

ComRAG: Retrieval-Augmented Generation with Dynamic Vector Stores for Real-time Community Question Answering in Industry

Qinwen Chen^{†*}, Wenbiao Tao^{†*}, Zhiwei Zhu[†], Mingfan Xi[‡], Liangzhong Guo[‡]

Yuan Wang[‡], Wei Wang[†], Yunshi Lan[†](✉)

[†]School of Data Science and Engineering, East China Normal University

[‡]Alibaba Group

{qwchen, wbtao, 51255903077}@stu.ecnu.edu.cn

{mingfan.xmf, liangzhong.glz, jingxuan.wy}@alibaba-inc.com

{wwang, yslan}@dase.ecnu.edu.cn

Abstract

Community Question Answering (CQA) platforms can be deemed as important knowledge bases in community, but effectively leveraging historical interactions and domain knowledge in real-time remains a challenge. Existing methods often underutilize external knowledge, fail to incorporate dynamic historical QA context, or lack memory mechanisms suited for industrial deployment. We propose **ComRAG**, a retrieval-augmented generation framework for real-time industrial CQA that integrates static knowledge with dynamic historical QA pairs via a centroid-based memory mechanism designed for retrieval, generation, and efficient storage. Evaluated on three industrial CQA datasets, ComRAG consistently outperforms all baselines—achieving up to **25.9%** improvement in vector similarity, reducing latency by **8.7%–23.3%**, and lowering chunk growth from **20.23%** to **2.06%** over iterations.

1 Introduction

Community Question Answering (CQA) is a collaborative question-and-answer paradigm where users post questions on online platforms (e.g., Stack Overflow¹ and AskUbuntu²) and community members contribute answers. This paradigm leverages collective intelligence, allowing users to refine answers through voting, commenting, and editing, ultimately enhancing the quality of shared knowledge (Roy et al., 2023). With the rise of ChatGPT (Achiam et al., 2023), DeepSeek (DeepSeek-AI et al., 2025), and other foundation models, Large Language Models (LLMs) have become powerful tools for CQA. However, existing CQA methods focus on static community knowledge, limiting their applicability to real-world scenarios.

We categorize existing CQA methods as follows:

(1) Retrieval-based methods: Retrievers or rankers

identify the most relevant answers from the community. Question-answer cross-attention networks with knowledge augmentation are utilized for answer selection (Hu, 2023), while structured information is leveraged to enhance ranking (Askari et al., 2024; Ghasemi and Shakery, 2024). (2) Generation-based methods: LLMs serve as community experts to answer professional questions. Techniques such as instruction tuning (Yang et al., 2023), reinforcement learning (Gorbatovski and Kovalchuk, 2024) and contrastive learning (Yang et al., 2025) equip LLMs with domain and community knowledge.

However, these methods have the following weaknesses: 1) They often overlook external domain knowledge, limiting their applicability for domain-specific industrial applications. 2) Real-time CQA presents a continuous stream of questions rather than a static pool, requiring systems to reflect historical interactions. 3) A suitable memory mechanism is needed for real-time industrial CQA, but existing methods overlook this issue.

Domain knowledge and community interaction history play key roles in industrial CQA, shaping the professionalism and relevance of responses, respectively. Thus, the following two key questions require our attention. **(Q1) How can we build a CQA system that combines static knowledge with dynamic reflection on the disparate quality of historical answers?** Domain knowledge serves as an authoritative benchmark, while community QA history links user queries to relevant insights. Combining both enhances LLMs’ ability to generate reliable CQA responses. **(Q2) How can a real-time CQA system manage both the rapidly growing volume of historical QA data and the wide variance in response quality?** The evolution of the community leads to varying quality in historical responses due to the open and collaborative nature of CQA. Efficiently identifying, organizing, and leveraging high- and low-quality QA records

*Equal Contribution

¹<https://stackoverflow.com/>

²<https://askubuntu.com/>

becomes crucial for maintaining reliable generation.

To tackle these challenges, we propose ComRAG, a retrieval-augmented generation framework that integrates static domain knowledge with dynamic historical QA interactions to enhance real-time CQA in industrial settings. In the query phase, ComRAG supports three strategies based on query characteristics: directly reusing answers from high-quality QA pairs, generating responses with reference to high-quality content, and generating responses while explicitly avoiding low-quality ones. During generation, an adaptive temperature tuning mechanism ensures more confident responses. In the update phase, the system dynamically manages high- and low-quality CQA vector stores using a centroid-based memory mechanism, optimizing retrieval efficiency for continuous question streams.

In summary, the contributions are as follows.

- We propose ComRAG, a novel retrieval-augmented generation framework that jointly integrates static domain knowledge and dynamic community history to address real-time industrial CQA.
- We develop a centroid-based memory mechanism for efficient retrieval and an adaptive temperature tuning mechanism for confident generation.
- We extensively evaluate our framework on MSQA, ProCQA and PolarDBQA, demonstrating its effectiveness and efficiency for real-time industrial CQA.

2 Related Work

2.1 Community Question Answering

The CQA task centers on improving the relevance and quality of answers. We categorize existing approaches to CQA into two main paradigms: retrieval-based and generation-based.

Retrieval-based methods aim to identify relevant community answers using retrievers or rankers. Some enhance answer selection by integrating cross-attention networks with LLM-augmented knowledge (Hu, 2023) or incorporate structured metadata into cross-encoder re-ranking (Askari et al., 2024). Expert finding in CQA is supported by modeling user interactions via topic-based multi-layer graphs (Amen-dola et al., 2024) while modality-agnostic con-

trastive pretraining is proposed for aligning code-question pairs (Li et al., 2024). Expanding queries and computing translation-based similarity using category-specific dictionaries improve question retrieval (Ghasemi and Shakery, 2024).

Generation-based methods rely on LLMs acting as community experts to generate answers. Prior work explores strategies such as pretraining a small expert model on documentation and CQA data to inject domain knowledge (Yang et al., 2023), reinforcement learning from human feedback (RLHF) using community voting signals as rewards (Gorbatovski and Kovalchuk, 2024), and aligning LLMs via multi-perspective ranking and contrastive learning to better satisfy diverse user preferences (Yang et al., 2025).

Despite strong performance on static benchmarks, existing CQA methods largely overlook the dynamic nature of community content and the inconsistency of historical responses.

2.2 Retrieval-Augmented Generation

Retrieval-Augmented Generation (RAG) (Lewis et al., 2021) has become a promising framework for enhancing LLMs with external knowledge access. RAG augments the input to generation models by retrieving relevant documents, which helps mitigate hallucinations, knowledge staleness, and limited interpretability (Gao et al., 2024). RAG has shown effectiveness across a range of knowledge-intensive tasks, including open-domain QA and summarization (Siriwardhana et al., 2022). Although current RAG implementations rely on static corpora, its retrieval-generation paradigm naturally lends itself to addressing challenges in real-time industrial CQA by enabling fast access to dynamically updated data and supporting the design of customizable retrieval strategies.

