RAG-EF: Train-free RAG Enhancement from Expert Feedbacks

Anonymous ACL submission

Abstract

As large language models (LLMs) become increasingly integral to digital interactions, their susceptibility to generating inaccurate or nonsensical content, called *hallucination*, poses significant challenges. Retrieval-Augmented Generation (RAG) has emerged as a promising technique to curb hallucinations by leveraging external databases to inform response generation. However, the RAG framework is not without limitations, often requiring computationally expensive methods like domain-specific retrieval-augmented finetuning. We introduce a novel and efficient enhancement RAG framework, RAG-EF (RAG with Expert Feedback), which incorporates expert-provided feedback composed of problematic Q&A and context pairs. Also, we present a new retrieval strategy that utilizes contexts alongside Q&A pairs to optimize information selection and prevent incorrect responses. To show the effectiveness of RAG-EF, we establish three new benchmarks with three datasets, and demonstrate adding relevant feedback into the database greatly improves the performance.

1 Introduction

The growth of data on the internet has spurred the advancement of large language models (LLMs) (Dubey et al., 2024; Achiam et al., 2023; Team et al., 2024). These models are transforming our interaction with digital content, proving invaluable in various domains, such as personalization (Bang et al., 2024), text summarization (Basyal and Sanghvi, 2023), etc. Despite their advanced capabilities, LLMs are prone to producing content that are inaccurate or nonsensical, an issue referred to as *hallucination* (Ji et al., 2023).

To alleviate this problem, Retrieval-Augmented Generation (RAG) (Lewis et al., 2020) has been introduced, which is a technique that improves LLMs by incorporating information from external databases (*e.g.,* Wikipedia) during the inference.

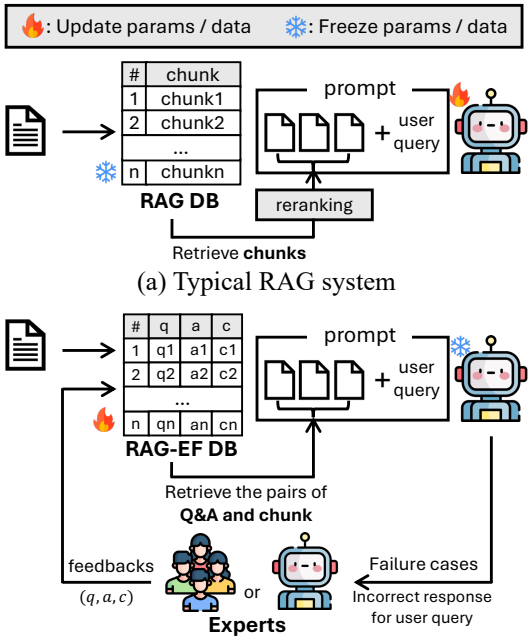

(a) Typical RAG system

(b) RAG system leveraging expert's feedbacks

Figure 1: **Overall system of RAG-EF.** For enhancing the model for failure cases, typical RAG system trains the LLM with training set augmented with feedbacks for failure cases, yet it is inefficient. RAG-EF is the new RAG system for enhancing the feedback (q, a, c) from the experts without model training, and it is efficient.

This approach has been effective in reducing hallucinations by providing a factual basis for response generation. However, it is inevitable for a retriever to select the irrelevant (*i.e.,* noisy) context to the user query. Additionally, the prior work (Liu et al., 2024) found that ordering retrieved contexts can severely affect the RAG performance (*i.e., lost-in-the-middle*). As such, many works (Yan et al., 2024; Yu et al., 2024b; Dong et al., 2024) proposed the methods for reranking retrieved contexts using LLM, but it still has difficulties for specific domain that the LLM struggles with.

To adapt the model for the specific domain, fine-tuning methods have been revealed. RAFT (Zhang et al., 2024) was proposed to make the robust gen-

erator in irrelevant contexts for a specific domain. Additionally, the paper (Siriwardhana et al., 2023) was introduced to adapt both a retriever and a generator into the domain-specific document. Although they improve the performance for the specific domain, those kind of methods require to train the model whenever updating the database, and it is impractical framework for maintaining the model.

