

DynamicRAG: Leveraging Outputs of Large Language Model as Feedback for Dynamic Reranking in Retrieval-Augmented Generation


Jiashuo Sun, Xianrui Zhong, Sizhe Zhou, Jiawei Han*

University of Illinois Urbana-Champaign
{jiashuo5, hanj}@illinois.edu

RAG Framework




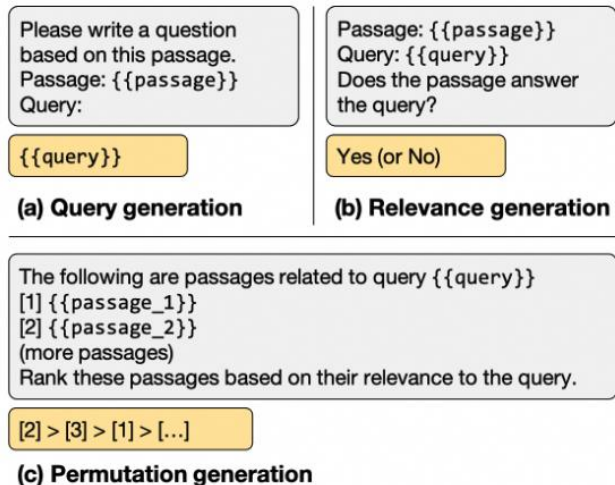
Reranker Component



Determining the **optimal value of k** is a challenging task

1. If k is too small, important information may be overlooked.
2. If k is set too large, irrelevant information may be introduced, long contexts can reduce the efficiency of the entire RAG system.

Recent studies have explored the use of LLMs as rerankers, leveraging their ability to generate rankings of documents. [1]




However, they mainly focus on leveraging the **internal knowledge of LLMs**, without fully utilizing the supervisory signals LLMs can provide.

Motivation

For instance, when a question and the ranked documents are input into the LLM, **the quality of the LLM's response** can directly reflect the relevance of the documents in the current ranking, offering a valuable feedback signal for optimization.


Data Construction



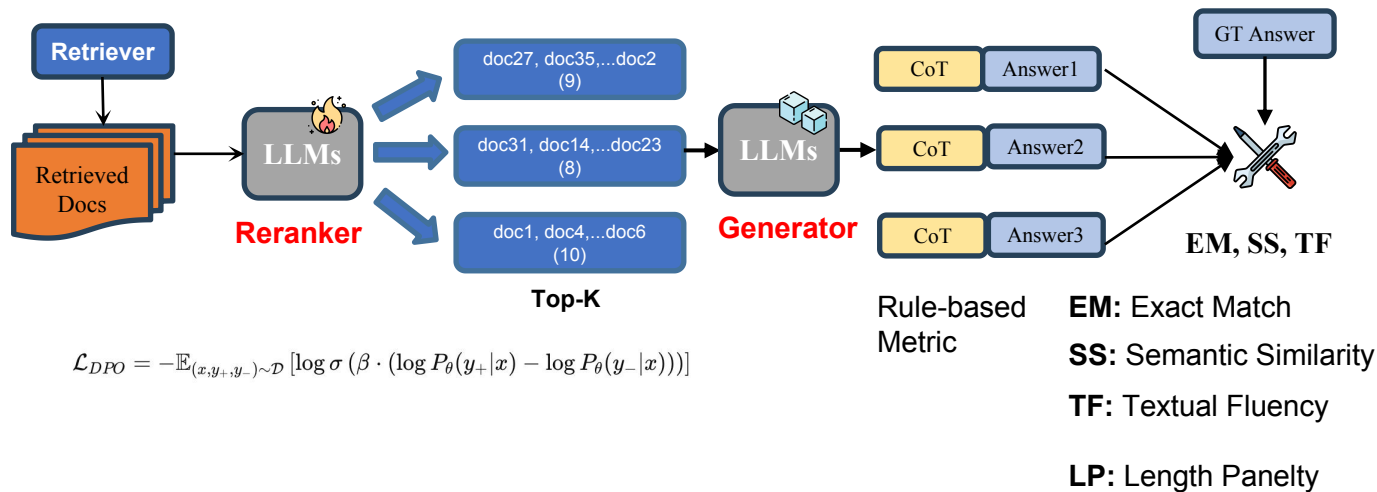
Corpus

Dynamic Reranker Agent

Top-K Docs



How to train LLM-ranker? Leveraging Large Language Model Outputs as Feedback




1. Behavior Cloning (SFT): LLMs to mimic the expert trajectory.

Taking Top10 as an example, all possible sequences are $P(10,1) + P(10,2) + \dots + P(10,10) = \mathbf{9,874,080}$,

$$P(10, k) = \frac{10!}{(10 - k)!}$$

2. Reinforcement Learning: We can use **DPO**, where we sampling different trajectories, and use the two with the **largest** and **smallest** rewards as positive and negative examples for training.

