studies (Huang et al. 2023; Mo et al. 2024d) aim to augment existing datasets with the aid of large language models (LLMs) for scalable generation, following the principle of reusing existing manually annotated data in conversations. However, these approaches simply instruct the LLMs for data generation from a single aspect, e.g., providing the few-shot examples, which cannot control and improve the diversification of the generated data. Chen et al. (2024) further develops a generalized data augmentation framework for conversational dense retrieval via LLM-cognition with a quality control mechanism to filter out those noisy and intent shift turns. However, it encounters the data distribution shift issue when combining multiple types of annotated data, i.e., the original human-annotated data and the LLM-generated ones (in different criteria), for model training. This should be alleviated by considering the near-distribution signals for direct utilized impact from the generated data samples (Peng et al. 2024; Dietz et al. 2025). Thus, an effective framework for augmenting

*Co-first authors.

†Corresponding authors.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

System	Multi-aspect Augmentation	LLM for Scalable	Quality Control	Near-Dist. Supervision
DialogInpaint	✓	✗	✗	✗
Converser	✗	✓	✗	✗
ConvTrans	✗	✗	✗	✗
ConvSDG	✓	✓	✗	✗
ConvAUG	✗	✓	✓	✗
ConvMix (Ours)	✓	✓	✓	✓

Table 1: The contained criteria comparison between ours and existing systems DialogInpaint (Dai et al. 2022), Converser (Huang et al. 2023), ConvTrans (Mao et al. 2022), ConvSDG (Mo et al. 2024d), ConvAUG (Chen et al. 2024), including whether support for various criteria.

conversational dense retrieval should mix multiple principles, including multi-aspect augmentation, scalable generation, quality control, and near-distribution supervision to apply to conversational context-dependent queries and their associated relevance judgments.

To address the data scarcity issue and the difficulty of data augmentation for conversational dense retrieval, we propose **ConvMix**, a mixed-criteria framework to augment conversational dense retrieval, whose supported criteria are overlooked in previous studies, as shown in Table 1. Since generating the search-oriented conversational session and its corresponding relevant documents simultaneously is infeasible for LLMs (Clarke and Dietz 2024; Balog, Metzler, and Qin 2025), we design a bidirectional relevance judgment augmentation from both query and document aspects: (i) generating augmented data from the existing manual relevance judgments by reformulating the corresponding conversational context-dependent queries; and (ii) generating pseudo generative relevance feedback (Mackie, Chatterjee, and Dalton 2023) by rewriting the relevant documents in terms of associated queries. This is achieved by deploying a generative LLM for scalable generation. Then, a quality control mechanism is designed to select qualified samples based on the aspects of semantic diversity. Besides, a near-distribution supervision mechanism is applied to select the model parameter-sensitive samples from the bidirectional augmented and the original data based on utilization estimation.

With the augmented data via scalable generation and quality selection from our ConvMix framework, we leverage these generated data with the original ones to train conversational dense retrievers and evaluate on five widely used conversational search benchmarks. The experimental results show that our ConvMix outperforms previous strong baselines with various techniques and demonstrate the superior effectiveness of our proposed framework.

Our contributions are summarized as follows:

(1) We propose a mixed-criteria data augmentation framework for conversational dense, which supports scalable multi-aspect augmentation by reusing existing conversational session data and relevance judgments from query and document sides based on LLMs.

(2) We design a quality control mechanism to select qualified samples from the aspects of increasing the diversification

of generated content and a near-distribution supervision mechanism to select the model parameter-sensitive sample based on utilization estimation.

(3) We conduct experiments on five widely used conversational search benchmarks with different settings and demonstrate the effectiveness of our approach by outperforming existing baseline methods.

Related Work

Conversational Search

Two main research lines in the literature are conducted for conversational search: (i) conversational query rewriting (CQR) and (ii) building a conversational dense retriever (CDR). The CQR methods aim to convert the context-dependent queries into stand-alone ones and reuse any existing ad-hoc search retrievers (Elgohary, Peskov, and Boyd-Graber 2019; Meng et al. 2025). Earlier CQR studies attempted to select useful tokens from the conversation context (Voskarides et al. 2020; Fang et al. 2022) or train a query rewriter with conversational sessions to mimic the human rewrites (Yu et al. 2020; Lin et al. 2020; Vakulenko et al. 2021). To mitigate the discrepancy between rewriter training and the downstream search task, some studies adopt reinforcement learning (Wu et al. 2022; Chen et al. 2022) or exploit the ranking signals with the rewriting model training (Mo et al. 2023a; Mao et al. 2023b; Mo et al. 2023b; Yoon et al. 2024). Recent LLM-based methods generate the query directly rewrites for retrieval (Ye et al. 2023; Mao et al. 2023a; Mo et al. 2024e, 2026) or to generate high-quality pseudo queries as supervision signals for rewrite model training (Jang et al. 2023; Mo et al. 2024a; Lai et al. 2024). The CDR methods (Yu et al. 2021; Lin, Yang, and Lin 2021; Mao, Dou, and Qian 2022; Jin et al. 2023; Cheng, Mao, and Dou 2024; Mao et al. 2024; Lupart, Aliannejadi, and Kanoulas 2025; Cheng et al. 2024; Mo et al. 2025a) directly encode the whole conversational search session into the model to perform end-to-end dense retrieval (Kim and Kim 2022; Mo et al. 2024c). This paradigm highly relies on the quantity and quality of the training samples with relevance judgment, and a series of studies aim to augment the training data from different aspects (Dai et al. 2022; Mao et al. 2022; Huang et al. 2023; Mo et al. 2024d; Chen et al. 2024).

Data Augmentation in Dense Retrieval

The quantity and quality of data hold significant value for dense retrieval. The development of LLMs provides the opportunity for scalable approaches to generate data for ad-hoc search models: to generate queries from a document (Bonifacio et al. 2022; Dai et al. 2023), to generate a document from a query (Gao et al. 2023; Mackie, Chatterjee, and Dalton 2023; Wang, Yang, and Wei 2023), etc. In the context of conversational search, the data generation framework should also consider the consistency of the conversational sessions (Joko et al. 2024; Li et al. 2025). In addition, a crucial point is to consider the data distribution between the original data sample and the augmented parts from different annotated criteria (Mo et al. 2024d), which requires avoiding generating

homogeneous data and obtaining near-distribution supervision (Peng et al. 2024). Besides, applying quality control mechanisms to select qualified samples is effective to reduce noise (Chen et al. 2024). Our proposed ConvMix is a mixed-criteria framework that addresses and covers all these aspects, which are overlooked in previous studies.