This paper aims to address the efficient RAG enhancement framework without additional extensive training when the database is updated. Inspired by the paper that argues LLM cannot correct reasoning itself and need external feedback (Huang et al., 2024), we firstly propose the novel RAG framework based on feedback from the experts, as illustrated in Fig. 1. More specifically, when the LLM generates the incorrect answer for the user query, the experts (*e.g.,* human or LLM) generate the feedback to correct the answer for the user query, and this feedback is inserted into the database. In this context, we also introduce a new retrieval strategy that utilizes both the contexts and the Q&A pairs (*i.e.,* feedback) in the database to refine the RAG process. The contributions are three folds:

- We firstly propose RAG-EF, a novel efficient RAG enhancement framework by updating the database (*i.e.,* feedback) without training.
- We develop a unique retrieval method considering both context and Q&A that are stored pairly in the database for effective sampling information in our framework.
- We show that RAG system is enhanced by incorporating Q&As from the LLM-generated database, even without expert's feedback.

2 Methodology

2.1 Overview of Existing RAG

Typically, RAG utilizes two models; Retriever($\mathcal{R}$) and Generator($\mathcal{G}$). At first, given a document d, we need to split it into several chunks($C = \{c_1, c_2, ..., c_n\}$) due to max token constraint of LLM. Then, retriever $\mathcal{R}$ generates the embeddings for each chunk and the user query, and selects the top-K chunks from C that are most relevant chunks to the given user query q by calculating the similarity. The similarity $\mathbf{s}_i$ between each chunk c_i and user query q as below:

$$\mathbf{s}_i = e_{\mathrm{q}}^{\top} e_{c_i}, \;\; e_{\mathrm{q}} = \mathcal{R}(\mathrm{q}), \;\; e_{c_i} = \mathcal{R}(c_i) \quad (1)$$

Then, both a query and selected chunks are fed into the prompt as the input of the generator $\mathcal{G}$.

2.2 Configuration of RAG-EF Database

Different from conventional RAG, RAG-EF database (DB) consists of both Q&A and chunk that contains the answer the question. There are two approaches to collect the Q&A-Chunk pairs.

Feedback from Expert. In some cases, RAG may generate the wrong responses by several reasons. For failure cases, the experts can annotate the answer a and golden chunk c for the user query q, where c is part of the document d but is not necessary to be same as one of the c_i. After constructing the (q, a, c) pair, it is inserted into the database.

Generation from LLM. When the chunk c_i is given, we can ask the LLM to generate the Q&A pairs to validate the understanding of c_i. By doing so, we can collect the set of Q&A and chunk c_i pairs (*i.e.,* $\{(q_j, a_j, c_i) | j \in [1, m_i]\}$, and add them into the database (see Appendix A for details). Generally, we utilize this approach when building the database at the initial stage.

2.3 Retrieval Method of RAG-EF

Different from typical RAG solutions, we need to consider both Q&A and chunk pair $p_i = (q_i, a_i, c_i)$ that is stored in the database. To do so, we propose a new retrieving method by reformulating the similarity from Eq. (1).

$$\mathbf{s}_i = (e_{\mathrm{q}}^{\top} e_{q_i})^{\gamma} (e_{\mathrm{q}}^{\top} e_{c_i})^{1-\gamma} \quad (2)$$

where $e_{\mathrm{q}} = \mathcal{R}(\mathrm{q})$, $e_{q_i} = \mathcal{R}(q_i)$, $e_{c_i} = \mathcal{R}(c_i)$ and hyperparameter $\gamma \in [0, 1]$. Given similarities $\mathbf{s}_i$ for all the pairs p_i from Eq. (2), we select top-K pairs.

After sampling the Q&A-Chunk pairs p_i, both a user query q and selected pairs p_i are added into the input of the generator $\mathcal{G}$. Since the chunks can be duplicated, the number of chunks can be less than the number of pairs (See Appendix B for details).