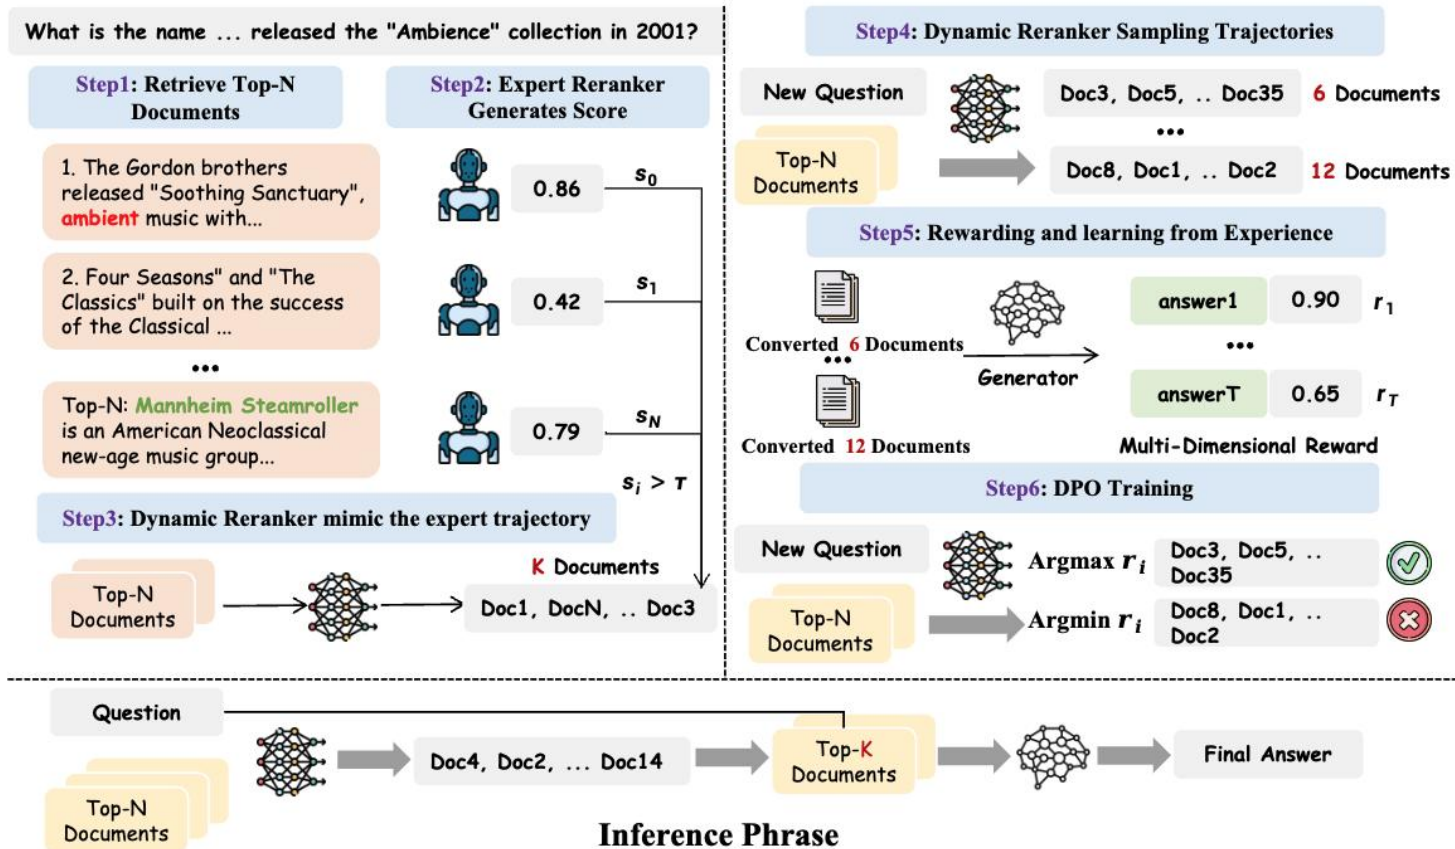
How to train LLM-generator?