Motivation

Task Formulation

Conversational search systems aim to identify relevant documents d^+ from a large collection \mathcal{C} in response to the current query (n -th) q_n . This process is based on the context provided by the ongoing conversation session $\mathcal{X} = \{q_i\}_{i=1}^{n-1}$. Each query turn within a session depends on the preceding context, necessitating the conversational retriever to possess the capability to comprehend the user’s search intent. Thus, having access to high-quality and adequate conversational search session data is important for model training. The goal of augmenting conversational dense retrieval is to generate new session data $\mathcal{X}^{\text{aug}} = \{q'_i\}_{i=1}^n$ with corresponding relevant query-document pairs $\{(q'_i, d'_i)\}_{i=1}^n$. Then, the combined annotated data $\{\mathcal{X}\} \cup \{\mathcal{X}^{\text{aug}}\}$ from existing data $\{\mathcal{X}\}$ and the automatically augmented data $\{\mathcal{X}^{\text{aug}}\}$ could be used for fine-tuning conversational dense retrieval.

Diversity-Augmented Conversational Sessions

We conduct preliminary experiments to investigate whether augmenting the data from existing annotated datasets by increasing the diversity of conversational sessions can impact the dense retrieval performance. The principle is motivated by data augmentation in the field of computer vision, which involves rotating the original photo (Mikołajczyk and Grochowski 2018). In a similar pattern, we treat each conversational session as a photo, and achieving rotation for data augmentation by reordering the query turns in topic level. Concretely, given an existing conversational session $s = \{\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_t\}$ with t topics, where each topic $\mathcal{T}_i \in s$ would contain several query turns q_i , we apply random shuffle operation $\pi(\cdot)$ as a permutation function to the topic order in this conversational session to obtain a new one s' , i.e., $s = \pi(s')$. As a result, the new session $s' = \{\mathcal{T}_{\pi(1)}, \mathcal{T}_{\pi(2)}, \dots, \mathcal{T}_{\pi(t)}\}$, $\pi \neq \text{id}$, is a permutation of the original session s , i.e., they are content-equivalent but order-different. By applying this operation on all conversational sessions \mathcal{X} in the original dataset, we obtain double the data $\mathcal{X}_{\text{DA}}^{\text{aug}} = \{\mathcal{X}\} \cup \{\mathcal{X}'_{\text{DA}}\}$ by augmenting from the content diversification aspect. Then, we conduct conversational dense retriever fine-tuning with augmented data $\mathcal{X}_{\text{DA}}^{\text{aug}}$ and dense model ANCE (Xiong et al. 2020), which is evaluated on two datasets with available topic information, TopiOCQA (Adlakha et al. 2022) and QReCC (Anantha et al. 2021). The results are shown in Table 2, where we can observe that augmenting existing datasets by increasing the diversity of conversational sessions can improve the performance of the conversational dense retriever. With the validated conjecture, we continue to investigate the augmenting conversational dense retrieval by increasing content diversification and achieving quality control with mixed-criteria.

	TopiOCQA		QReCC	
	MRR	R@10	MRR	R@10
Original	22.9	43.0	46.3	70.4
+ Augmentation	29.8	49.2	49.1	72.5

Table 2: Preliminary experiments on diversity-oriented augmented conversational search data on two datasets.

Methodology

Overview

Augmenting conversational dense retrieval training data with multiple criteria is a non-trivial task. The augmented data generation should be aware of various aspects, including multi-aspect augmentation (Balog, Metzler, and Qin 2025) in a scalable way (Mo et al. 2024d), data quality control by filtering noisy samples (Chen et al. 2024), and obtain near-distribution supervision to fit different types of annotated data (Peng et al. 2024). By considering all these aspects, we propose **ConvMix**, a mixed-criteria framework to augment conversational dense retrieval data, whose overview is presented in Figure 1. Several components are integrated in our **ConvMix**, including (i) augmenting relevance judgment from two-sides by reformulating context-dependent query and rewriting its relevant documents, (ii) selecting semantic diversity samples by clustering to increase the diversification of generated data content, and (iii) selecting model parameter-sensitive sample with near-distribution supervision by utilization estimation.

Bidirectional Relevance Judgment Augmentation

One crucial criteria is to conduct data augmentation from multiple aspects to increase the diversification of the generated content in a scalable manner. However, even with the advanced LLMs, it is still infeasible to generate the search-oriented conversational session and its corresponding relevant documents simultaneously (Clarke and Dietz 2024; Sekulić, Alinanejadi, and Crestani 2024) and increasing the diversity of the generated results (Peng et al. 2024). To reduce the difficulty for LLMs and fully leverage the existing annotated datasets, we conduct data augmentation from two directions to separate the generation of conversational sessions and relevance judgment annotations, which can also generate data from multiple aspects to increase the diversity.

On one hand, we augment the data from the query side as **ConvMix-Q** variant, which aims to reuse the annotated relevance judgments in existing datasets by reformulating the context-dependent queries. We expect the reformulated query turns q_n^{ref} into another expression could keep the same search intent as the original one q_n , and then enhance the diversity of the generated new session after applying on every query turn. On the other hand, we conduct data augmentation from the existing relevant documents side as **ConvMix-D** variant, where we rewrite each relevant document d_n^+ for its associated query turn q_n and its conversational context $\{q_i\}_{i=1}^{n-1}$. The rewritten relevant documents d_n^{ref} serve as generative relevance feedback (Mackie, Chatterjee, and Dalton 2023),