3 Experiment

3.1 RAG-EF Database Benchmark

Database. This work is the first to suggest the use of feedback (*i.e.,* Q&A-Chunk pairs p_i) to enhance the RAG system. Since there is no standard for evaluating the efficacy of RAG-EF, we introduce three benchmarks to assess its performance. Based on Sec. 2.2, these benchmarks include one for generation using LLM ($\mathcal{S}$) and two that incorporate feedback from the expert ($\mathcal{E}$ and $\mathcal{O}$), to demonstrate

$\mathcal{R}$	Database	Natural Question EM	Natural Question F1	TriviaQA EM	TriviaQA F1	HotpotQA EM	HotpotQA F1	Average EM	Average F1
BM25	$\mathcal{S}$	26.01	39.78	75.56	81.79	21.83	30.10	41.13	50.56
	$\mathcal{S} \cup \mathcal{E}$	26.95	40.68	77.81	85.17	22.52	30.77	42.43	52.21
	$\mathcal{S} \cup \mathcal{O}$	27.11	40.77	90.31	93.54	23.23	31.35	46.88	55.22
DPR	$\mathcal{S}$	36.93	54.95	74.72	80.88	32.14	43.82	47.93	59.88
	$\mathcal{S} \cup \mathcal{E}$	49.96	68.52	81.60	88.61	40.04	52.67	57.20	69.93
	$\mathcal{S} \cup \mathcal{O}$	78.12	87.38	93.26	95.97	81.27	86.38	84.22	89.91
bge-m3	$\mathcal{S}$	36.19	55.19	76.83	83.02	32.78	44.76	48.60	60.99
	$\mathcal{S} \cup \mathcal{E}$	49.80	68.92	83.85	91.35	40.49	53.17	58.05	71.15
	$\mathcal{S} \cup \mathcal{O}$	76.60	86.25	96.49	98.45	79.24	84.45	84.11	89.72

Table 1: **Main results of RAG-EF.** We first build the database with $\mathcal{S}$, generated by LLM with contexts. After then, we add feedback sets (*i.e.,* $\mathcal{E}$ or $\mathcal{O}$), which are relevant to the user query q, into the database to enhance the RAG performance. Note that the difference between $\mathcal{E}$ and $\mathcal{O}$ is that $\mathcal{O}$ is an exact match with a test set, whereas $\mathcal{E}$ is not.

Methods	Database Config Context	Database Config Q&A	Natural Question EM	Natural Question F1	TriviaQA EM	TriviaQA F1	HotpotQA EM	HotpotQA F1
w/o RAG (No RAG)	✗	✗	16.26	28.68	69.10	75.44	16.75	28.46
Typical RAG	✓	✗	28.74	46.62	67.13	75.09	26.40	39.09
RAG-EF w/o ctx	✗	✓	26.25	41.56	58.15	66.24	21.40	31.79
RAG-EF (Ours)	✓	✓	**36.93**	**54.95**	**74.72**	**80.88**	**32.14**	**43.82**

Table 2: **Fair comparison between typical RAG and RAG-EF with only $\mathcal{S}$ generated by LLM without a test set.** It shows the large gap between typical RAG and RAG-EF.

the performance improvement when feedback is included in the database (See Appendix C).

Synthesis($\mathcal{S}$) - It is created with the generative method described in Sec. 2.2, utilizing the LLaMA-3 8B with each chunk. As it is created without a test set, it is hard to collect the relevant Q&A-Chunk pairs to the user query q within a test set compared to the other benchmarks using the test set.

Expert($\mathcal{E}$) - Adhering to the expert feedback method detailed in Sec. 2.2, it is generated by the expert (employing a GPT) with a test set. Initially, the LLM rephrases the user query q, following which it generates answers for the rephrased query using a golden chunk. Consequently, each Q&A-Chunk pair p_i is composed of the rephrased query, the LLM's response, and a golden chunk.

Omniscient($\mathcal{O}$) - It also constitutes of the feedbacks from the expert, but it is totally same as a test set. Therefore, it is the most relevant to the test set among three benchmarks, but impractical.

Implementation Details. To show the effectiveness of RAG-EF for various retrievers, we adopt three retrievers $\mathcal{R}$ that are BM25 (Robertson et al., 2009), DPR (Karpukhin et al., 2020), and bge-m3 (Chen et al., 2024). All the retrievers select five pairs in the database using Eq. (2) with $\gamma = 0.5$. As a generator $\mathcal{G}$, we utilized LLaMA-3 8B (Dubey et al., 2024) (See Appendix D for more details).