Retrieval-augmented Training

$$\mathcal{L}_{CE} = -\frac{1}{N} \sum_{n=1}^N \log P_{\theta}(y_n | q_n, \{c_{n,1}, c_{n,2}, \dots, c_{n,k}\}),$$

Method



Experiment-Generator

Metrics	Extra Data for Training	NQ EM	TriviaQA EM/Acc	HotpotQA EM	2WikimQA EM	ASQA EM	FEVER Acc	ELI5 Rg
<i>Baseline Without Retrieval</i>								
GPT-3.5-Turbo [40]	N/A	38.6	70.7/74.3	29.9	23.9	68.3	82.7	27.5
GPT-4 [39]	N/A	40.3	73.3/78.4	34.5	29.8	71.9	87.7	30.3
GPT-4o [38]	N/A	40.0	74.0/79.2	36.1	33.3	74.1	86.3	30.2
<i>Baseline With Retrieval</i>								
IRCoT [55]	N/A	-	-	17.4	-	-	-	-
ReAct [58]	N/A	-	-	35.1	-	-	62.0	-
RA-DIT [29]	~ 1,129k	43.5	72.8/-	36.6	-	-	86.9	-
FLARE* [18]	N/A	-	-	-	51.0	41.3	-	-
Reward-RAG [36]	Unknown	42.2	75.6/80.4	-	-	-	89.8	-
LLaMA2-7B [54]	N/A	17.9	-/42.5	16.6	17.9	19.0	30.0	15.6
w/ Reranker	N/A	20.6	-/49.6	18.9	18.3	21.1	35.6	16.7
LLaMA2-7B-SFT	~ 130k	<u>29.1</u>	53.7/59.1	<u>27.1</u>	<u>18.9</u>	23.8	<u>40.6</u>	<u>18.6</u>
Self-RAG (LLaMA2-7B) [1]	~ 150k	36.4	<u>-/66.4</u>	-	-	<u>30.0</u>	-	-
DRAGIN [50]	N/A	-	-	23.2	22.0	-	-	-
Smart-RAG (LLaMA2-7B) [12]	~ 218k	-	-	-	-	26.6	-	-
Ours (LLaMA2-7B)	~ 150k	38.7	59.6/70.5	29.4	23.1	41.1	51.2	22.6
LLaMA2-13B [54]	N/A	23.6	-/47.0	17.7	18.7	20.5	30.2	19.9
w/ Reranker	N/A	26.5	-/53.2	20.4	18.8	23.6	37.1	<u>20.3</u>
LLaMA2-13B-SFT	~ 130k	<u>32.5</u>	60.1/66.2	<u>27.9</u>	<u>19.1</u>	28.4	<u>45.8</u>	20.1
Self-RAG (LLaMA2-13B) [1]	~ 150k	-	<u>-/69.3</u>	-	-	<u>31.7</u>	-	-
Ours (LLaMA2-13B)	~ 150k	39.1	62.3/72.6	30.1	25.0	46.4	77.2	23.3
LLaMA3-8B [13]	N/A	36.4	-/57.4	26.1	24.6	24.9	34.6	<u>24.0</u>
w/ Reranker	N/A	37.5	-/64.5	28.7	25.3	29.8	49.7	23.7
LLaMA3-8B-SFT	~ 130k	39.1	67.5/74.2	31.5	27.1	<u>46.8</u>	82.1	22.9
Auto-RAG (LLaMA3-8B-Instruct) [59]	Unknown	37.9	60.9/-	-	-	30.0	-	-
ChatQA-1.5 (LLaMA3-8B) [32]	~ 442k	42.4	81.0/87.6	33.4	26.8	-	90.9	-
RankRAG (LLaMA3-8B) [60]	~ 470k	50.6	82.9/89.5	<u>35.3</u>	<u>31.4</u>	-	93.8	-
Ours (LLaMA3-8B)	~ 150k	<u>48.4</u>	78.3/87.4	36.7	34.2	56.3	<u>91.4</u>	24.6

Experiment-Reranker

Table 2. The performance of different Reranker model. Results are directly from the original paper. Best results are in **bold** and the second results are underlined.

