

RetroLLM: Empowering Large Language Models to Retrieve Fine-grained Evidence within Generation

Xiaoxi Li¹, Jiajie Jin¹, Yujia Zhou², Yongkang Wu³, Zhonghua Li³,
Qi Ye³, Zhicheng Dou^{1,4*}

¹Gaoling School of Artificial Intelligence, Renmin University of China

²Tsinghua University ³Huawei Poisson Lab

⁴Beijing Key Laboratory of Research on Large Models and Intelligent Governance
{xiaoxi_li, dou}@ruc.edu.cn

Abstract

Large language models (LLMs) exhibit remarkable generative capabilities but often suffer from hallucinations. Retrieval-augmented generation (RAG) offers an effective solution by incorporating external knowledge, but existing methods still face several limitations: additional deployment costs of separate retrievers, redundant input tokens from retrieved text chunks, and the lack of joint optimization of retrieval and generation. To address these issues, we propose **RetroLLM**, a unified framework that integrates retrieval and generation into a single, auto-regressive process, enabling LLMs to directly generate fine-grained evidence from the corpus with constrained decoding. Moreover, to mitigate false pruning in the process of constrained evidence generation, we introduce (1) hierarchical FM-Index constraints, which generate corpus-constrained clues to identify a subset of relevant documents before evidence generation, reducing irrelevant decoding space; and (2) a forward-looking constrained decoding strategy, which considers the relevance of future sequences to improve evidence accuracy. Extensive experiments on five open-domain QA datasets demonstrate RetroLLM’s superior performance across both in-domain and out-of-domain tasks. The code is available at <https://github.com/sunnyxexus/RetroLLM>.

1 Introduction

Large language models (LLMs) exhibit remarkable capabilities and are widely applied in various domains (Zhao et al., 2023; Zhu et al., 2023; Qiao et al., 2024). However, due to their reliance on model internal memory, they often struggle with long-tail or newly updated knowledge, leading to the issue of “hallucinations” (Huang et al., 2023). To address this, retrieval-augmented generation (RAG) has emerged as a promising solution. By integrating the retrieval of external knowledge, RAG

enables models to access up-to-date and factual information, enhancing the accuracy and reliability of their responses (Lewis et al., 2020; Gao et al., 2024).

Existing RAG methods typically rely on a separate dense retriever to fetch top- k text chunks for LLMs to generate answers, as shown in Figure 1(a). However, these methods face several limitations: (1) Maintaining a separate retriever increases deployment costs (Zhang et al., 2024a); (2) Retrieved documents often contain redundant information, consuming vast input tokens and distracting the model’s attention from key information (Jiang et al., 2023b); (3) The fixed granularity and number of retrieved text chunks limits flexibility of RAG systems (Qian et al., 2024); and (4) Retrievers rely on standalone document indices, hindering joint optimization with generators. Since retrieval and generation are inherently interconnected, jointly learning these tasks can enhance the overall performance of RAG systems (Li et al., 2024c). Thus, we aim to develop a unified framework that seamlessly integrates retrieval and generation processes.

Recently, generative retrieval (GR) has emerged as a promising approach that leverages generative models to directly generate document identifiers (DocIDs), eliminating the need for document indices and making it possible for joint optimization (Li et al., 2024b; Tay et al., 2022; Li et al., 2024c,a). However, existing GR methods still require mapping the generated DocIDs back to the document content before these can be used by LLMs for answer generation, as depicted in Figure 1(b). This step disrupts the seamless integration of retrieval and generation processes.

To address the above limitations, we propose **RetroLLM**, which empowers LLMs to generate factual evidence from knowledge sources and a final answer within a unified, auto-regressive decoding process, as shown in Figure 1(c). RetroLLM enables the model to autonomously decide how

*Corresponding author.