3.2 Main Results

Effectiveness of RAG-EF. Our hypothesis is that RAG-EF will boost performance notably when the database includes Q&A-Chunk pairs p_i that are relevant to the user query q. To evaluate RAG-EF's efficacy, we start with a database initialized with $\mathcal{S}$. Subsequently, we insert feedback sets $\mathcal{E}$ or $\mathcal{O}$ into the database. It is important to note that the difference between $\mathcal{E}$ and $\mathcal{O}$ is that $\mathcal{O}$ is an exact match with the test set, whereas $\mathcal{E}$ is not. In Tab. 1, the $\mathcal{S} \cup \mathcal{O}$ achieves the highest performance by a significant margin, due to the database containing pairs that exactly match the test set. However, having a database $\mathcal{O}$, which perfectly mirrors the test set, is an unrealistic scenario. In contrast, the $\mathcal{S} \cup \mathcal{E}$, which is both relevant and realistic, shows a substantial improvement over the baseline set $\mathcal{S}$. From these findings, we can conclude that RAG-EF significantly improves performance when the database includes Q&A-Chunk pairs p_i that are relevant to the user query q.

Various Retrievers. We also conducted experiments with various retrievers. Tab. 1 shows the performance trends across three different retrievers are consistent. Notably, BM25 exhibited a narrower performance gap compared to the other retrievers. This can be attributed to BM25's susceptibility to paraphrased text, stemming from its reliance on

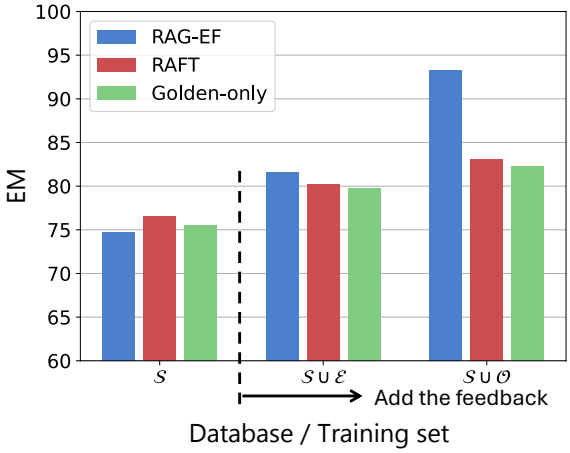

Figure 2: **Performance comparison between RAG-EF(ours) and train-based RAG in TriviaQA.** Despite of necessities of extra computational resources, train-based RAG (See Appendix E) fails to meet RAG-EF after adding feedback sets into the database.

a traditional retrieval scoring mechanism, which contrasts with the deep learning-based approaches employed by the other retrieval methods.

3.3 Further Analyses

Comparison of Typical RAG. Tab. 2 shows a comparison between RAG-EF and the typical RAG, highlighting the advantages of integrating retrieved contexts with Q&A pairs. To ensure a fair comparison with the typical RAG, we only employed the sequences generated by LLaMA-3 8B without incorporating the test set. We find that employing the database (rows 2–4) yields superior performance compared to the absence of RAG (row 1). Furthermore, the typical RAG demonstrates enhanced effectiveness over the use of Q&A pairs alone (row 3). The database with $\mathcal{S}$, generated by LLM without a test set, encounters challenges in existing Q&A-Chunk pairs relevant to the test set, leading to marginally inferior outcomes when compared to the typical RAG. Nonetheless, the combined use of contexts and Q&A pairs from the database significantly enhances performance relative to the other methods. Despite the potential disconnect between the question in the retrieved Q&A-Chunk pairs p_i and the user query q, the associated chunk may still possess high relevance to the user query q. This relevance can potentially empower the generator $\mathcal{G}$ to produce more accurate response.