3 Task Definition

We formally define the task of answering real-time community questions with external knowledge. Given a collection of documents $D = \{d_i\}_{i=1}^{|D|}$ as the external knowledge, the community questions arrive as a continuous stream. Suppose that at this moment, we have collected the community history denoted as $H = \{(q_i, a_i)\}_{i=1}^{|H|}$ where the historical questions are associated with the historical responses. When there is a new question q , we can extract the answer \hat{a} either from external knowledge or from the community history to ensure that

\hat{a} equals the ground truth answer a^* . Furthermore, we should determine how to organize (q, \hat{a}) in the community history H to meet memory constraints and accommodate future follow-up questions.

4 Our System: ComRAG

For real-time CQA in industry, questions can be answered using external knowledge, community history, or a combination of both. Specifically, the external knowledge is static and filled with domain-specific information. The community history is dynamic and characterized by accumulated question-answer pairs. To improve retrieval efficiency and response quality, ComRAG is built upon the RAG framework and contains a *static knowledge vector store* and two *dynamic CQA vector stores*. The knowledge vector store retrieves relevant domain-specific documents, while the CQA vector stores dynamically maintain and retrieve historical community QA pairs. They work together to either retrieve or generate answers based on high-quality community QA pairs, or alternatively generate answers by avoiding low-quality QA pairs and incorporating external knowledge as additional context. The overview of ComRAG is shown in Figure 1.

4.1 Static Knowledge Vector Store

Following existing RAG methods (Gao et al., 2024; Guo et al., 2024), we embed the documents in the external knowledge via a static vector database. Specifically, each document is converted into a vector using an embedding model, allowing us to retrieve relevant documents by computing their similarity to the embedded representation of a given question. The above procedure can be formulated as follows:

$$\mathcal{V}_{kb} = \{(d_i, \text{Emb}(d_i)) \mid d_i \in \mathcal{D}\},$$

$$\{\hat{d}_i\}_{i=1}^k = \arg \text{top-}k_{d_i \in \mathcal{D}} \text{CosSim}(\text{Emb}(q), \text{Emb}(d_i)).$$

Here, $\{\hat{d}_i\}_{i=1}^k$ are the retrieved knowledge documents that serve as evidence of the question. We define *arg top-k* as the function that returns the documents with the top- k similarities. Then a frozen LLM generates the predicted answers given the instruction and retrieved $\{\hat{d}_i\}_{i=1}^k$ as follows:

$$\hat{a} = \text{LLM}(q, \{\hat{d}_i\}_{i=1}^k).$$

When a question-answer pair is produced, we measure the quality of the answer using a predefined metric, which can be based on either manual or automatic scoring. Automatic scoring can

be implemented using either LLMs or various evaluation metrics, or by combining both. A score is computed for each QA pair, denoted as $s = \text{Scorer}(q, \hat{a})$.

4.2 Dynamic CQA Vector Store

While the static knowledge vector store effectively handles domain-specific questions in industrial settings, it fails to reflect the answers of varying quality in the community history. Hence, we propose a dynamic Community Question-Answer (CQA) vector store consisting of a *high-quality CQA vector store* and a *low-quality CQA vector store*, which are based on a *centroid-based memory mechanism*.

Centroid-Based Memory Mechanism. Since the community history H will continuously increase, we adopt a centroid-based memory mechanism to maintain it within a limited memory size. This mechanism partitions similar historical questions into clusters and only retains the representative questions in each cluster to avoid memory overflow. Formally, assume we have m clusters $\{C_1, C_2, \dots, C_m\}$ in the memory. Each cluster contains a set of questions belonging to the same topic $C = \{q_i\}_{i=1}^{|C|}$ and its centroid is computed as:

$$\mathbf{c} = \frac{1}{|C|} \sum_{q_i \in C} \text{Emb}(q_i).$$

Given a new question, we embed it into a vector, then assign it to the most relevant cluster if the similarity exceeds a threshold τ :

$$C = \arg \max_{C \in \{C_1, C_2, \dots, C_m\}} \text{CosSim}(\text{Emb}(q), \mathbf{c}).$$

The centroid of cluster C is then updated accordingly.

If the similarity is smaller than τ , we believe a question should be derived from a new topic, a new cluster is created in memory with an initial centroid represented as $\mathbf{c} = \text{Emb}(q)$.

To maintain a fixed memory size, we allow removing questions from a cluster when necessary. When a newly assigned question is highly similar to an existing one in the cluster (similarity $> \delta$), we compare their answer quality. If the existing question’s answer has a lower evaluation score, we consider the new question to be of higher quality. Hence, we remove the existing question from the cluster and replace it with the new one. In this case, we can effectively control the size of each cluster and avoid memory overflow due to the accumulated questions.

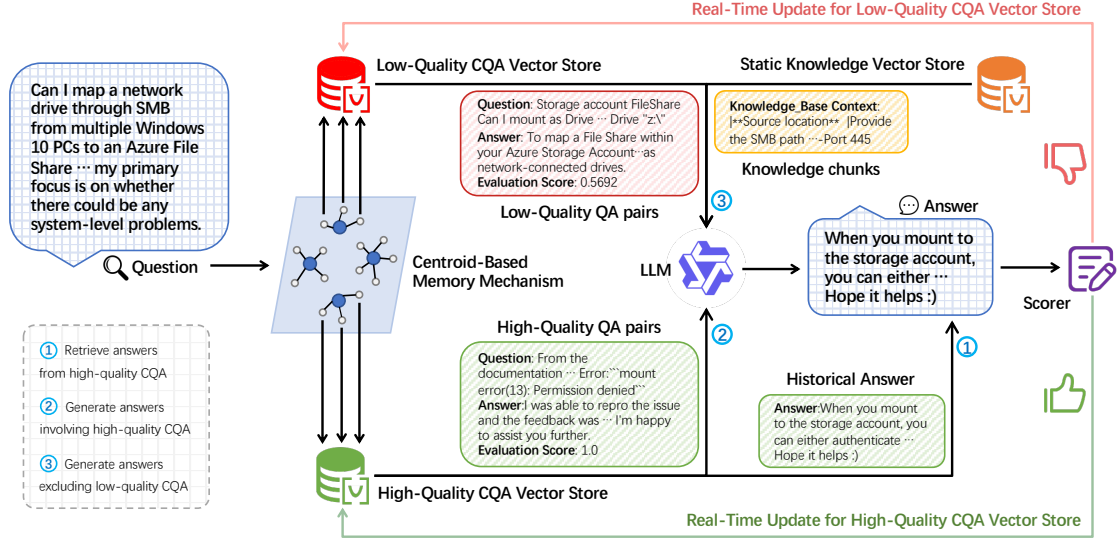


Figure 1: ComRAG architecture for real-time CQA. The system integrates a static knowledge vector store and two dynamic CQA vector stores (high- and low-quality), with the latter managed via a centroid-based memory mechanism. When a question is posed, ComRAG follows one of three paths: ① retrieving answers from the high-quality CQA vector store, ② generating answers using high-quality CQA, or ③ generating answers by excluding low-quality CQA and incorporating static knowledge. Real-time updates to either the high- or low-quality CQA vector store ensure efficient memory management and scalable deployment.