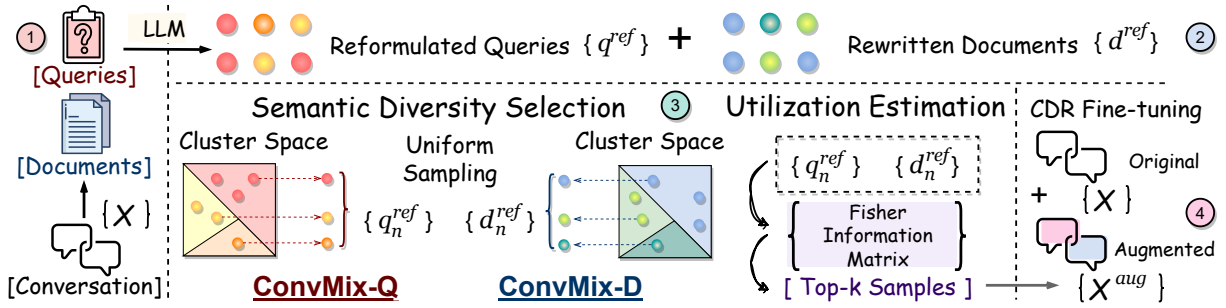


Figure 1: Illustration of our ConvMix framework, including bidirectional augmentation, two quality control mechanisms with semantic and utilization selection, and mixing original and augmented data for conversational dense retrieval fine-tuning.

which do not exist in the collection but share the same semantic information as the original documents. Although the rewritten relevant documents might contain wrong information or hallucination, it can still benefit to learn the semantic representation match for retriever model training (Bonifacio et al. 2022; Peng et al. 2024). Eventually, the relevance judgments are augmented from both sides based on the existing datasets, which can be formulated as

$$q_n^{\text{ref}} = \mathcal{LLM}(q_n | \{q_i\}_{i=1}^{n-1}), \quad d_n^{\text{ref}} = \mathcal{LLM}(d_n^+ | \{q_i\}_{i=1}^{n-1}, q_n)$$

where the $(\{q_i\}_{i=1}^{n-1} \circ q_n, d_n^+)$ is the available relevance judgment in the existing datasets.

To scale up the size of the augmented data, an intuitive approach is to produce multiple reformulated queries or rewritten documents, i.e., m variant of q_n^{ref} and d_n^{ref} would be generated as $\{q_n^{\text{ref}}\}_1^m$ and $\{d_n^{\text{ref}}\}_1^m$ ($m > 1$), in **ConvMix-Q** and **ConvMix-D**, respectively. The generated data from both sides corresponding to all turns consists of the augmented set \mathcal{X}^{aug} . Then, the follow-up question is how to control the quality and select the utilized ones of the generated data.

Clustering for Semantic Diversity Sample Selection

The pseudo data generated by LLM might be homogeneous and thus using them for model training would result in overfitting and compromise the model’s robustness (Peng et al. 2024). Therefore, increasing the diversification of the generated sample is crucial to mitigate this issue (Wang et al. 2025). However, the common practice to improve the diversity of the generated content is to adjust the temperature hyperparameter within the generative models, which does not directly impact the semantic diversity. To this end, we leverage clustering to select the generated data sample from the semantic diversity aspect. Specifically, with m generated content, e.g., reformulated queries $\{q_n^{\text{ref}}\}_1^m$ or rewritten relevant documents $\{d_n^{\text{ref}}\}_1^m$, for each (n -th) turn, a k -means clustering algorithm is used to partition them into k disjoint clusters as

$$\mathcal{C} = \{C_1, C_2, \dots, C_k\}, \quad \bigcup_{i=1}^k C_i = \{q_n^{\text{ref}}\}_1^m \text{ or } \{d_n^{\text{ref}}\}_1^m$$

where $C_i \cap C_j = \emptyset$ for $i \neq j$.

Then, from each cluster C_i , one item is randomly sampled as $x_{ni}^{\text{Div}} \sim \text{Uniform}(C_i)$. The final output set x_n^{Div}

$\subseteq \{q_n^{\text{ref}}\}_1^m / \{d_n^{\text{ref}}\}_1^m$ contains k augmented data samples and the whole procedure is formulated as

$$x_n^{\text{Div}} = \{x_{ni}^{\text{Div}} \sim \text{Uniform}(C_i) \mid C_i \in \text{Cluster}_k(\mathcal{X}^{\text{gen}})\}$$

where \mathcal{X}^{gen} could be either $\{q_n^{\text{ref}}\}_1^m$ or $\{d_n^{\text{ref}}\}_1^m$. Based on the hyperparameter k , we can obtain k -fold augmented data from both the reformulated query q_n^{ref} and the rewritten relevant document d_n^{ref} sides.

Utilization Estimation for Near-Distribution Supervision Selection via Fisher Information

Mixing data with various annotated types, i.e., the bidirectional augmented data and the original data, for model training might result in data distribution drift (Guilhaumon et al. 2024; Zhang et al. 2025). To mitigate this potential issue, the crucial solution is to obtain near-distribution supervision (Peng et al. 2024) by selecting those samples that are “most useful for model update” to ensure that the training with augmented data is both efficient and less noisy, and better impact on the utilized fitting of model parameters (Liu et al. 2020). Intuitively, a data sample that is highly sensitive to the model’s output is considered to carry significant information and thus such a sample is closer to the ideal data distribution and should thus be selected.

To achieve it, we leverage Fisher Information Matrix (FIM) (Fisher 1922) to measure the sensitivity of model parameter estimation in terms of a given augmented data sample from the previous selected set $\{\mathcal{X}^{\text{Div}}\}$. As shown in Algorithm 1, we compute the FIM score for each pseudo-labeled relevant query-document pair $(q_n^{\text{ref}}, d_n^{\text{ref}})$ by estimating the squared gradient norm of the prediction loss based on an ad-hoc search model used for conversational dense retrieval fine-tuning later. The top- k pairs with the highest scores are retained for downstream training, as they are assumed to provide the most informative updates to the model parameters. Eventually, the final mixed augmented datasets with near-distribution supervision are used for conversational dense retriever training as a combination **ConvMix-Combine**.