Comparison of Train-based RAG. Fig. 2 illustrates a comparison between RAG-EF and other train-based RAG methods, such as RAFT and Golden-Only. We analyze that better alignment between training datasets (or database) and a test set (*i.e.,* more relevant pairs p_i in the database), all the methods boost their performance. However, RAG-EF benefits more significantly from this alignment, increasing its lead over the others. This indicates that RAG-EF is particularly effective when the database contains pairs p_i closely aligned with user queries q, making it a preferable choice for enhancing the RAG system with updated databases.

4 Related Works

Train-based RAG. Several papers (Lin et al., 2024; Wang et al., 2024) have been proposed in terms of improving RAG performance with training. They focused on instruction tuning pretrained LLM to be familiar with RAG tasks, preparing the training set with several domains for generability. On the contrary, RAFT (Zhang et al., 2024) focused on finetuning the generator(*i.e.,* LLM) for adapting the specific domain, and RAG-end2end (Siriwardhana et al., 2023) introduced the joint training method of the retriever and the generator. Recently, Chain-of-Note (Yu et al., 2024a) proposed the reasoning method by training the generator to generate both summaries for each context and an answer inspired by Chain-of-Thought (Wei et al., 2022), and showed the performance improvement.

Train-free RAG. While RAG mostly solves the hallucination of pretrained LLM, RAG has the severe problem called *lost-in-the-middle* (Liu et al., 2024), which LLM generates the wrong answers with high probability when the golden context is located in the middle of the prompt. To solve this problem, *Reranking* concept was proposed. CRAG (Yan et al., 2024) proposed a *retrieval evaluator*, which decides whether to put contexts selected from the retriever into the LLM's prompt or not. Rank-RAG (Yu et al., 2024b) proposed unified LLM that can conduct ranking the contexts and generating the answer, and G-RAG (Dong et al., 2024) proposed graph-based reranking system. Similarly, Self-RAG (Asai et al., 2024) proposed *critique* to determine whether each context is dropped or not.

5 Conclusions

We introduced RAG-EF, a novel RAG enhancement system by integrating feedback mechanisms with an updated database. Our unique retrieval method utilizes context and Q&A to effectively select pairs in the database. The efficacy of RAG-EF was validated, making it the optimal RAG enhancement system for handling updated databases.

6 Limitations

This work acknowledges that the input length of the generator $\mathcal{G}$ increases compared to conventional RAG due to adding the pairs in RAG-EF database that contains both Q&As and contexts. However, most of the Q&A pairs are less than 100 tokens, which is largely less than the length of retrieved contexts. Moreover, RAG-EF is vulnerable in storage constraint (*e.g.,* edge devices) since database size increases by adding the feedbacks continuously. However, large-sized database might have the relevant Q&A-Chunk pairs (*i.e.,* feedback) to the user query q with high probability, resulting in high performance. Additionally, we can concern the management of database (*e.g.,* cache management), which can be the future research.

7 Potential Risks

Retrieval-Augmented Generation (RAG) models pose potential privacy risks when the document incorporates the sensitive information such as personal information. RAG-EF focuses on enhancement of RAG system for failure cases, and it is not directly related to the negative ethical and societal impacts. Nonetheless, when RAG-EF is applied to the application service, it needs to consider carefully whether external database (*i.e.,* document) contains the private information or not.

References

Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *ArXiv:2303.08774*.

Akari Asai, Zeqiu Wu, Yizhong Wang, Avirup Sil, and Hannaneh Hajishirzi. 2024. Self-rag: Learning to retrieve, generate, and critique through self-reflection. In *Proceedings on International Conference on Learning Representations*.

Jihwan Bang, Juntae Lee, Kyuhong Shim, Seunghan Yang, and Simyung Chang. 2024. Crayon: Customized on-device llm via instant adapter blending and edge-server hybrid inference. In *Proceedings on Association for Computational Linguistics*, pages 3720–3731.

Lochan Basyal and Mihir Sanghvi. 2023. Text summarization using large language models: a comparative study of mpt-7b-instruct, falcon-7b-instruct, and openai chat-gpt models. *ArXiv:2310.10449*.