Based on this principle, we introduce two vector stores to maintain high-quality and low-quality community history, enabling reflection in follow-up real-time CQA. For high-quality CQA, our system continues to generate similar answers for subsequent questions. For low-quality CQA, our system avoids generating similar answers for follow-up questions.

High-Quality CQA Vector Store. We leverage the evaluation score of the answers to decide whether QA pairs are updated into the high-quality or low-quality CQA vector stores. The high-quality CQA vector store maintains historical QA pairs with scores above γ , where each answer and its score are stored as metadata.

$$\mathcal{V}_{\text{high}} = \{(q, \text{Emb}(q), \hat{a}, s) \mid s \geq \gamma\}.$$

To maintain the vector store in a controllable size, we apply the centroid-based memory mechanism to cluster the high-quality question-answer pairs.

Low-Quality CQA Vector Store. Similarly, we maintain a low-quality CQA vector store consisting of question-answer pairs with a score lower than γ and apply the centroid-based memory mechanism.

$$\mathcal{V}_{\text{low}} = \{(q, \text{Emb}(q), \hat{a}, s) \mid s < \gamma\}.$$

4.3 Query and Update

In the query phase, ComRAG retrieves relevant historical QA pairs and domain knowledge to answer the current question. The system supports three query strategies:

- ① **Retrieve answers from high-quality CQA.** If the question already exists in the high-quality CQA vector store—that is, the most similar historical question $\tilde{q} = \arg \max_{q_i \in \mathcal{V}_{\text{high}}} \text{CosSim}(\text{Emb}(q), \text{Emb}(q_i))$ satisfies $\text{CosSim} \geq \delta$ —we directly reuse the corresponding historical answer:

$$\hat{a} = \mathcal{V}_{\text{high}}[\tilde{q}].\hat{a}$$

where $\mathcal{V}_{\text{high}}[\tilde{q}]$ returns the stored tuple (\tilde{q}, \hat{a}, s) .

- ② **Generate answers involving high-quality CQA.** If the similarity satisfies $\tau \leq \text{CosSim}(\text{Emb}(q), \text{Emb}(\tilde{q})) < \delta$, the retrieved QA pairs still serve as useful references for LLM generation:

$$\hat{a} = \text{LLM}(q, \{\mathcal{V}_{\text{high}}[\tilde{q}_i]\}_{i=1}^k)$$

where each $\mathcal{V}_{\text{high}}[\tilde{q}_i]$ returns the tuple $(\tilde{q}_i, \hat{a}_i, s_i)$, which is used as input evidence. Similar to document retrieval, we select the top- k relevant question-answer pairs as context.

- ③ **Generate answers involving low-quality CQA and external knowledge.** If no sufficiently similar question is found in the high-quality CQA vector store, we retrieve evidence from both the static knowledge vector store and the low-quality CQA vector store to guide the LLM away from repeating inaccurate historical answers:

$$\hat{a} = \text{LLM}(q, \{\hat{d}_i\}_{i=1}^k, \{\mathcal{V}_{\text{low}}[\tilde{q}_j]\}_{j=1}^k)$$

where each $\mathcal{V}_{\text{low}}[\tilde{q}_j]$ returns $(\tilde{q}_j, \hat{a}_j, s_j)$ for contrastive referencing.

After the predicted answers are generated, we score each answer as described in Section 4.1, and assign the resulting question-answer pair to either the high-quality or low-quality CQA vector store, as detailed in Section 4.2. Pseudocode for both the query and update phases is provided in Appendix B.

4.4 Adaptive Temperature Tuning for Generation

ComRAG introduces an *adaptive temperature tuning mechanism* to dynamically adjust the LLM’s decoding temperature, balancing response diversity and consistency. Specifically, for the evidence retrieved from high-quality or low-quality vector stores, we store the answer scores as metadata. If these scores exhibit low variance, this indicates that the historical answers are highly similar; thus we prompt the LLMs with a higher temperature to encourage exploration. In contrast, when the scores have high variance, we use a lower temperature to ensure consistency with reliable historical answers.

Assume we collect l QA pairs as evidence for prompting, each with an annotated score. After sorting these scores in ascending order (s_1, s_2, \dots, s_l) , we define the adaptive temperature (with scaling factor k) as:

$$T(\Delta) = |\exp(-k \cdot \min_{1 \leq i \leq l-1} (s_{i+1} - s_i))|_{[T_{\min}, T_{\max}]},$$

where $|\bullet|_{[T_{\min}, T_{\max}]}$ is a clamp function restricting the temperature to the predefined range $[T_{\min}, T_{\max}]$. Then $T(\Delta)$ is set as the argument for the final answer generation.

5 Experimental Setup

5.1 Dataset Collection

We conduct experiments on three community QA datasets: Microsoft QA (MSQA)(Yang et al.,

2023), ProCQA(Li et al., 2024), and PolarDBQA. MSQA is a question-answering dataset collected from the Microsoft Q&A forum. ProCQA consists of structured programming QA pairs extracted from StackOverflow. PolarDBQA is constructed from Alibaba Cloud’s official PolarDB documentation, containing question-answer pairs generated by LLM to simulate typical user inquiries in specialized database domains.

For each dataset, there is an associated set of documents as external knowledge. MSQA utilizes Azure documentation. ProCQA provides an official external knowledge corpus for retrieval. We collect data and construct the PolarDBQA dataset from the Alibaba Cloud platform³, where PolarDB documentation is utilized as external knowledge.

The question-answer pairs in training sets are initially stored as high-quality CQA vectors. To simulate real-time CQA where questions arrive sequentially, we paraphrase each question in the test sets into multiple versions using LLMs. We then shuffle all the questions and split them into several iterations for evaluation. An overview of the datasets is provided in Appendix A.

5.2 Baselines

We use qwen2.5:14b-instruct-fp16(Bai et al., 2023) as the LLM and compare **ComRAG** with several baselines differing in the external context provided. **Raw LLM** generates answers without any additional input. **BM25** and **DPR** retrieve historical QA pairs using BM25 (Robertson and Zaragoza, 2009) and DPR (Karpukhin et al., 2020), respectively. **Vanilla RAG** uses only documents retrieved from the static knowledge vector store. **RAG+BM25** and **RAG+DPR** extend Vanilla RAG by additionally retrieving historical QA pairs via BM25 or DPR. **LLM+EXP** (Yang et al., 2023) follows MSQA’s expert-guided interaction paradigm by aligning knowledge with a domain-specific model and incorporating it into the LLM.