Conversational Dense Retrieval Fine-Tuning

With the augmented data $\{\mathcal{X}^{\text{aug}}\}$, we perform conversational dense retrieval fine-tuning together with the original data $\{\mathcal{X}\}$ in a semi-supervised learning paradigm. “Semi-supervised” means we combine the original data points with

Algorithm 1: Utilized Selection via Fisher Information

Require: Pseudo-labeled query-document pairs $\mathcal{X}_n^{\text{Div}} = \{(q_1^{\text{ref}}, d_1^{\text{ref}}), \dots, (q_m^{\text{ref}}, d_m^{\text{ref}})\}$ by data augmentation for a query turn q_n , retrieval model $f(\cdot, \cdot; \theta)$, top- k value k

Ensure: Selected final augmented set $\mathcal{X}_n^{\text{aug}}$

- 1: Initialize empty list $\mathcal{L} \leftarrow []$
- 2: **for** each $(q_i^{\text{ref}}, d_i^{\text{ref}}) \in \mathcal{X}_n^{\text{Div}}$ **do**
- 3: Compute similarity score $s_i = f(q_i^{\text{ref}}, d_i^{\text{ref}}; \theta)$
- 4: Set retrieval score $r_i = f(q_i, d_i^{\text{ref}}; \theta)$ or $f(q_i^{\text{ref}}, d_i^+; \theta)$
- 5: Compute loss $\ell_i = \text{Loss}(s_i, r_i)$
- 6: Compute gradient $\nabla_{\theta} \ell_i$ via backpropagation
- 7: Estimate Fisher information $F_i = \|\nabla_{\theta} \ell_i\|^2$
- 8: Append $(q_i^{\text{ref}}, d_i^{\text{ref}}, F_i)$ to \mathcal{L}
- 9: **end for**
- 10: Sort \mathcal{L} in descending order by F_i
- 11: Select top k pairs to form $\mathcal{X}_n^{\text{aug}}$
- 12: **return** $\mathcal{X}_n^{\text{aug}}$

manual labels and the generated data points without manual labels for model training, which enables the augmented parts with consistency regularization. We employ a widely used ad-hoc search retriever, ANCE (Xiong et al. 2020), which serves as both the query and document encoder, denoted as \mathcal{F}_Q and \mathcal{F}_D , respectively. In this process, we consider all preceding queries within the same session to concatenate with the current query turn as the session representation q_n^{concat} as

$$q_n^{\text{concat}} = q_1 \circ \dots \circ q_i \circ \dots \circ q_{n-1} \circ q_n, \quad i \in [1, n-1] \quad (1)$$

where \circ denotes concatenation. Then, a similarity function \mathcal{S} based on dot product is applied to score a candidate document d as Eq. 2. During the training phase, we update only the query encoder, while the document encoder is frozen. We follow the simplest fine-tuning principle with only in-batch negatives, which aims to directly reflect the impact of the augmented data rather than involving other advanced techniques in fine-tuning, e.g., hard negatives mining (Kim and Kim 2022) or injecting other signals (Mo et al. 2024b).

$$\mathcal{S}(q_n^{\text{concat}}, d) = \mathcal{F}_Q(q_n^{\text{concat}})^T \cdot \mathcal{F}_D(d) \quad (2)$$

The final training objective follows the widely used contrastive learning loss as

$$\mathcal{L} = \frac{e^{\mathcal{S}(q_n^{\text{concat}}, d^+)}}{e^{\mathcal{S}(q_n^{\text{concat}}, d^+)} + \sum_{d_n^- \in \mathcal{C}^-} e^{\mathcal{S}(q_n^{\text{concat}}, d^-)}} \quad (3)$$

where the d^+ and d^- denote the positive and negative samples for each query turn. During the inference, we perform the Approximate Nearest Neighbors search based on the dense index using Faiss (Johnson, Douze, and Jégou 2019).

Experiments

Experimental Setup

Datasets and Evaluation Metric. We use five conversational search datasets to evaluate our methods in different settings. All datasets contain relevance judgments for each original query, and the detailed statistical information is provided in

the Appendix. The TopiOCQA (Adlakha et al. 2022) and QReCC (Anantha et al. 2021) datasets are used to produce augmented data on top of their original relevance judgment annotations and conversational sessions. Then, the original data and the augmented data are combined to train conversational dense retrievers for the usual training-test evaluation. The three datasets (Dalton, Xiong, and Callan 2020, 2021, 2022) with only a few samples are used for out-of-domain evaluation (i.e., directly apply the models trained on TopiOCQA to the CAsT datasets). For a thorough comparison with existing systems, we use four standard metrics: MRR, NDCG@3, Recall@10, and Recall@100 to evaluate the retrieval effectiveness. Significance tests are conducted using paired t-tests at $p < 0.05$ level compared to the main competitors. We adopt the *pytrecc_eval* tool (Van Gysel and de Rijke 2018) for metric computation.

Baselines Methods. We compare our method with two categories of baselines: data augmentation methods for conversational dense retrieval (Augmented CDR), and conversational dense retrieval approaches with advanced fine-tuning techniques (Advanced CDR). The first category serves as main competitors, which includes the following systems: (1) **ConvTrans** (Mao et al. 2022): An approach to transform session search data into conversational search data for data augmentation. (2) **ConvSDG** (Mo et al. 2024d): A data augmentation approach to generate session-level conversational sessions with pseudo-relevant feedback as relevance judgments, and (3) **ConvAUG** (Chen et al. 2024): a framework to generate multi-level augmented conversations to capture the diverse nature of conversational contexts via LLM with human cognition mechanism. The second line of advanced CDR includes previous systems: (4) **InstructorR** (Jin et al. 2023): An unsupervised conversational dense retrieval method by instructing LLMs to produce LLM-based representation of the documents to serve as additional features. (5) **SDRConv** (Kim and Kim 2022): a fine-tuning method with different mined hard negatives. (6) **ConvDR** (Yu et al. 2021): A method to leverage both conversational search data and manually rewritten queries as supervision signals for conversational dense retrieval fine-tuning. (7) **HACovDR** (Mo et al. 2024c): An approach fine-tuned on context-denoising reformulated query and additional signals from historical turns based on the impact of retrieval effectiveness, and (8) **QRACDR** (Mo et al. 2024b): A query representation alignment conversational dense retriever to align the query representation with those of rewritten queries and relevant documents.