Jianlv Chen, Shitao Xiao, Peitian Zhang, Kun Luo, Defu Lian, and Zheng Liu. 2024. Bge m3-embedding: Multi-lingual, multi-functionality, multi-granularity text embeddings through self-knowledge distillation. *ArXiv:2402.03216*.

Jialin Dong, Bahare Fatemi, Bryan Perozzi, Lin F Yang, and Anton Tsitsulin. 2024. Don't forget to connect! improving rag with graph-based reranking. *ArXiv:2405.18414*.

Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *ArXiv:2407.21783*.

Edward J Hu, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. 2022. Lora: Low-rank adaptation of large language models. In *Proceedings on International Conference on Learning Representations*.

Jie Huang, Xinyun Chen, Swaroop Mishra, Huaixiu Steven Zheng, Adams Wei Yu, Xinying Song, and Denny Zhou. 2024. Large language models cannot self-correct reasoning yet. In *Proceedings on International Conference on Learning Representations*.

Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. 2023. Survey of hallucination in natural language generation. *ACM Computing Surveys*, 55(12):1–38.

Mandar Joshi, Eunsol Choi, Daniel S Weld, and Luke Zettlemoyer. 2017. Triviaqa: A large scale distantly supervised challenge dataset for reading comprehension. In *Proceedings on Association for Computational Linguistics*, pages 1601–1611.

Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. Dense passage retrieval for open-domain question answering. In *Proceedings on Empirical Methods in Natural Language Processing*, pages 6769–6781.

Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, et al. 2019. Natural questions: a benchmark for question answering research. *Transactions of the Association for Computational Linguistics*, 7:453–466.

Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. In *Proceedings on Advances in Neural Information Processing Systems*, volume 33, pages 9459–9474.

Xi Victoria Lin, Xilun Chen, Mingda Chen, Weijia Shi, Maria Lomeli, Richard James, Pedro Rodriguez, Jacob Kahn, Gergely Szilvasy, Mike Lewis, et al.

2024. Ra-dit: Retrieval-augmented dual instruction tuning. In *Proceedings on International Conference on Learning Representations*.

Nelson F Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2024. Lost in the middle: How language models use long contexts. *Transactions of the Association for Computational Linguistics*, 11:157–173.

Stephen Robertson, Hugo Zaragoza, et al. 2009. The probabilistic relevance framework: Bm25 and beyond. *Foundations and Trends® in Information Retrieval*, 3(4):333–389.

Shamane Siriwardhana, Rivindu Weerasekera, Elliott Wen, Tharindu Kaluarachchi, Rajib Rana, and Suranga Nanayakkara. 2023. Improving the domain adaptation of retrieval augmented generation (rag) models for open domain question answering. *Transactions of the Association for Computational Linguistics*, 11:1–17.

Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, et al. 2024. Gemma 2: Improving open language models at a practical size. *ArXiv:2408.00118*.

Boxin Wang, Wei Ping, Lawrence McAfee, Peng Xu, Bo Li, Mohammad Shoeybi, and Bryan Catanzaro. 2024. Instructretro: Instruction tuning post retrieval-augmented pretraining. In *Proceedings on International Conference on Machine Learning*.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. In *Proceedings on Advances in Neural Information Processing Systems*, volume 35, pages 24824–24837.

Shi-Qi Yan, Jia-Chen Gu, Yun Zhu, and Zhen-Hua Ling. 2024. Corrective retrieval augmented generation. *ArXiv:2401.15884*.

Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D Manning. 2018. Hotpotqa: A dataset for diverse, explainable multi-hop question answering. In *Proceedings on Empirical Methods in Natural Language Processing*, pages 2369–2380.

Wenhao Yu, Hongming Zhang, Xiaoman Pan, Kaixin Ma, Hongwei Wang, and Dong Yu. 2024a. Chain-of-note: Enhancing robustness in retrieval-augmented language models. In *Proceedings on Empirical Methods in Natural Language Processing*.

Yue Yu, Wei Ping, Zihan Liu, Boxin Wang, Jiaxuan You, Chao Zhang, Mohammad Shoeybi, and Bryan Catanzaro. 2024b. Rankrag: Unifying context ranking with retrieval-augmented generation in llms. *ArXiv:2407.02485*.