5.3 Evaluation Metrics

We evaluate the generated answers using both lexical and semantic metrics. For lexical alignment, we use **BLEU** (Papineni et al., 2002) and **ROUGE-L** (Lin, 2004) to measure n-gram overlap with reference answers. For semantic evaluation, we adopt **BERT-Score** and report the F1 score. Additionally,

³<https://docs.polardbpg.com/1733726429035/>,
<https://help.aliyun.com/zh/polardb/polardb-for-xscale/>

			MSQA				
Methods	Doc	ComQA	Avg Time	BERT-Score	SIM	BLEU	ROUGE-L
Raw LLM	✗	✗	12.70	54.70	80.58	10.46	15.07
BM25	✗	✓	15.07	54.98	80.41	10.44	15.03
DPR	✗	✓	13.91	55.01	80.67	10.65	15.09
Vanilla RAG	✓	✗	13.86	54.43	80.73	10.09	14.82
RAG+BM25	✓	✓	15.47	54.95	80.79	10.31	15.09
RAG+DPR	✓	✓	14.08	55.01	80.50	10.62	15.21
LLM+EXP	✓	✓	20.23	55.79	76.70	11.13	15.66
Ours	✓	✓	11.60	<u>55.47</u>	94.70	11.61	16.66

			ProCQA				
Methods	Doc	ComQA	Avg Time	BERT-Score	SIM	BLEU	ROUGE-L
Raw LLM	✗	✗	12.77	56.16	74.88	12.17	15.49
BM25	✗	✓	13.99	56.21	75.68	11.41	15.81
DPR	✗	✓	14.11	56.08	75.73	11.20	15.56
Vanilla RAG	✓	✗	16.97	57.76	75.59	14.13	18.10
RAG+BM25	✓	✓	14.11	56.20	75.30	11.22	15.69
RAG+DPR	✓	✓	13.79	56.06	74.83	10.62	15.21
LLM+EXP	✓	✓	22.69	<u>58.40</u>	67.78	<u>14.36</u>	16.70
Ours	✓	✓	10.42	58.41	95.31	14.37	18.13

			PolarDBQA				
Methods	Doc	ComQA	Avg Time	BERT-Score	SIM	BLEU	ROUGE-L
Raw LLM	✗	✗	4.63	60.34	93.51	1.60	9.40
BM25	✗	✓	5.54	63.39	94.06	4.42	20.15
DPR	✗	✓	5.45	64.01	94.08	5.72	21.52
Vanilla RAG	✓	✗	9.67	64.78	92.27	5.21	23.45
RAG+BM25	✓	✓	10.60	65.86	92.98	6.71	24.83
RAG+DPR	✓	✓	22.98	66.55	93.45	6.66	28.47
LLM+EXP	✓	✓	8.15	<u>67.00</u>	90.11	8.04	33.61
Ours	✓	✓	3.55	67.39	96.04	<u>7.81</u>	30.19

Table 1: Performance comparison of different methods. "Doc" refers to documents retrieved from the static knowledge vector store, and "ComQA" refers to historical QA pairs retrieved from dynamic CQA vector stores. ✓ indicates the source is retrieved; ✗ indicates it is not.

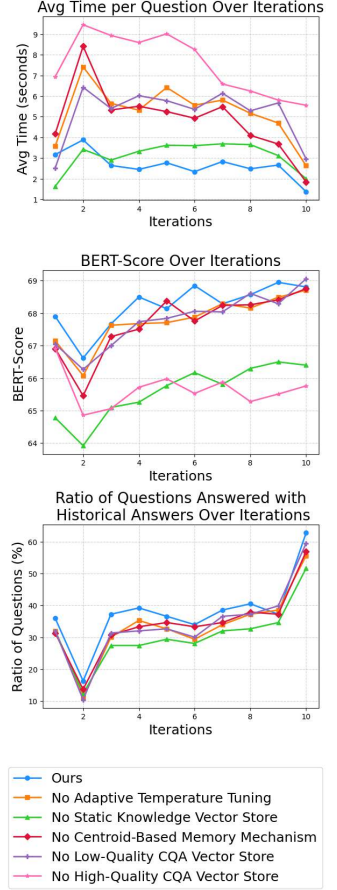


Figure 2: Ablation study on PolarDBQA under a 10-round iterative evaluation setting.

we compute the cosine similarity between the embeddings of the generated and reference answers, denoted as the **SIM** metric. We also report **Avg Time**, defined as the average processing time per question in seconds.

5.4 Implementation Details

We use the sentence embedding model bge-large-en-v1.5-f32 (Xiao et al., 2023) for MSQA and ProCQA, and bge-large-zh-v1.5-f32 for PolarDBQA. SIM is computed using GPT-2 (Radford et al., 2019), following MSQA. For re-ranking, we use BAAI/bge-reranker-large. Vector storage and retrieval are managed with ChromaDB v0.6.3. All experiments are run on a Linux server with PyTorch 2.6.0 (CUDA 12.4) and Python 3.10.16. For ComRAG, the core hyperparameters τ , δ , and γ , introduced in Sections 4.2 and 4.3, are set to (0.75, 0.9, 0.6) for MSQA and ProCQA, and (0.75, 0.8, 0.7) for PolarDBQA. The scoring function $\text{Scorer}(\cdot)$ used to evaluate answer quality

is implemented via BERT-Score (Zhang et al., 2020), measuring semantic similarity between generated answers and references. For the adaptive temperature tuning mechanism, we set $k=250$, $T_{\min}=0.7$ and $T_{\max}=1.2$.

6 Results and Analysis

6.1 Main Results

To ensure a fair comparison under the same initial conditions of real-time QA, we evaluate all baselines and ComRAG on the first iteration of the sequential question-answering setting described in Section 5.1. As shown in Table 1, ComRAG consistently outperforms all baselines in both answer quality and response efficiency. Compared to the second-best method on each dataset, ComRAG achieves improvements in the SIM metric of **2.1%-25.9%**, and reduces average query latency by **8.7%-23.3%**. These results demonstrate ComRAG’s effectiveness in balancing response quality and latency in real-time CQA.

6.2 Ablation Study

We conduct ablation experiments on PolarDBQA under the iterative evaluation setting described in Section 5.1, which simulates real-time CQA by processing questions over multiple rounds. We evaluate the effect of removing each module introduced in Section 4. Removing any module increased latency and reduced accuracy, which highlights their necessity as illustrated in Figure 2. The high-quality CQA vector store had the most significant impact, delaying responses by **4.9s** and lowering BERT-Score by **2.6**. Similarly, removing the centroid-based memory mechanism increased delays by **2.2s** and reduced BERT-Score by **0.5**, demonstrating its importance in dynamically updating historical QA pairs. Additionally, removing the static knowledge vector store and adaptive temperature tuning mechanism significantly decreased the proportion of directly answerable test questions, indicating that these modules play a crucial role in improving response quality, thereby indirectly enhancing answer reuse efficiency.

6.3 Real-time QA Evaluation

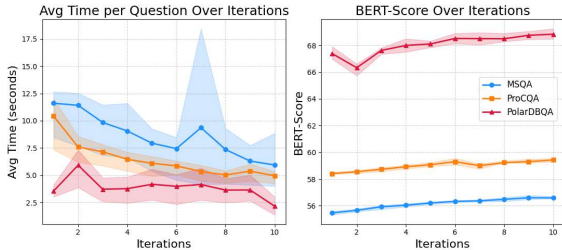


Figure 3: Avg Time and BERT-Score over iterations. ComRAG improves efficiency and response quality as historical QA interactions accumulate.