Implementation Details. We implement the conversational dense retriever based on ANCE (Xiong et al. 2020) using PyTorch and Huggingface libraries, ensuring all baseline systems are in a comparable setting. For data augmentation, each query turn is reformulated and relevant documents are rewritten on Mistral-7B-Instruct-v0.3 (Jiang et al. 2023), whose instruction templates are provided in the Appendix. The scalable relevance judgment augmentation is investigated from 1-fold to 10-fold. The top- k sample selection for semantic diversity clustering and the FIM score for near-distribution are both set to 3. For conversational dense retrieval fine-tuning, following previous studies (Yu et al. 2021; Mao et al. 2023c), the lengths of the query turn, session query, and document

Category	Method	TopiOCQA			QReCC		
		MRR	NDCG@3	Recall@10	MRR	NDCG@3	Recall@10
Advanced CDR	InstructorR	25.3	23.7	45.1	43.5	40.5	66.7
	SDRConv	26.1	25.4	44.4	47.3	43.6	69.8
	ConvDR	27.2	26.4	43.5	38.5	35.7	58.2
	HACConvDR	30.1	28.5	50.8	48.5	45.6	72.4
	QRACDR	31.8	30.6	50.0	47.1	44.5	71.4
Augmented CDR	ConvTrans	28.6	27.4	48.1	48.2	46.9	70.6
	ConvSDG	28.4	27.0	47.9	49.2	47.7	73.4
	ConvAUG	35.0	33.3	57.9	52.7	50.4	75.6
	ConvMix-Q	36.9	35.8	57.9	51.3	49.8	74.7
	ConvMix-D	34.4	33.3	55.2	50.8	49.2	74.0
	ConvMix-Combine	37.7	36.6	58.7	52.9	50.7	75.6

Table 3: Performance of different dense retrieval methods on two datasets. The number of training samples for our ConvMix-Q and ConvMix-D is double that of the original datasets, as the methods in Advanced CDR. The training sample for the other three main competitors, ConvTrans, ConvSDG, and ConvAUG, is the same as shown in Table 4. **Bold** indicates the best results.

Method	Training Sample	CASt-19			CASt-20			CASt-21		
		MRR	N@3	R@100	MRR	N@3	R@100	MRR	N@3	R@100
No Augmentation	45,450	66.2	40.1	29.2	42.5	26.5	31.6	46.7	29.1	37.4
ConvTrans	408,389	73.2	45.3	36.0	<u>45.9</u>	31.2	38.7	47.9	35.2	38.8
ConvAUG	136,858	-	-	-	45.0	30.7	38.4	54.8	36.8	45.9
ConvSDG	89,400	63.2	37.2	36.9	40.4	29.7	36.5	49.3	33.4	38.2
ConvMix-Q	90,900	64.5	37.8	36.6	45.4	<u>32.6</u>	39.1	<u>49.9</u>	34.6	42.5
ConvMix-D	90,900	67.4	41.4	36.3	44.7	31.5	38.3	<u>58.3</u>	<u>37.6</u>	<u>46.3</u>
ConvMix-Combine	136,350	<u>68.2</u>	<u>44.3</u>	<u>36.8</u>	46.1	33.0	<u>39.7</u>	60.3	38.4	47.2

Table 4: Out-of-Domain performance on three CASt datasets. **Bold** and underline indicate the best and the second-best results.

are truncated into 64, 512, and 384, respectively, to fit the majority of examples in the dataset. The batch size is set to 32 in accordance to our computational resources. We use Adam optimizer with 1e-5 learning rate and set the training epoch to 10. All experiments are conducted on four Nvidia V100 32G GPUs. More details are provided in our code repository: <https://github.com/fengranMark/ConvMix>.

Overall Performance

We first conduct the standard training-testing evaluation on TopiOCQA and QReCC, combining with the augmented data based on the original dataset.

The overall performance is shown in Table 3. We can observe that our ConvMix-Combine with data augmentation from combining reformulating query turns and rewriting relevant documents achieves the best performance on both datasets. Specifically, our ConvMix-Combine outperforms the main competitors ConvTrans, ConvSDG, and ConvAUG with 9.1%, 9.2%, and 3.3% absolute gains, respectively, on TopiOCQA in terms of MRR score, which demonstrates the effectiveness of augmenting the existing dataset from various aspects. The superior performance might be attributed to the quality control from the aspects of increasing semantic diversity and utilization estimation via clustering and FIM scores. The ConvMix-Q variant is consistently better than ConvMix-

D on two datasets, and they also outperform existing methods except for comparing to ConvAUG on QReCC, which indicates the importance of quality control mechanisms in ConvAUG and our ConvMix, while lacking in previous augmented CDR methods. Besides, we can find that ConvAUG and our ConvMix perform better than the previous advanced CDR methods, which emphasizes the potential of augmenting conversational dense retrieval via existing datasets. Further investigation could find a better way to connect augmenting CDR with advanced fine-tuned techniques.

Out-of-Domain Performance

We then conduct out-of-domain testing on three CASt datasets to evaluate the robustness of our ConvMix. The results are shown in Table 4, where the used training sample for each compared baseline is reported.

We can see that our ConvMix outperforms previous methods on two newer and more challenge datasets, CASt-20 and CASt-21, and is on par with those that have similar training samples on CASt-19. The better performance achieved by ConvMix with less training samples, which demonstrate its efficiency in data usage and the higher generated data in our mix-criteria framework. Besides, we observe similar results that better performance could be obtained by combining reformulating query turns and rewriting relevant documents,

Ablation	TopiOCQA		QReCC		CAsT-19		CAsT-20		CAsT-21	
	MRR	N@3	MRR	N@3	MRR	N@3	MRR	N@3	MRR	N@3
Full Model (ConvMix-Combine)	37.7	36.6	52.9	50.7	68.2	44.3	46.1	33.0	60.3	38.4
w/o Semantic Selection	36.2	35.1	50.6	48.9	66.4	40.2	45.2	31.9	57.8	37.0
w/o Utilization Selection	33.0	32.0	47.5	46.8	63.7	36.5	42.4	29.8	54.2	35.9
w/o Query-side Augmentation	34.4	33.3	50.8	49.2	67.4	41.4	44.7	31.5	58.3	37.6
w/o Doc-side Augmentation	36.9	35.8	51.3	49.8	64.5	37.8	45.4	32.6	49.9	34.6
w/o Original Data	32.2	32.9	46.2	45.3	60.1	35.1	36.7	25.1	48.0	31.1

Table 5: Ablation studies on five datasets within various components of our ConvMix.

i.e., ConvMix-Combine performs better than using either ConvMix-Q or ConvMix-D alone, which indicates our data augmentation framework can produce more diverse generated results from various aspects and effectively to mix them via near-distribution supervision.

Ablation Studies

In this section, we conduct ablation studies to investigate the effects of incorporating various components into our ConvMix framework, from the aspects of the quality control selection and the used annotated data types. The results are reported in Table 5.