Tianjun Zhang, Shishir G Patil, Naman Jain, Sheng Shen, Matei Zaharia, Ion Stoica, and Joseph E Gonzalez. 2024. Raft: Adapting language model to domain specific rag. *ArXiv:2403.10131*.

-Supplementary Material-

RAG-EF: Train-free RAG Enhancement from Expert Feedbacks

A Q&A Generation from LLM

To build the synthesis dataset $\mathcal{S}$, we ask LLM (*i.e.,* Llama-3 8B) for each chunk from the document by utilizing following prompt template. After getting responses from Llama-3 8B, we parse them into the question-and-answer pairs. In other words, we can get the Q&A-Chunk pairs p_i for synthesis set $\mathcal{S}$ by conducting those process.

Prompt template for Q&A Generation from LLM

```
please generate question and answer pairs to check
the understandability of following context.
Format should be Q: question A: answer.
{CONTEXT}
```

B Prompt Template for RAG

Prompt engineering with retrieved chunks or Q&As is the another important factor to improve the performance. For fair comparison between RAG-EF and conventional RAG, we tried to unify the prompt template as shown below. Especially, RAG-EF retrieves not only chunks but also question and answer pairs, they added Q&As into the prompts before the contexts. Since it can be possible to duplicate the chunks, the number of retrieved contexts, denoted as N, can be five or fewer. Aside from this difference, the prompt templates employed by RAG-EF and the conventional RAG remain largely consistent.

Prompt template for RAG-EF

```
Question: {QUESTION #1}
Answer: {ANSWER #1}
Question: {QUESTION #2}
Answer: {ANSWER #2}
...
Question: {QUESTION #5}
Answer: {ANSWER #5}
Context1: {CONTEXT #1}
Context2: {CONTEXT #2}
...
ContextN: {CONTEXT #N}
Please answer the below question based on given
above question and answer pairs and contexts.
Note that you should generate the response only
for answering the question within a few words.
Do not contain extra comments.
Question: {TARGET QUESTION}
```

Prompt template for conventional RAG

```
Context1: {CONTEXT #1}
Context2: {CONTEXT #2}
...
Context5: {CONTEXT #5}
Please answer the below question based on given
above contexts. Note that you should generate
the response only for answering the question
within a few words. Do not contain extra comments.
Question: {TARGET QUESTION}
```

C RAG-EF Database Benchmark

C.1 Examples in Three Types of Database

Tab. 3 shows the examples of three database benchmarks $\mathcal{S}$, $\mathcal{E}$, and $\mathcal{O}$ that are most relevant examples to the test set example using the DPR retriever. As described in Sec. 3.1, since $\mathcal{S}$ is built by LLaMA-3 8B with the chunk c_i, it is hard to generate the Q&A that are exactly same as the test set example. Nonetheless, the example from $\mathcal{S}$ shows the correlation of test set example, and it is sufficient to generate the correct answer for the question of the test set by referring the Q&A of $\mathcal{S}$. In $\mathcal{E}$, the only difference from the test set example is the question, and you can observe that they have the same meaning. Sometimes, it can get the different (*i.e.,* incorrect) answers due to answer generation from LLM (See the LLM's performance in below section). Lastly, the example of $\mathcal{O}$ is same as the test set example.

C.2 Performance of LLM when Building $\mathcal{E}$

To build the dataset $\mathcal{E}$ that is relevant and realistic, we make GPT paraphrase the question and generate the answer given rewritten question and whole document to regard the GPT as a human. As such, Tab. 4 shows the results of GPT for various datasets. Due to constrained performance of GPT, Tab. 1 shows that there is a gap between $\mathcal{S} \cup \mathcal{E}$ and $\mathcal{S} \cup \mathcal{O}$. We believe that the results of $\mathcal{S} \cup \mathcal{E}$ can significantly be improved when the performance of GPT increases (The humans can have more accurate results than GPT).