We further evaluate ComRAG under the iterative question-answering setting, where questions arrive in sequential batches. As historical QA records accumulate over iterations, ComRAG exhibits consistent improvements in both efficiency and answer quality. As shown in Figure 3, query latency drops substantially, with the most notable reduction on ProCQA: average processing time decreases from 10.42s in the first iteration to 4.95s in the final iteration, yielding a **52.5%** improvement. Alongside these efficiency gains, response quality also improves, with BERT-Score increasing steadily—most significantly on MSQA, where it rises by **2.25%** over time. These results highlight ComRAG’s effectiveness in real-time applications,

balancing low-latency generation with progressive quality refinement.

6.4 Effect of Memory Size

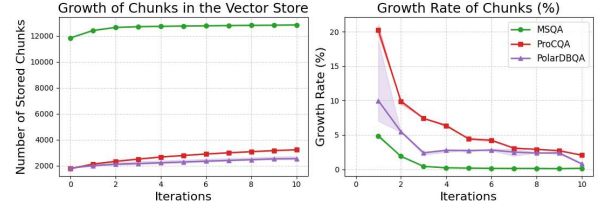


Figure 4: Total stored chunks and growth rate over iterations across all dynamic CQA vector stores. ComRAG efficiently manages memory, preventing excessive storage expansion.

To evaluate ComRAG’s memory adaptation, we analyze the growth rate of stored chunks across dynamic CQA vector stores over iterations. As shown in Figure 4, the growth rate peaks early and then gradually declines as the system stabilizes. Notably, ProCQA shows the most significant initial expansion, with a **20.23%** increase in iteration 1, dropping to just **2.06%** by iteration 10. This sharp decline suggests that most necessary knowledge is integrated early, which helps reduce redundant storage in later iterations and stabilizes memory growth over time. These results demonstrate that ComRAG effectively manages historical QA storage, preventing uncontrolled expansion while maintaining efficient retrieval. Such controlled memory usage contributes to scalable deployment in real-time industrial CQA systems.

7 Conclusion

We present **ComRAG**, a retrieval-augmented generation framework for real-time industrial CQA. By combining static domain knowledge with dynamic QA history, ComRAG improves response accuracy, latency, and adaptability. It employs a centroid-based memory mechanism to control storage growth and an adaptive temperature tuning mechanism to balance consistency and diversity of generated answers. Experiments on multiple CQA benchmarks demonstrate its practical effectiveness in retrieval and generation for real-world CQA scenarios. Furthermore, ComRAG’s modular design supports scalable deployment by enabling replacement of core components such as the LLM backbone, scoring strategy, and retrieval modules, allowing it to accommodate different computational budgets and deployment environments.

Limitations

While ComRAG demonstrates strong performance in real-time industrial CQA, several limitations remain. First, the centroid-based memory mechanism relies on fixed similarity thresholds and does not consider topic relevance or usage frequency, which may hinder memory efficiency in dynamic environments. Second, low-quality QA pairs are handled via simple avoidance through prompt design. More advanced filtering or correction mechanisms may enhance reliability. Lastly, the current query and generation paths are rule-based. Incorporating learning-based routing strategies could improve adaptability to diverse question types and knowledge needs.

Acknowledgments

The authors would like to thank the anonymous reviewers for their insightful comments. This work is supported by the National Key Research & Develop Plan (Project No. 2023YFF0725100) and Natural Science Foundation of China (Project No. 62137001 and U23A20298).

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altmenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. [arXiv preprint arXiv:2303.08774](#).
- Maddalena Amendola, Andrea Passarella, and Raffaele Perego. 2024. Leveraging topic specificity and social relationships for expert finding in community question answering platforms. [arXiv preprint arXiv:2407.04018](#).
- Arian Askari, Zihui Yang, Zhaochun Ren, and Suzan Verberne. 2024. Answer retrieval in legal community question answering. In *European Conference on Information Retrieval*, pages 477–485. Springer.
- Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, Binyuan Hui, Luo Ji, Mei Li, Junyang Lin, Runji Lin, Dayiheng Liu, Gao Liu, Chengqiang Lu, Keming Lu, Jianxin Ma, Rui Men, Xingzhang Ren, Xuancheng Ren, Chuanqi Tan, Sinan Tan, Jianhong Tu, Peng Wang, Shijie Wang, Wei Wang, Shengguang Wu, Benfeng Xu, Jin Xu, An Yang, Hao Yang, Jian Yang, Shusheng Yang, Yang Yao, Bowen Yu, Hongyi Yuan, Zheng Yuan, Jianwei Zhang, Xingxuan Zhang, Yichang Zhang, Zhenru Zhang, Chang Zhou, Jingren Zhou, Xiaohuan Zhou, and Tianhang Zhu. 2023. [Qwen technical report](#). [Preprint](#), arXiv:2309.16609.
- DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, Bing Xue, Bingxuan Wang, Bochao Wu, Bei Feng, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, Damai Dai, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, Guangbo Hao, Guanting Chen, Guowei Li, H. Zhang, Han Bao, Hanwei Xu, Haocheng Wang, Honghui Ding, Huajian Xin, Huazuo Gao, Hui Qu, Hui Li, Jianzhong Guo, Jiashi Li, Jiawei Wang, Jingchang Chen, Jingyang Yuan, Junjie Qiu, Junlong Li, J. L. Cai, Jiaqi Ni, Jian Liang, Jin Chen, Kai Dong, Kai Hu, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang, Lecong Zhang, Liang Zhao, Litong Wang, Liyue Zhang, Lei Xu, Leyi Xia, Mingchuan Zhang, Minghua Zhang, Minghui Tang, Meng Li, Miaojun Wang, Mingming Li, Ning Tian, Panpan Huang, Peng Zhang, Qiancheng Wang, Qinyu Chen, Qiusi Du, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang, R. J. Chen, R. L. Jin, Ruyi Chen, Shanghao Lu, Shangyan Zhou, Shanhuang Chen, Shengfeng Ye, Shiyu Wang, Shuiping Yu, Shunfeng Zhou, Shuting Pan, S. S. Li, Shuang Zhou, Shaoqing Wu, Shengfeng Ye, Tao Yun, Tian Pei, Tianyu Sun, T. Wang, Wangding Zeng, Wanjia Zhao, Wen Liu, Wenfeng Liang, Wenjun Gao, Wenqin Yu, Wentao Zhang, W. L. Xiao, Wei An, Xiaodong Liu, Xiaohan Wang, Xiaokang Chen, Xiaotao Nie, Xin Cheng, Xin Liu, Xin Xie, Xingchao Liu, Xinyu Yang, Xinyuan Li, Xuecheng Su, Xuheng Lin, X. Q. Li, Xiangyue Jin, Xiaojin Shen, Xiaosha Chen, Xiaowen Sun, Xiaoxiang Wang, Xinnan Song, Xinyi Zhou, Xianzu Wang, Xinxia Shan, Y. K. Li, Y. Q. Wang, Y. X. Wei, Yang Zhang, Yanhong Xu, Yao Li, Yao Zhao, Yaofeng Sun, Yaohui Wang, Yi Yu, Yichao Zhang, Yifan Shi, Yiliang Xiong, Ying He, Yishi Piao, Yisong Wang, Yixuan Tan, Yiyang Ma, Yiyuan Liu, Yongqiang Guo, Yuan Ou, Yuduan Wang, Yue Gong, Yuheng Zou, Yujia He, Yunfan Xiong, Yuxiang Luo, Yuxiang You, Yuxuan Liu, Yuyang Zhou, Y. X. Zhu, Yanhong Xu, Yanping Huang, Yaohui Li, Yi Zheng, Yuchen Zhu, Yunxian Ma, Ying Tang, Yukun Zha, Yuting Yan, Z. Z. Ren, Zehui Ren, Zhangli Sha, Zhe Fu, Zhean Xu, Zhenda Xie, Zhengyan Zhang, Zhewen Hao, Zhicheng Ma, Zhigang Yan, Zhiyu Wu, Zihui Gu, Zijia Zhu, Zijun Liu, Zilin Li, Ziwei Xie, Ziyang Song, Zizheng Pan, Zhen Huang, Zhipeng Xu, Zhongyu Zhang, and Zhen Zhang. 2025. [DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning](#). [Preprint](#), arXiv:2501.12948.
- Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, Meng Wang, and Haofen Wang. 2024. [Retrieval-augmented generation for large language models: A survey](#). [Preprint](#), arXiv:2312.10997.
- Shima Ghasemi and Azadeh Shakery. 2024. [Harnessing the power of metadata for enhanced question retrieval](#)