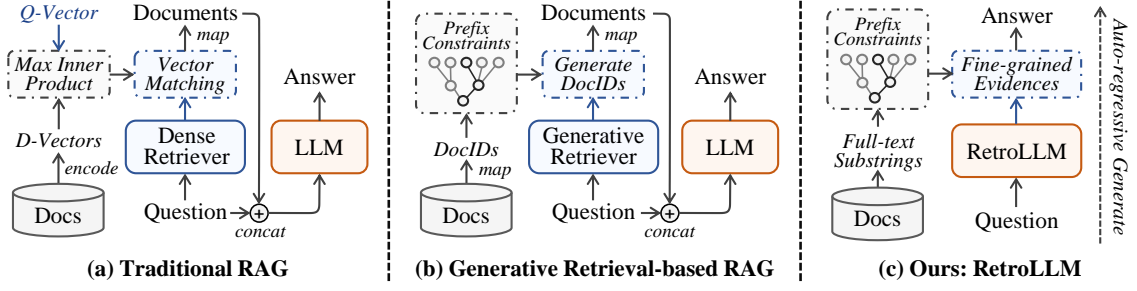


Figure 1: Comparison of retrieval-augmented generation frameworks. (a) Traditional RAG uses a dense retriever for document matching, while (b) generative RAG relies on constrained DocID generation. Both require feeding retrieved document text into the LLM for answer generation. (c) Our RetroLLM unifies retrieval and generation in a single auto-regressive decoding process, leveraging FM-Index constraints to retrieve fine-grained evidence.

much evidence to retrieve and when to generate the final response, eliminating the need for a separate embedding model and enhancing the flexibility of the RAG system. Furthermore, RetroLLM achieves joint optimization of retrieval and generation, facilitating a deeper understanding of their relationships and improving overall performance.

To achieve this, a simple approach is to apply constrained beam search based on FM-Index to generate factual evidence contained in corpus (Jain et al., 2024). However, the prefix-constrained approach suffers from a severe **false pruning** problem, where correct evidence sequences are often pruned due to errors in early decoding steps (Zhang et al., 2023). While initial prefixes may appear relevant, subsequent decoded sequences often reveal that they are grounded in irrelevant documents, leading to failure in generating relevant evidence. This issue arises from two main challenges: (1) Large corpora result in a vast number of prefix choices during early constrained decoding steps, making it difficult to choose the correct one; (2) It is also difficult to predict the relevance of subsequent sequences based solely on a short prefix. (See Appendix B for details)

To address these challenges, we propose two key strategies: (1) We construct a **hierarchical FM-Index**, which first enables the LLM to generate corpus-constrained clues to identify a subset of candidate documents. Evidence is then generated under the constraints of this subset’s FM-Index, significantly reducing the irrelevant decoding space, especially in the early steps. (2) For evidence generation, we introduce **forward-looking constrained decoding**, which utilizes the clues and the document FM-Indexes to identify future windows in the candidate documents. A relevance model then scores these windows, guiding the generation of

relevant evidence with future window relevance.

We conduct extensive experiments on five open-domain QA datasets, testing both in-domain and out-of-domain tasks, different base LLMs, and parameter sizes. Our results demonstrate the superior performance of RetroLLM compared to traditional RAG and complex RAG strategies.

The main contributions of this paper are: (1) We propose RetroLLM, a unified framework that unifies retrieval and generation into a single auto-regressive process, eliminating the need for a separate retriever and enabling joint optimization of RAG tasks. (2) To reduce irrelevant decoding paths in constrained evidence generation, we propose hierarchical FM-Index constraints and first predict clues with corpus FM-Index constraints to identify a document subset. (3) We introduce forward-looking constrained decoding, which identifies candidate windows based on clues and leverages future window relevance scores to guide the model in generating relevant evidence.

2 Preliminary

2.1 Task Formulation

Retrieval-augmented generation (RAG) leverages external knowledge sources to enhance the accuracy of language model generations. In this work, we formulate RAG within a generative framework. We divide the task into two sub-problems:

Constrained Evidence Generation: This involves retrieving relevant evidence from a large corpus in a generative manner with pre-built constraints. Formally, let \mathcal{C} be the corpus of documents and q be the input query. The constrained evidence generation process can be formulated as:

$$P(e|q, \mathcal{C}) = \prod_{t=1}^{T_e} P(e_t | e_{<t}, q, \mathcal{I}_c), \quad (1)$$