D Implementation Details

D.1 Datasets

TriviaQA (Joshi et al., 2017). It offers a complex question answering dataset featuring 950,000

Set	Question	Answer	Chunk
test	Who was known by his stage name Aladin and helped organizations improve their performance as a consultant?	Eenasul Fateh	Eenasul Fateh (born 3 April 1959), also known by his stage name Aladin, is a Bangladeshi-British cultural practitioner, magician, live artist and former international management consultant.
$\mathcal{S}$	Who is Eenasul Fateh also known as?	Aladin	Eenasul Fateh (born 3 April 1959), also known by his stage name Aladin, is a Bangladeshi-British cultural practitioner, magician, live artist and former international management consultant.
$\mathcal{E}$	Who was recognized under the stage name Aladin and provided consultancy services to enhance organizational performance?	Eanasul Fateh	Eenasul Fateh (born 3 April 1959), also known by his stage name Aladin, is a Bangladeshi-British cultural practitioner, magician, live artist and former international management consultant.
$\mathcal{O}$	Who was known by his stage name Aladin and helped organizations improve their performance as a consultant?	Eenasul Fateh	Eenasul Fateh (born 3 April 1959), also known by his stage name Aladin, is a Bangladeshi-British cultural practitioner, magician, live artist and former international management consultant.

Table 3: **Examples of three database benchmarks that are the most relevant to ground truth in HotpotQA.**

Dataset	EM	F1
Natural Question	54.64	72.98
TriviaQA	72.05	82.57
HotpotQA	35.10	46.18

Table 4: **GPT performance in various datasets given whole document.** Since some examples do not have the key passages that contain the answers, GPT has especially poor performance in HotpotQA.

Parameters	Value
Rank	8
Alpha	16
Train on Inputs	True
LoRA Modules	Q, K, V, O
Dropout	0.05
Batch Size	128
Loss Function	Cross Entropy Loss
Learning Rate	3e-4
Scheduler	Cosine Annealing
Optimizer	AdamW
Epochs	1

Table 5: **Training hyperparameters.**

pairs of questions and answers sourced from over 662,000 documents on Wikipedia and the internet. Unlike simpler QA benchmarks like SQuAD, TriviaQA presents a tougher challenge because answers often require more than just predicting a text span from lengthy contexts. The dataset includes both subsets that are verified by humans and those created by machines, and we only utilized verified datasets in this experiment.

Natural Questions (Kwiatkowski et al., 2019). It is designed for training question answering systems, featuring over 300,000 training, nearly 8,000 development, and approximately 8,000 test examples, each pairing a Google search query with a relevant Wikipedia article. These articles include marked sections that provide a detailed response to the query, as well as shorter snippets that directly answer the question, although some annotations may be absent, indicating no answer is available. Additionally, a small fraction of the dataset, about 1%, contains binary "yes" or "no" answers instead of detailed excerpts.

HotpotQA (Yang et al., 2018). It is a dataset featuring roughly 113,000 questions based on English Wikipedia, designed to need the lead sections from two articles for answers. Accompanying each question are the key paragraphs and a selection of sentences marked by crowdworkers as essential facts for responding. However, there are several examples that are missing some key paragraphs, and it causes the performance drop in RAG.

D.2 Evaluation Metrics

EM. Exact Match is a strict metric that measures the percentage of predictions that match the ground truth exactly. It is often used in the context of machine comprehension and question answering tasks where the goal is to produce an exact answer. For a given dataset, the Exact Match score is calculated by dividing the number of predictions that are exactly the same as the true answers by the total

number of predictions made.

F1. The F1 score is a more nuanced metric that considers both precision and recall. It is the harmonic mean of precision and recall, providing a balance between the two.

E Training Configuration

To validate the superiority of our method compared to train-based RAG, RAFT and golden-only finetuning methods followed the parameters as described in Tab. 5. We adopted LoRA tuning (Hu et al., 2022) that is one of the parameter efficient finetuning methods. As a training set, we utilized $\mathcal{S}$, $\mathcal{E}$, and $\mathcal{O}$ of which sizes are 147,160, 725, and 725, respectively. Despite of imbalance problem of three datasets, both RAFT and golden-only methods get higher EM as using more relevant dataset as a training set, but they have lower performance than RAG-EF that is only updating the dataset into the database. Lastly, all the experiments for training are conducted on two Nvidia A5000s.