- in community question answering. [IEEE Access](#), 12:65768–65779.
- Alexey Gorbatoyski and Sergey Kovalchuk. 2024. Reinforcement learning for question answering in programming domain using public community scoring as a human feedback. In [Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems](#), pages 2294–2296.
- Zirui Guo, Lianghao Xia, Yanhua Yu, Tu Ao, and Chao Huang. 2024. [Lightrag: Simple and fast retrieval-augmented generation](#). Preprint, arXiv:2410.05779.
- Xinghang Hu. 2023. Enhancing answer selection in community question answering with pre-trained and large language models. [arXiv preprint arXiv:2311.17502](#).
- Vladimir Karpukhin, Barlas Oğuz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen tau Yih. 2020. [Dense passage retrieval for open-domain question answering](#). Preprint, arXiv:2004.04906.
- 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. 2021. [Retrieval-augmented generation for knowledge-intensive nlp tasks](#). Preprint, arXiv:2005.11401.
- Zehan Li, Jianfei Zhang, Chuantao Yin, Yuanxin Ouyang, and Wenge Rong. 2024. Procqa: A large-scale community-based programming question answering dataset for code search. In [Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation \(LREC-COLING 2024\)](#), pages 13057–13067.
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In [Text Summarization Branches Out](#), pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In [Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics](#), pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. [OpenAI blog](#), 1(8):9.
- Stephen Robertson and Hugo Zaragoza. 2009. [The probabilistic relevance framework: Bm25 and beyond](#). [Found. Trends Inf. Retr.](#), 3(4):333–389.
- Pradeep Kumar Roy, Sunil Saumya, Jyoti Prakash Singh, Snehasish Banerjee, and Adnan Gutub. 2023. Analysis of community question-answering issues via machine learning and deep learning: State-of-the-art review. [CAAI Transactions on Intelligence Technology](#), 8(1):95–117.
- Shamane Siriwardhana, Rivindu Weerasekera, Elliott Wen, Tharindu Kaluarachchi, Rajib Rana, and Suranga Nanayakkara. 2022. [Improving the domain adaptation of retrieval augmented generation \(rag\) models for open domain question answering](#). Preprint, arXiv:2210.02627.
- Shitao Xiao, Zheng Liu, Peitian Zhang, and Niklas Muennighoff. 2023. [C-pack: Packaged resources to advance general chinese embedding](#). Preprint, arXiv:2309.07597.
- Fangkai Yang, Pu Zhao, Zezhong Wang, Lu Wang, Bo Qiao, Jue Zhang, Mohit Garg, Qingwei Lin, Saravan Rajmohan, and Dongmei Zhang. 2023. Empower large language model to perform better on industrial domain-specific question answering. In [Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: Industry Track](#), pages 294–312.
- Hongyu Yang, Jiahui Hou, Liyang He, and Rui Li. 2025. [Multi-perspective preference alignment of LLMs for programming-community question answering](#). In [Proceedings of the 31st International Conference on Computational Linguistics](#), pages 1667–1682, Abu Dhabi, UAE. Association for Computational Linguistics.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. [Bertscore: Evaluating text generation with bert](#). Preprint, arXiv:1904.09675.

A Dataset Overview

	MSQA	ProCQA	PolarDBQA
Number of KB Chunks	557,235	14,478	1,403
Train Set Size	9,518	3,107	1,395
Test Set Size	571	346	153

Table 2: Overview of datasets used in our experiments. “Number of KB Chunks” refers to the total number of knowledge base document chunks used as external context. “Train Set Size” denotes the number of QA pairs initially loaded into the high-quality CQA vector store. “Test Set Size” is the total number of test questions evaluated. For ProCQA, we use its Lisp programming language subset.

B Query and Update Algorithms

Algorithm 1 outlines the complete query phase in ComRAG. Given an input question, the system first checks whether any high-quality QA pair in the CQA vector store can be directly reused. If not, it retrieves relevant high-quality QA pairs as references for generation. If no suitable high-quality QA pairs are found, counter-examples are retrieved from the low-quality store, and relevant documents are retrieved from the knowledge base to guide answer generation.

Algorithm 1 Query Phase in ComRAG

Input: Question q , thresholds τ , δ , and γ , high- and low-quality CQA vector stores V_{high} and V_{low} , static knowledge vector store D , high-quality centroid vector store C_{high} , low-quality centroid vector store C_{low} , number of retrieved candidates k

Output: Answer \hat{a}

```

1:  $\mathbf{q} = \text{Emb}(q)$ 
2:  $\hat{c}_{\text{high}} \leftarrow \arg \text{top-}k\text{CosSim}(\mathbf{q}, C_{\text{high}})$ 
3:  $\hat{C}_{\text{high}} \leftarrow \text{RetrieveCQA}(V_{\text{high}}, \hat{c}_{\text{high}})$ 
4: if  $\max(\hat{C}_{\text{high},i}.\text{sim}) \geq \delta$  then
5:   return  $\hat{a} \leftarrow \hat{C}_{\text{high},i}.\text{answer}$ 
6: else if  $\tau \leq \hat{C}_{\text{high},i}.\text{sim} < \delta$  then
7:    $\hat{a} \leftarrow \text{LLM}(q, \{\hat{C}_{\text{high},i}\})$ 
8: else
9:    $\hat{c}_{\text{low}} \leftarrow \arg \text{top-}k\text{CosSim}(\mathbf{q}, C_{\text{low}})$ 
10:   $\hat{C}_{\text{low}} \leftarrow \text{RetrieveCQA}(V_{\text{low}}, \hat{c}_{\text{low}})$ 
11:   $\{\hat{C}_{\text{low},i}\}_{i=1}^k \leftarrow \arg \text{top-}k\text{CosSim}(\mathbf{q}, \hat{C}_{\text{low}})$ 
12:   $\hat{D} \leftarrow \arg \text{top-}k\text{CosSim}(\mathbf{q}, D)$ 
13:   $\hat{a} = \text{LLM}(q, \{\hat{C}_{\text{low},i}\}_{i=1}^k, \hat{D})$ 
14: end if
15: return  $\hat{a}$ 

```

Algorithm 2 describes the complete process of the update phase. After evaluation, ComRAG determines whether the new QA pair should replace an existing entry in the CQA vector store and updates both the vector store and the corresponding centroid. Otherwise, it adds the QA pair to an existing or newly created cluster.