Metric	Training Data	NQ			HotpotQA			Avg.
		R@5	R@10	R@20	R@5	R@10	R@20	
Close-Sourced Models								
GPT-3.5-Turbo (Ouyang et al., 2022)	Unknown	77.8	82.5	85.7	52.1	56.6	62.4	69.5
GPT-4 (OpenAI et al., 2024b)	Unknown	79.3	83.2	85.1	53.2	57.0	61.0	69.8
Open-Sourced Rerank Models								
BM25 (Robertson & Zaragoza, 2009)	N/A	38.0	50.7	60.1	57.5	63.0	67.5	56.1
Contriever (Izacard et al., 2022)	Unknown	73.6	80.2	84.8	53.1	58.7	62.4	68.8
monoT5 (Nogueira et al., 2020)	~ 503k	75.6	80.9	84.9	54.8	60.2	63.3	70.0
RankLLaMA (Ma et al., 2024)	~ 503k	77.8	83.1	86.0	57.1	62.1	64.8	71.8
ChatQA-1.5 (LLaMA3-8B) (Liu et al., 2024b)	N/A	68.2	75.7	82.0	37.4	45.0	53.6	60.3
RankRAG (LLaMA3-8B) (Yu et al., 2024b)	~ 50k	80.3	84.0	86.3	57.6	61.8	65.2	72.5
Open-Sourced Generative Models								
GENRE (Cao et al., 2021)	~ 406k	61.4	-	-	34.0	-	-	-
Re3eval (Song et al., 2024)	~ 240k	65.4	-	-	44.2	-	-	-
SEAL (Bevilacqua et al., 2022)	Unknown	68.2	-	-	51.0	-	-	-
DynamicRAG (LLaMA3-8B)	~ 20k	79.3	83.7	86.8	59.1	63.7	67.2	73.7

Ablation Study

Effect of Key Components on DynamicRAG Performance

Table 4. The impact of different key components in DynamicRAG among different benchmarks. We use Exact Match as the metric.

	NQ	HotpotQA	ASQA	Avg.
	EM	EM	EM	
DynamicRAG	48.4	36.7	56.3	47.1
w/o Retrieval	25.0	25.6	15.7	22.1
w/o Reranking	36.4	27.2	39.8	34.5
w/o RL	44.6	29.6	45.5	39.9

Impact of Reranker and Generator Size

Reranker	DynamicRAG Generator	NQ EM	HotpotQA EM	Avg.
LLaMA2-7B	LLaMA2-13B	37.6	28.6	33.1
LLaMA2-13B	LLaMA2-13B	38.7	29.8	34.3
LLaMA2-13B*	LLaMA2-13B*	39.1	30.1	34.6
LLaMA3-8B	LLaMA2-13B	39.4	30.2	34.8

Further Analysis

We have fewer LLM calls and higher accuracy.

4.6.2. EFFICIENCY




Figure 5. Comparison of different RAG models in terms of efficiency and effectiveness. Models closer to the top-left corner achieve better overall performance.

Case Study1

Table 7. Case Study for DynamicRAG. We compare our method with Vanilla-RAG with Reranker. The retrieved content is same for both method. We take Top-40 documents as the input.