In terms of the quality control mechanism, the semantic diversity selection and the utilization estimation selection are both effective across all datasets, indicating again the necessity of selecting high-quality samples in the data augmentation schema. The utilization estimation selection contributes more than the semantic diversity selection. The reason might be that the utilization estimation for near-distribution supervision based on the FIM information can more directly impact the model parameter updating and better fit the ideal data distribution than enhancing the semantic diversity.

In the context of the usage of various types of annotated data, we can observe that the original data is important to prevent data distribution drift, while using the augmented data alone can also achieve a competitive performance. Besides, the augmentation from the query and document sides are both effective. They can complement each other and further boost the performance with a combination. This phenomenon reveals that our ConvMix is robust to mix various annotated data and fully exploits their contribution for obtaining clean and near-distribution augmented data as supervision signals.

Ratio of Augmented Training Samples Impact

To study the impact of the amount of training data, we test the performance of two variants of our ConvMix with different augmented ratios. The results are shown in Figure 2.

We find that by increasing the augmentation scale of the generated training data, the retrieval performance shows an upward trend, but gradually becomes slower. The observation indicates that although increasing the amount of augmented training data can improve the effectiveness of the retrieval model, simply increasing it cannot achieve better results. This is because the noise and homogeneous information are also increasing among the generated data. Thus, it is necessary to improve the data quality for conversational dense retrieval

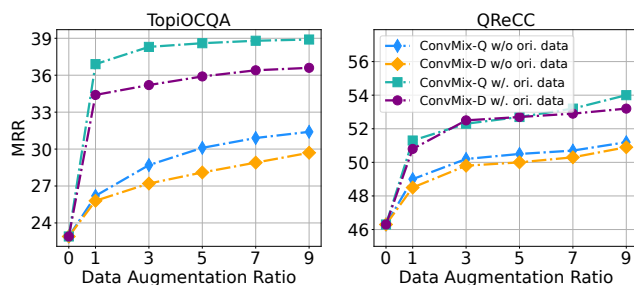


Figure 2: Model performance at the MRR score with various ratios of augmented training samples on two datasets.

by quality selection mechanisms, which demonstrate the importance of our designed approaches via increasing semantic diversity and estimating utilization scores for high-quality sample selection.

Besides, we can also observe that the effectiveness of data augmentation is more obvious on the TopiOCQA dataset than on QReCC. This might be related to the original construction of the dataset, where the TopiOCQA contains multiple topics within a conversational session than QReCC. Thus, the original TopiOCQA provides more abundant semantic information and makes it more possible to achieve effective data augmentation than QReCC. These phenomenon emphasizes the importance of content diversity in augmented data.

Conclusion

In this paper, we introduce **ConvMix**, a mixed-criteria framework to augment conversational dense retrieval with multi-aspect augmentation in a scalable manner. To ensure the highly-quality of the generated training sample, we improve the content diversification from semantic aspect. We further design a utilization estimation mechanism to obtain near-distribution supervision to combine various annotated data for effective conversational dense retrieval fine-tuning. The experimental results on five widely used conversational search benchmarks demonstrate the effectiveness of our framework by outperforming different existing baseline methods. Our study shows the potential of using an LLM to generate additional training data with multiple criteria for fine-tuning a dense retriever, which enriches the already extensive body of studies trying to exploit LLMs for search. More research is still required to find the best way to leverage LLMs for enhancing conversational search.