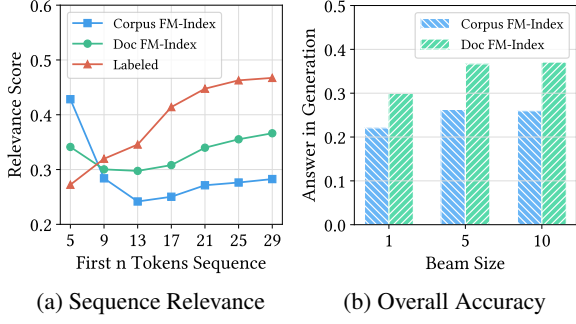


Figure 2: Empirical Study on false pruning problem in constrained evidence generation, comparing corpus-level and document-level FM-Index approaches. The relevance score is calculated by bge-reranker-large.

where e is the generated evidence sequence with length T_e , $e_{<t}$ is the tokens generated before position t , and $\mathcal{I}_c = \text{FM-Index}(\mathcal{C})$ represents the FM-Index built on corpus \mathcal{C} .

Answer Generation: Based on the retrieved evidence, the language model continues to generate the final answer, which can be expressed as:

$$P(a|q, e) = \prod_{t=1}^{T_a} P(a_t|a_{<t}, q, e), \quad (2)$$

where a is the generated answer with length T_a and $a_{<t}$ denotes the tokens generated before position t .

2.2 Empirical Study

To enable language models to generate evidence in a knowledge corpus, a natural approach is to apply FM-Index constraints over the entire corpus. However, our preliminary experiments reveal a critical limitation: while the initially generated evidence sequence usually appears relevant, later generated tokens often reveal that it has been grounded to irrelevant documents under FM-Index constraints, resulting in incorrect evidence prediction. This phenomenon is known as **false pruning**, where relevant sequences are eliminated prematurely during beam search (Detailed analysis in Appendix B).

To quantify this issue, we conducted an empirical study. Figure 2(a) illustrates how the relevance between query and generated evidence prefix changes during decoding. The results show that compared to labeled evidence sequences, the prefix relevance under corpus FM-Index constraints experiences a significant decline, particularly severe within the first 13 tokens. When we restrict the FM-Index constraints to only relevant documents, this degradation is substantially reduced and evidence generation accuracy improves over different beam sizes (Figure 2(b)). This finding suggests that

constraining the search space to a curated subset of relevant documents effectively mitigates false pruning, guiding the development of our strategies.

3 RetroLLM: Retrieval in Generation

In this section, we introduce RetroLLM, a unified LLM for RAG via auto-regressive decoding. The decoding process includes a clue and evidence stage for retrieval and an answer generation stage.

3.1 Hierarchical FM-Index Constraints

Before model generation, we construct hierarchical FM-Indexes for different levels of constraints for clue and evidence generation stages, including: **(1) a corpus-level global FM-Index** \mathcal{I}_c built from the entire corpus: $\mathcal{I}_c = \text{FM-Index}(\mathcal{C})$; and **(2) a document-level FM-Index manager** \mathcal{I}_d built for each document: $\mathcal{I}_d = \text{FM-Index}(d) : d \in \mathcal{C}$. The global index \mathcal{I}_c is primarily used during the clue generation stage to ensure generated phrases exist in the corpus, while document-level indexes \mathcal{I}_d are employed during evidence generation to constrain outputs to specific document d .

3.2 Clue Generation and Document Scoring

As discussed in Section 2.2, generating evidence with relevant document FM-Indexes could reduce the decoding paths and enhance accuracy. Therefore, we propose that the LLM first predict key phrases, or “clues,” that are likely to appear in relevant documents to retrieve subsets of documents.

Clue Generation. Given a query q , we first generate a set of clues \mathcal{C}_{gen} under corpus FM-Index constraints to predict key topics of relevant documents. For each clue $c_i \in \mathcal{C}_{\text{gen}}$, its generation probability can be formulated as:

$$P(c_i|q, c_{<i}, \mathcal{I}_c) = \prod_{t=1}^{T_i} P(c_{i,t}|c_{i,<t}, q, c_{<i}, \mathcal{I}_c) \quad (3)$$

where $c_{i,t}$ represents the t -th token of the i -th clue, $c_{i,<t}$ represents all previously generated tokens for the i -th clue, \mathcal{I}_c is the corpus-level FM-Index, and T_i is the length of the i -th clue. Clues are generated between the special tokens `<clue>` and `</clue>`, separated by the special token `<sepl>`.