Question	What is the name of the American Neoclassical new-age music group that released the "Ambience" collection in 2001?
GPT-4o	The Ambient. ✗
Vanilla-RAG w/Reranker	<p>(Reranker Results:)</p> <p>Doc 2: Pure Moods Pure Moods was the first United States release of a series of compilation albums of new-age music released by Virgin Records. The original was titled "Moods – A contemporary Soundtrack" and released in the UK in 1991. This was followed by "Moods 2" in 1992. The series focuses on the genres of new-age, ambient, world music, and to a lesser extent, downtempo, trip-hop and smooth jazz.</p> <p>Doc 8: Christina Lux, Alquimia and Gleisberg; and several compilations of meditative and world-influenced New Age music including "Musical Healing" in 2001 and "Perfect Balance" in 2006. In 2006, the Gordon brothers released "Soothing Sanctuary", ambient music with nature sounds, a re-version of the earlier release "Sanctuary" (so stated in the liner notes) with new arrangements and remastering.</p> <p>Doc 14: number that expanded to 200 as of March 2010 with the release of "Shaman", a duo recording featuring Jean Derome and Joane Héту. Presently DAME has a catalogue of over 450 titles. Ambiances Magnétiques began in 1982 when René Lussier, a guitarist, and Jean Derome, who plays saxophone and flute, performed live as a duo in Montreal under the name Ambiances Magnétiques</p> <p>(Total 50 Documents)</p> <p>Response: Ambient ✗</p>
DynamicRAG	<p>(Reranker Results:)</p> <p>Doc 5: over the fact that the album was not released as of May 2018, in the group's 2018 Mannheim Steamroller Christmas Tour announcement, founder Chip Davis announced that Exotic Spaces would be officially released upon the start of the tour, being sold at all concert locations as well as on www.mannheimsteamroller.com and Amazon.com. "Billboard"s Top New Age Albums chart became the New Age Albums chart in June 2009. Mannheim Steamroller is an American Neoclassical new-age music group founded by Chip Davis that is known primarily for its "Fresh Aire" series of albums, which blend classical music with elements of</p> <p>Doc 17: Neoclassical new-age music content: Neoclassical new-age music Within the broad movement of new-age music, neoclassical new-age music, or instrumental pop, is influenced by and sometimes also based upon early, baroque or classical music, especially in terms of melody and composition.</p> <p>Doc 3: to an increase CD sales and eventually awards with a Western Canadian Music Awards (WCMA) (best dance) and a Juno Award nomination (best instrumental). Canadian success and a trip to MIDEM in France led to a deal with Bay Area indie label XDOT25 which released the EP in the USA.</p> <p>(Total 7 Documents)</p> <p>Response: Mannheim Steamroller ✓</p>
Ground Truth	Mannheim Steamroller

Case Study2

Table 8. Case Study for DynamicRAG. We compare our method with Vanilla-RAG with Reranker. The retrieved content is same for both method. We take Top-40 documents as the input.

Question	where is arachidonic acid found in the body
GPT-4o	Cell Membranes, Brain, Liver, Platelets, Eicosanoid Production Sites.
Vanilla-RAG w/Reranker	<p>(Reranker Results:)</p> <p>Doc 1: present in the phospholipids (especially phosphatidylethanolamine, phosphatidylcholine, and phosphatidylinositides) of membranes of the body’s cells, and is abundant in the brain, muscles, and liver. Skeletal muscle is an especially active site of arachidonic acid retention, accounting for roughly 10-20% of the phospholipid fatty acid content typically. In addition to being involved in cellular signaling as a lipid second messenger involved in the regulation of signaling enzymes, such as PLC-, PLC, and PKC, and isoforms, arachidonic acid is a key inflammatory intermediate and can also act as a vasodilator. (Note separate synthetic pathways, as described in section below.)</p> <p>Doc 3: Arachidonic acid Arachidonic acid (AA, sometimes ARA) is a polyunsaturated omega-6 fatty acid 20:4(-6), or 20:4(5,8,11,14). It is structurally related to the saturated arachidic acid found in cupuaçu butter ("L. arachis" – peanut). In chemical structure, arachidonic acid is a carboxylic acid with a 20-carbon chain and four "cis"-double bonds; the first double bond is located at the sixth carbon from the omega end. Some chemistry sources define arachidonic acid to designate any of the eicosatetraenoic acids. However, almost all writings in biology, medicine, and nutrition limit the term to all "cis"-5,8,11,14-eicosatetraenoic acid. Arachidonic acid is a polyunsaturated fatty acid...</p> <p>Doc 5: arachidonic acid supplementation for Alzheimer’s patients are needed. Another study indicates that air pollution is the source of inflammation and arachidonic acid metabolites promote the inflammation to signal the immune system of the cell damage. Arachidonic acid is marketed as an anabolic bodybuilding supplement in a variety of products. Supplementation of arachidonic acid...</p> <p>(Total 50 Documents)</p> <p>Response: brain, muscles, and liver ✓</p>
DynamicRAG	<p>(Reranker Results:)</p> <p>None</p> <p>(Total 0 Documents)</p> <p>Response: brain, muscles, liver ✓</p>
Ground Truth	brain, muscles and liver.

Thanks!