Algorithm 2 Update Phase in ComRAG

Input: Evaluation result (q,a,s) with the question-answer pair and score, thresholds τ , δ , γ , high- and low-quality CQA vector stores V_{high} and V_{low} , C is the cluster and c is the centroid vector

```

1:  $q = \text{Emb}(q)$ 
2:  $\hat{V} \leftarrow V_{\text{high}}$  if  $s \geq \gamma$  else  $V_{\text{low}}$ 
3: if  $\max(\text{CosSim}(\mathbf{q}, \text{Emb}(\hat{V}_i.q))) \geq \delta$  then
4:   if  $s > \hat{V}_i.\text{score}$  then
5:      $\hat{V}.\text{add}((q, \text{Emb}(q), a, s))$ 
6:      $\hat{V}.\text{delete}(\hat{V}_i)$ 
7:      $\hat{C} \leftarrow \text{ClusterOf}(\hat{V}_i)$ 
8:      $\hat{C}.\text{delete}(\hat{V}_i.q)$ 
9:      $\hat{C}.\text{add}(q)$ 
10:     $\hat{c} \leftarrow \frac{1}{|\hat{C}|} \sum_{q' \in \hat{C}} \text{Emb}(q')$ 
11:   else
12:     Discard  $(q, a, s)$ 
13:   end if
14: else
15:   if  $\max(\text{CosSim}(\mathbf{q}, \text{Emb}(c_i))) \geq \tau$  then
16:      $C_i.\text{append}(q)$ 
17:      $c_i \leftarrow \frac{1}{|C_i|} \sum_{q' \in C_i} \text{Emb}(q')$ 
18:   else
19:      $C_{\text{new}} \leftarrow \{q\}$ 
20:      $c_{\text{new}} = \text{Emb}(q)$ 
21:   end if
22: end if
23: return

```

C Impact of Hyperparameters on ComRAG Performance over Iterations

We conduct a series of ablation experiments on the PolarDBQA dataset to analyze the sensitivity of ComRAG to three key hyperparameters: τ , δ , and γ . Their respective roles in the query and update phases are summarized in Table 3. In each experiment, we vary one hyperparameter while keeping the other two fixed at their default values ($\tau=0.75$, $\delta=0.8$, $\gamma=0.7$). For each setting, we track system performance over 10 iterations using four metrics: Avg Time, BERT-Score, ratio of historical answer reuse, and vector store chunk growth rate.

Hyperparameter	Role in ComRAG
τ	Used in both the update and query phases: <ul style="list-style-type: none"> • In the update phase, it determines whether a new question is similar enough to be assigned to an existing cluster in the centroid-based memory. • In the query phase, it sets the lower bound for retrieving similar high-quality QA pairs as reference for generation.
δ	Used in both the query and update phases to identify near-duplicate questions: <ul style="list-style-type: none"> • In the query phase, it decides whether to directly reuse historical answers. • In the update phase, it determines whether a newly added QA pair should replace a lower-quality one within a cluster.
γ	Used in the update phase to classify QA pairs based on answer quality: <ul style="list-style-type: none"> • QA pairs with scores $\geq \gamma$ are stored in the high-quality CQA vector store; others go into the low-quality CQA vector store.

Table 3: Roles of hyperparameters τ , δ , and γ in different phases of ComRAG.

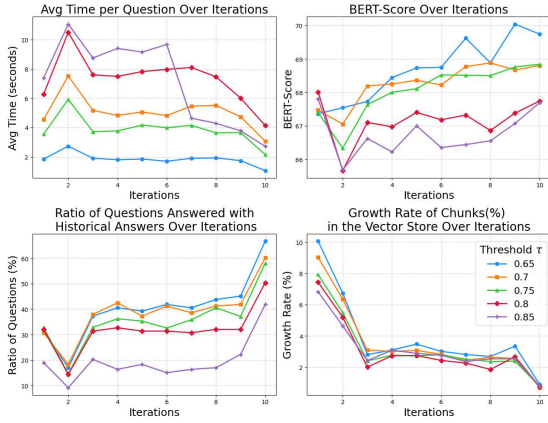


Figure 5: Impact of similarity threshold τ on ComRAG performance over iterations (with $\delta = 0.8$, $\gamma = 0.7$ fixed).

Impact of τ . As shown in Figure 5, lower values of τ (e.g., 0.65) lead to more aggressive matching with historical QA pairs, resulting in a higher reuse ratio of **66.67%** and reduced query latency down to **1.06s** by the final iteration. This also slows the growth rate of stored chunks, reflecting more efficient memory usage. Conversely, higher thresholds (e.g., 0.85) restrict reuse opportunities, leading to increased latency and memory expansion.

However, overly small τ values may introduce loosely related historical answers, slightly degrading generation quality as indicated by BERT-Score fluctuations. The default $\tau = 0.75$ provides a strong trade-off—ensuring stable semantic quality (e.g., BERT-Score **68.84**), moderate memory growth, and high efficiency. These findings highlight the role of τ in balancing reuse, precision, and storage efficiency.

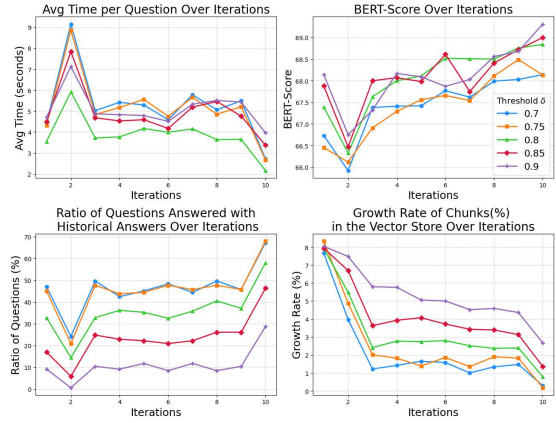


Figure 6: Impact of reuse threshold δ on ComRAG performance over iterations (with $\tau = 0.75$, $\gamma = 0.7$ fixed).

Impact of δ . Figure 6 shows the impact of δ , which controls the threshold for answer reuse and replacement. Setting $\delta=0.8$ yields the best balance across metrics, achieving the highest BERT-Score (68.84), lowest latency (2.16s), and stable chunk growth. A lower δ (e.g., 0.7) increases reuse ratio (67.32%) but risks low-quality matches. In contrast, higher values (e.g., 0.9) overly restrict reuse, leading to more generation, higher latency (3.96s), and greater chunk accumulation. This highlights the need for a moderate reuse threshold to ensure both efficiency and quality.