References

- Adlakha, V.; Dhuliawala, S.; Suleman, K.; de Vries, H.; and Reddy, S. 2022. TopiOCQA: Open-domain Conversational Question Answering with Topic Switching. *TACL*.
- Anantha, R.; Vakulenko, S.; Tu, Z.; Longpre, S.; Pulman, S.; and Chappidi, S. 2021. Open-Domain Question Answering Goes Conversational via Question Rewriting. In *NAACL*.
- Bajaj, P.; Campos, D.; Craswell, N.; Deng, L.; Gao, J.; Liu, X.; Majumder, R.; McNamara, A.; Mitra, B.; Nguyen, T.; et al. 2016. Ms marco: A human generated machine reading comprehension dataset. *arXiv preprint arXiv:1611.09268*.
- Balog, K.; Metzler, D.; and Qin, Z. 2025. Rankers, judges, and assistants: Towards understanding the interplay of llms in information retrieval evaluation. In *SIGIR*, 3865–3875.
- Bonifacio, L.; Abonizio, H.; Fadaee, M.; and Nogueira, R. 2022. Inpars: Data augmentation for information retrieval using large language models. *arXiv preprint arXiv:2202.05144*.
- Chen, H.; Dou, Z.; Mao, K.; Liu, J.; and Zhao, Z. 2024. Generalizing Conversational Dense Retrieval via LLM-Cognition Data Augmentation. *arXiv preprint arXiv:2402.07092*.
- Chen, Z.; Zhao, J.; Fang, A.; Fetahu, B.; Oleg, R.; and Malmasi, S. 2022. Reinforced Question Rewriting for Conversational Question Answering. In *EMNLP*.
- Cheng, Y.; Mao, K.; and Dou, Z. 2024. Interpreting Conversational Dense Retrieval by Rewriting-Enhanced Inversion of Session Embedding. In *ACL*.
- Cheng, Y.; Mao, K.; Zhao, Z.; Dong, G.; Qian, H.; Wu, Y.; Sakai, T.; Wen, J.-R.; and Dou, Z. 2024. CORAL: Benchmarking Multi-turn Conversational Retrieval-Augmentation Generation. In *NAACL*.
- Clarke, C. L.; and Dietz, L. 2024. LLM-based relevance assessment still can't replace human relevance assessment. *arXiv preprint arXiv:2412.17156*.
- Dai, Z.; Chaganty, A. T.; Zhao, V. Y.; Amini, A.; Rashid, Q. M.; Green, M.; and Guu, K. 2022. Dialog Inpainting: Turning Documents into Dialogs. In *PMLR*.
- Dai, Z.; Zhao, V. Y.; Ma, J.; Luan, Y.; Ni, J.; Lu, J.; Bakalov, A.; Guu, K.; Hall, K. B.; and Chang, M. 2023. Promptagator: Few-shot Dense Retrieval From 8 Examples. In *ICLR*.
- Dalton, J.; Xiong, C.; and Callan, J. 2020. TREC CAsT 2019: The conversational assistance track overview. In *TREC*.
- Dalton, J.; Xiong, C.; and Callan, J. 2021. CAsT 2020: The Conversational Assistance Track Overview. In *TREC*.
- Dalton, J.; Xiong, C.; and Callan, J. 2022. TREC CAsT 2021: The conversational assistance track overview. In *TREC*.
- Dietz, L.; Zendel, O.; Bailey, P.; Clarke, C. L.; Cotterill, E.; Dalton, J.; Hasibi, F.; Sanderson, M.; and Craswell, N. 2025. Principles and Guidelines for the Use of LLM Judges. In *ICTIR*, 218–229.
- Elgohary, A.; Peskov, D.; and Boyd-Graber, J. 2019. Can You Unpack That? Learning to Rewrite Questions-in-Context. In *EMNLP-IJCNLP*, 5918–5924.
- Fang, H.-C.; Hung, K.-H.; Huang, C.-W.; and Chen, Y.-N. 2022. Open-Domain Conversational Question Answering with Historical Answers. In *Findings of AACL-IJCNLP*.
- Fisher, R. A. 1922. On the mathematical foundations of theoretical statistics. *Philosophical transactions of the Royal Society of London. Series A, containing papers of a mathematical or physical character*, 222(594-604): 309–368.
- Gao, J.; Xiong, C.; Bennett, P.; and Craswell, N. 2022. Neural approaches to conversational information retrieval. *arXiv preprint arXiv:2201.05176*.
- Gao, L.; Ma, X.; Lin, J.; and Callan, J. 2023. Precise Zero-Shot Dense Retrieval without Relevance Labels. In *ACL*.
- Guilhaumon, C.; Hascoët, N.; Chinesta, F.; Lavarde, M.; and Daim, F. 2024. Data augmentation for regression machine learning problems in high dimensions. *Computation*.
- Huang, C.-W.; Hsu, C.-Y.; Hsu, T.-Y.; Li, C.-A.; and Chen, Y.-N. 2023. CONVERSER: Few-shot Conversational Dense Retrieval with Synthetic Data Generation. In *SIGDial*.
- Jang, Y.; Lee, K.-i.; Bae, H.; Won, S.; Lee, H.; and Jung, K. 2023. IterCQR: Iterative Conversational Query Reformulation without Human Supervision. *arXiv preprint arXiv:2311.09820*.
- Jiang, A. Q.; Sablayrolles, A.; Mensch, A.; Bamford, C.; Chaplot, D. S.; de las Casas, D.; Bressand, F.; Lengyel, G.; Lample, G.; Saulnier, L.; Lavaud, L. R.; Lachaux, M.-A.; Stock, P.; Scao, T. L.; Lavril, T.; Wang, T.; Lacroix, T.; and Sayed, W. E. 2023. Mistral 7B. *arXiv preprint arXiv:2310.06825*.
- Jin, Z.; Cao, P.; Chen, Y.; Liu, K.; and Zhao, J. 2023. InstructoR: Instructing Unsupervised Conversational Dense Retrieval with Large Language Models. In *Findings of EMNLP*.
- Johnson, J.; Douze, M.; and Jégou, H. 2019. Billion-scale similarity search with gpus. *IEEE Transactions on Big Data*.
- Joko, H.; Chatterjee, S.; Ramsay, A.; De Vries, A. P.; Dalton, J.; and Hasibi, F. 2024. Doing personal laps: Llm-augmented dialogue construction for personalized multi-session conversational search. In *SIGIR*, 796–806.
- Kim, S.; and Kim, G. 2022. Saving Dense Retriever from Shortcut Dependency in Conversational Search. In *EMNLP*.
- Lai, Y.; Wu, J.; Zhang, C.; Sun, H.; and Zhou, D. 2024. AdaCQR: Enhancing Query Reformulation for Conversational Search via Sparse and Dense Retrieval Alignment. In *COLING*.
- Li, L.; Cai, P.; Rossi, R. A.; Démoncourt, F.; Kveton, B.; Wu, J.; Yu, T.; Song, L.; Yang, T.; Qin, Y.; et al. 2025. A Personalized Conversational Benchmark: Towards Simulating Personalized Conversations. *arXiv preprint arXiv:2505.14106*.
- Lin, S.-C.; Yang, J.-H.; and Lin, J. 2021. Contextualized Query Embeddings for Conversational Search. In *EMNLP*.
- Lin, S.-C.; Yang, J.-H.; Nogueira, R.; Tsai, M.-F.; Wang, C.-J.; and Lin, J. 2020. Conversational question reformulation via sequence-to-sequence architectures and pretrained language models. *arXiv preprint arXiv:2004.01909*.
- Liu, J.; Yuan, H.; Lu, X.-M.; and Wang, X. 2020. Quantum Fisher information matrix and multiparameter estimation. *Journal of Physics A: Mathematical and Theoretical*, 53(2).
- Lupart, S.; Aliannejadi, M.; and Kanoulas, E. 2025. DiSCo Meets LLMs: A Unified Approach for Sparse Retrieval and Contextual Distillation in Conversational Search. In *SIGIR*.