With the predicted clues, we can obtain the appearance frequency $\text{CF}(c_i)$ of clue c_i in the corpus based on the corpus FM-Index, along with $\text{DF}(c_i)$ which is the document frequency, and $\text{TF}(c_i, d)$ which is the term frequency in document d . Drawing inspiration from TF-IDF (Robertson

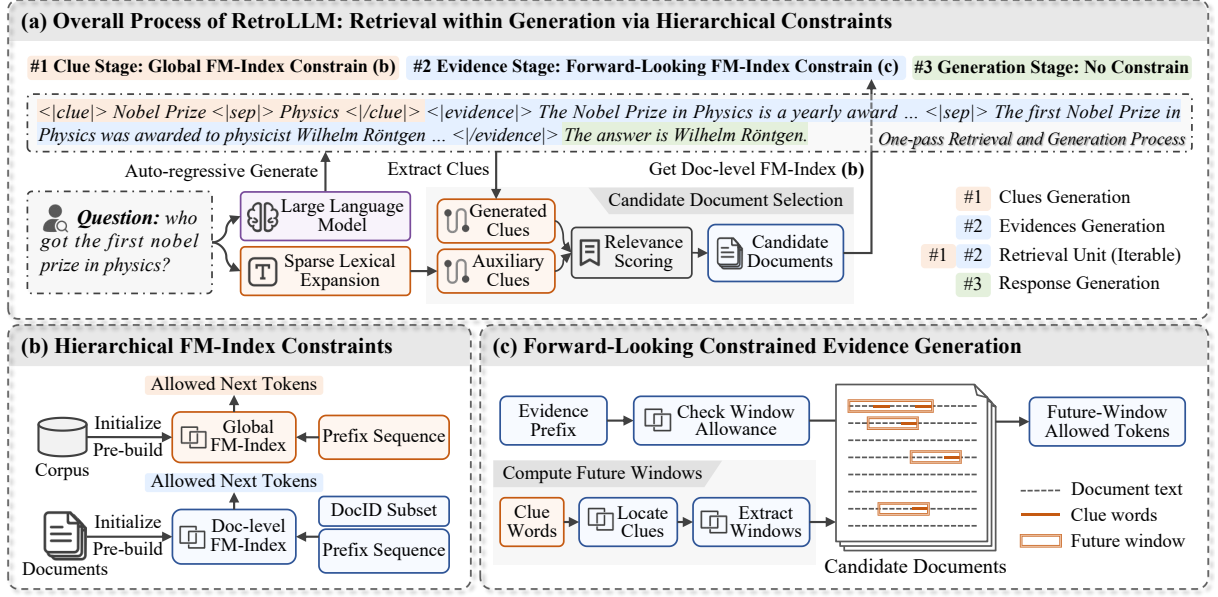


Figure 3: Overview of the RetroLLM Framework, which retrieves fine-grained evidence through a hierarchical, forward-looking FM-Index constrained generation process. During generation, the model autonomously determines whether to generate additional evidence or provide the final answer, based on the sufficiency of the current context.

and Zaragoza, 2009), we assign higher weights to clues that appear less frequently in the corpus and are present in fewer documents. For a document d , we calculate the relevance score as:

$$\mathcal{S}_{\text{gen}}(d) = \sum_{i=1}^{|\mathcal{C}_{\text{gen}}|} w_i \cdot f(c_i, d), \quad (4)$$

where w_i is the weight of the i -th clue, defined as:

$$w_i = \log \frac{N}{\text{CF}(c_i)} + \log \frac{N}{\text{DF}(c_i)}, \quad (5)$$

and $f(c_i, d)$ scores the document d for clue c_i :

$$f(c_i, d) = \log(1 + \text{TF}(c_i, d)), \quad (6)$$

where N is the total number of documents. Based on Equation (4), we form the ranking list $R_1(d)$ by selecting the top- k_{gen} documents from those containing at least one $c \in \mathcal{C}_{\text{gen}}$.