Impact of γ . As shown in Figure 7, lower values of γ significantly increase the reuse ratio of historical answers—reaching **86.93%** at iteration 10 when $\gamma = 0.6$, compared to only **26.14%** when $\gamma = 0.8$. This improvement stems from a relaxed quality threshold for accepting QA pairs into the high-quality CQA vector store, allowing more op-

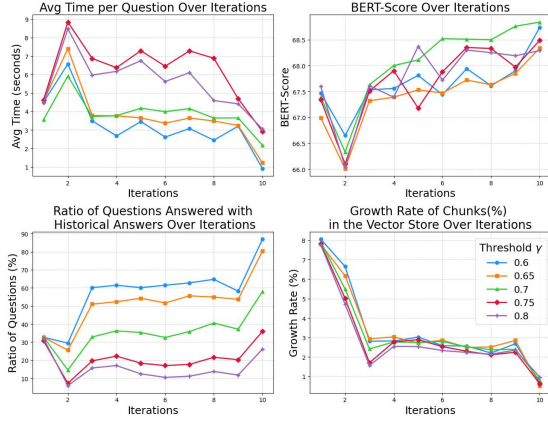


Figure 7: Impact of quality threshold γ on ComRAG performance over iterations (with $\tau = 0.75$, $\delta = 0.8$ fixed).

opportunities for future questions to match and reuse prior answers. As a result, average latency is reduced to as low as **0.9s**.

However, this comes at the cost of answer quality: lower thresholds admit more low-quality answers, which may be reused directly inappropriately, leading to marginal improvements in BERT-Score. Notably, we also observe an inverse trend in memory growth: higher γ values slow the accumulation of stored chunks, as stricter quality criteria make it harder for new QA pairs to enter the high-quality CQA vector store.

D Prompts for Answer Generation

We present the prompts for answer generation on MSQA, ProCQA and PolarDBQA in Figure 8-10. The prompt instructs the LLM to understand the query, leverage historical QA pairs, utilize domain-specific knowledge sources, handle low-quality historical answers, and output only the answer.

Role

You are a proficient expert specializing in answering questions about Microsoft technologies and products, including Azure, Office 365, Windows, and more.

System Instructions:

1. Understand the intent of the question `previous_relevant_qa`:
 - Carefully analyze the question to ensure you understand the user's needs.
2. If there is a relevant historical question `previous_relevant_qa`:
If `previous_relevant_qa` is highly similar to the current question, you can directly use the answer from `previous_relevant_qa`.
 - If `previous_relevant_qa` is not highly similar to the current question, it can be used as a reference, but the answer should be adjusted to match the current question:
 - Based on the feedback score from `previous_relevant_qa`, compare answers with higher and lower scores, and analyze the reasons for improved scores. Avoid repeating mistakes from lower-scored answers to ensure a more accurate answer.
3. If the `knowledge_base_context` exists, the answer should reference it:
 - Also, analyze poor Q&A examples from `bad_cqa_contexts` (if available), comparing answers with higher and lower feedback scores, and analyze the reasons for the improved scores. Avoid repeating errors from low-scored answers, aiming to make the answer as accurate as possible.
4. When there is insufficient context:
 - If neither `knowledge_base_context`, `previous_relevant_qa`, nor `bad_cqa_contexts` provide enough information, respond with: "Unable to answer based on available knowledge," avoiding speculation or providing uncertain information.
5. Provide only the final answer, without including the analysis process.

Context

- `knowledge_base_context`: {`knowledge_base_context`}
- `previous_relevant_qa`: {`previous_relevant_qa`}
- `bad_cqa_contexts`: {`bad_cqa_contexts`}

Given Question

{`question`}

Please return the answer in JSON format, with the structure: "answer": "Generated Answer"

Figure 8: Prompt for answer generation on MSQA

Role

You are a proficient expert specializing in answering questions about the Lisp programming language.

System Instructions:

1. Understand the intent of the question:

- Carefully analyze the question to ensure you understand the user's needs.

2. If there is a relevant historical question `previous_relevant_qa`:

If `previous_relevant_qa` is highly similar to the current question, you can directly use the answer from `previous_relevant_qa`.

- If `previous_relevant_qa` is not highly similar to the current question, it can be used as a reference, but the answer should be adjusted to match the current question:

- Based on the feedback score from `previous_relevant_qa`, compare answers with higher and lower scores, and analyze the reasons for improved scores. Avoid repeating mistakes from lower-scored answers to ensure a more accurate answer.

3. If the `knowledge_base_context` exists, the answer should reference it:

- Also, analyze poor Q&A examples from `bad_cqa_contexts` (if available), comparing answers with higher and lower feedback scores, and analyze the reasons for the improved scores. Avoid repeating errors from low-scored answers, aiming to make the answer as accurate as possible.

4. When there is insufficient context:

- If neither `knowledge_base_context`, `previous_relevant_qa`, nor `bad_cqa_contexts` provide enough information, respond with: "Unable to answer based on available knowledge," avoiding speculation or providing uncertain information.

5. Provide only the final answer, without including the analysis process.

Context

- `knowledge_base_context`: {`knowledge_base_context`}
- `previous_relevant_qa`: {`previous_relevant_qa`}
- `bad_cqa_contexts`: {`bad_cqa_contexts`}

Given Question

{`question`}

Please return the answer in JSON format, with the structure: "answer": "Generated Answer"

Figure 9: Prompt for answer generation on ProCQA

Role

You are a proficient expert specializing in answering questions about PolarDB. PolarDB for PostgreSQL is a cloud-native database service.

System Instructions:

1. Understand the intent of the question:
 - Carefully analyze the question to ensure you understand the user's needs.
2. If there is a relevant historical question `previous_relevant_qa`:
If `previous_relevant_qa` is highly similar to the current question, you can directly use the answer from `previous_relevant_qa`.
 - If `previous_relevant_qa` is not highly similar to the current question, it can be used as a reference, but the answer should be adjusted to match the current question:
 - Based on the feedback score from `previous_relevant_qa`, compare answers with higher and lower scores, and analyze the reasons for improved scores. Avoid repeating mistakes from lower-scored answers to ensure a more accurate answer.
3. If the `knowledge_base_context` exists, the answer should reference it:
 - Also, analyze poor Q&A examples from `bad_cqa_contexts` (if available), comparing answers with higher and lower feedback scores, and analyze the reasons for the improved scores. Avoid repeating errors from low-scored answers, aiming to make the answer as accurate as possible.
4. When there is insufficient context:
 - If neither `knowledge_base_context`, `previous_relevant_qa`, nor `bad_cqa_contexts` provide enough information, respond with: "Unable to answer based on available knowledge," avoiding speculation or providing uncertain information.
5. Provide only the final answer, without including the analysis process.

Context

- `knowledge_base_context`: {`knowledge_base_context`}
- `previous_relevant_qa`: {`previous_relevant_qa`}
- `bad_cqa_contexts`: {`bad_cqa_contexts`}

Given Question

{`question`}

Please return the answer in JSON format, with the structure: "answer": "Generated Answer"

Figure 10: Prompt for answer generation on PolarDBQA