- Mackie, I.; Chatterjee, S.; and Dalton, J. 2023. Generative relevance feedback with large language models. In *SIGIR*.
- Mao, K.; Deng, C.; Chen, H.; Mo, F.; Liu, Z.; Sakai, T.; and Dou, Z. 2024. ChatRetriever: Adapting Large Language Models for Generalized and Robust Conversational Dense Retrieval. In *EMNLP*.
- Mao, K.; Dou, Z.; Chen, H.; Mo, F.; and Qian, H. 2023a. Large Language Models Know Your Contextual Search Intent: A Prompting Framework for Conversational Search. In *Findings of EMNLP*.
- Mao, K.; Dou, Z.; Liu, B.; Qian, H.; Mo, F.; Wu, X.; Cheng, X.; and Cao, Z. 2023b. Search-Oriented Conversational Query Editing. In *Findings of ACL*, 4160–4172.
- Mao, K.; Dou, Z.; and Qian, H. 2022. Curriculum Contrastive Context Denoising for Few-shot Conversational Dense Retrieval. In *SIGIR*, 176–186.
- Mao, K.; Dou, Z.; Qian, H.; Mo, F.; Cheng, X.; and Cao, Z. 2022. ConvTrans: Transforming Web Search Sessions for Conversational Dense Retrieval. In *EMNLP*, 2935–2946.
- Mao, K.; Qian, H.; Mo, F.; Dou, Z.; Liu, B.; Cheng, X.; and Cao, Z. 2023c. Learning Denoised and Interpretable Session Representation for Conversational Search. In *Proceedings of the ACM Web Conference 2023*, 3193–3202.
- Meng, C.; Tonolini, F.; Mo, F.; Aletras, N.; Yilmaz, E.; and Kazai, G. 2025. Bridging the gap: From ad-hoc to proactive search in conversations. In *SIGIR*, 64–74.
- Mikołajczyk, A.; and Grochowski, M. 2018. Data augmentation for improving deep learning in image classification problem. In *2018 international interdisciplinary PhD workshop (IIPhDW)*, 117–122. IEEE.
- Mo, F.; Gao, Y.; Meng, C.; Liu, X.; Wu, Z.; Mao, K.; Wang, Z.; Chen, P.; Li, Z.; Li, X.; et al. 2025a. UniConv: Unifying Retrieval and Response Generation for Large Language Models in Conversations. In *ACL*, 6936–6949.
- Mo, F.; Gao, Y.; Wu, Z.; Liu, X.; Chen, P.; Li, Z.; Wang, Z.; Li, X.; Jiang, M.; and Nie, J.-Y. 2026. Leveraging historical information to boost retrieval-augmented generation in conversations. *Information Processing & Management*.
- Mo, F.; Ghaddar, A.; Mao, K.; Rezagholizadeh, M.; Chen, B.; Liu, Q.; and Nie, J.-Y. 2024a. CHIQ: Contextual History Enhancement for Improving Query Rewriting in Conversational Search. In *EMNLP*, 2253–2268.
- Mo, F.; Mao, K.; Zhao, Z.; Qian, H.; Chen, H.; Cheng, Y.; Li, X.; Zhu, Y.; Dou, Z.; and Nie, J.-Y. 2025b. A survey of conversational search. *ACM TOIS*, 43(6): 1–50.
- Mo, F.; Mao, K.; Zhu, Y.; Wu, Y.; Huang, K.; and Nie, J.-Y. 2023a. ConvGQR: Generative Query Reformulation for Conversational Search. In *ACL*, 4998–5012.
- Mo, F.; Nie, J.-Y.; Huang, K.; Mao, K.; Zhu, Y.; Li, P.; and Liu, Y. 2023b. Learning to Relate to Previous Turns in Conversational Search. In *ACM SIGKDD*.
- Mo, F.; Qu, C.; Mao, K.; Wu, Y.; Su, Z.; Huang, K.; and Nie, J.-Y. 2024b. Aligning query representation with rewritten query and relevance judgments in conversational search. In *CIKM*, 1700–1710.
- Mo, F.; Qu, C.; Mao, K.; Zhu, T.; Su, Z.; Huang, K.; and Nie, J.-Y. 2024c. History-Aware Conversational Dense Retrieval. In *Findings of ACL*, 13366–13378.
- Mo, F.; Yi, B.; Mao, K.; Qu, C.; Huang, K.; and Nie, J.-Y. 2024d. ConvSDG: Session Data Generation for Conversational Search. In *Companion Proceedings of the ACM on Web Conference 2024*, 1634–1642.
- Mo, F.; Zhao, L.; Huang, K.; Dong, Y.; Huang, D.; and Nie, J.-Y. 2024e. How to Leverage Personal Textual Knowledge for Personalized Conversational Information Retrieval. In *CIKM*, 3954–3958.
- Peng, L.; Gu, Y.; Dong, C.; Wang, Z.; and Shang, J. 2024. Text Grafting: Near-Distribution Weak Supervision for Minority Classes in Text Classification. In *EMNLP*, 3741–3752.
- Sekulić, I.; Alinanejadi, M.; and Crestani, F. 2024. Analysing utterances in llm-based user simulation for conversational search. *ACM TOIS*, 15(3): 1–22.
- Vakulenko, S.; Longpre, S.; Tu, Z.; and Anantha, R. 2021. Question rewriting for conversational question answering. In *WSDM*, 355–363.
- Van Gysel, C.; and de Rijke, M. 2018. Pytrec_eval: An Extremely Fast Python Interface to trec_eval. In *SIGIR*. ACM.
- Voskarides, N.; Li, D.; Ren, P.; Kanoulas, E.; and de Rijke, M. 2020. Query resolution for conversational search with limited supervision. In *SIGIR*, 921–930.
- Wang, L.; Yang, N.; and Wei, F. 2023. Query2doc: Query expansion with large language models. *arXiv preprint arXiv:2303.07678*.
- Wang, Z.; Zhang, J.; Zhang, X.; Liu, K.; Wang, P.; and Zhou, Y. 2025. Diversity-oriented data augmentation with large language models. *arXiv preprint arXiv:2502.11671*.
- Wu, Z.; Luan, Y.; Rashkin, H.; Reitter, D.; and Tomar, G. S. 2022. CONQRR: Conversational Query Rewriting for Retrieval with Reinforcement Learning. In *EMNLP*.
- Xiong, L.; Xiong, C.; Li, Y.; Tang, K.-F.; Liu, J.; Bennett, P. N.; Ahmed, J.; and Overwijk, A. 2020. Approximate Nearest Neighbor Negative Contrastive Learning for Dense Text Retrieval. In *ICLR*.
- Ye, F.; Fang, M.; Li, S.; and Yilmaz, E. 2023. Enhancing Conversational Search: Large Language Model-Aided Informative Query Rewriting. In *Findings of EMNLP*, 5985–6006.
- Yoon, C.; Kim, G.; Jeon, B.; Kim, S.; Jo, Y.; and Kang, J. 2024. Ask Optimal Questions: Aligning Large Language Models with Retriever’s Preference in Conversational Search. In *Findings of NAACL*.
- Yu, S.; Liu, J.; Yang, J.; Xiong, C.; Bennett, P.; Gao, J.; and Liu, Z. 2020. Few-shot generative conversational query rewriting. In *SIGIR*, 1933–1936.
- Yu, S.; Liu, Z.; Xiong, C.; Feng, T.; and Liu, Z. 2021. Few-shot conversational dense retrieval. In *SIGIR*, 829–838.
- Zhang, W.; Chen, X.; Sun, Y.; Guan, W.; and Nie, L. 2025. Batch Training for Streaming Time Series: A Transferable Augmentation Framework to Combat Distribution Shifts. *Transactions on Machine Learning Research*.