Auxiliary Clues. Although the generated clues could locate relevant documents, they typically contain only 1-3 key phrases, which may limit comprehensive document recall. To enhance retrieval robustness, we obtain auxiliary clues by employing a sparse lexical model f_{lex} that takes query q as input and assigns importance weights to each word in its vocabulary \mathcal{V}_{lex} :

$$w_{\text{lex}}(v) = f_{\text{lex}}(q) : v \in \mathcal{V}_{\text{lex}}. \quad (7)$$

Subsequently, we select the top- k_{aux} words from \mathcal{V}_{lex} as auxiliary clues set \mathcal{C}_{aux} . Now we form a

combined clue set $\mathcal{C}_{\text{all}} = \mathcal{C}_{\text{gen}} \cup \mathcal{C}_{\text{aux}}$. Additionally, we obtain a document ranking list $R_2(d)$ consisting of the top- k_{lex} documents retrieved by f_{lex} .

Rank Fusion. The final candidate documents are determined by combining the ranking lists from both generated and expanded clues using weighted reciprocal rank fusion, which can be expressed as:

$$\mathcal{S}(d) = w_1 \sum_{r \in R_1(d)} \frac{1}{r} + w_2 \sum_{r \in R_2(d)} \frac{1}{r}, \quad (8)$$

where w_1 and w_2 are the respective weights for $R_1(d)$ and $R_2(d)$, $\frac{1}{r}$ represents the reciprocal rank score. Finally, the top- k ranked documents form the candidate set \mathcal{D}_c for evidence generation.

3.3 Forward-Looking Constrained Evidence Generation

Now we have candidate documents, but a key challenge still remains: the model cannot foresee the relevance of future sequences when predicting the current token, making it difficult to decode tokens that lead to correct evidence sequences. To address this challenge, we propose a forward-looking constrained decoding strategy that enables the model to be aware of future sequence relevance.

Our strategy consists of three key components: (1) locating potential future windows that contain query-relevant information, (2) scoring the relevance of these windows, and (3) adjusting decoding logits based on future relevance. Given a candidate

document set \mathcal{D}_c , $\mathcal{I}_{\mathcal{D}_c}$ is its FM-Indexes, the evidence generation process can be formulated as:

$$P(e_i|q, e_{<i}, \mathcal{C}_{\text{all}}, \mathcal{I}_{\mathcal{D}_c}) = \prod_{t=1}^{T_i} P(e_{i,t}|e_{i,<t}, q, e_{<i}, \mathcal{C}_{\text{all}}, \mathcal{I}_{\mathcal{D}_c}, \mathcal{W}), \quad (9)$$

where e_i represents the i -th evidence sequence, \mathcal{C}_{all} contains the generated clues, and $\mathcal{W}_{\text{info}}$ encapsulates future window information.

Locate Future Windows. First, we identify window sequences containing clues \mathcal{C}_{all} in the candidate document set \mathcal{D}_c , as these contexts typically exhibit high relevance to the query. We obtain all future window sequences \mathcal{W}_{raw} through document-specific FM-Indexes:

$$\mathcal{W}_{\text{raw}} = \bigcup_{d \in \mathcal{D}_c} \bigcup_{c \in \mathcal{C}_{\text{all}}} \text{Ext}(\mathcal{I}_d, \text{Loc}(\mathcal{I}_d, c), l_w). \quad (10)$$

Here, $\text{Loc}(\mathcal{I}_d, c)$ locates clue positions in document d 's FM-Index \mathcal{I}_d , and $\text{Ext}(\mathcal{I}_d, p_c, l_w)$ extracts sequences of length l_w around these positions. We then merge overlapping windows from \mathcal{W}_{raw} to create the candidate future window set \mathcal{W} , with each merged window not exceeding length l_{max} .

Future Window Relevance. We employ a reranker model f_{rel} to efficiently evaluate the relevance between each future window $w \in \mathcal{W}$ and the query q :

$$\mathcal{S}_w(w) = f_{\text{rel}}(q, w) : w \in \mathcal{W}. \quad (11)$$

Now for each $w \in \mathcal{W}$, we have its document source d , position p_c , and relevance score $\mathcal{S}_w(w)$.

Logits Adjustment. During decoding, we adjust token logits to favor sequences from highly relevant future windows. At each step, for allowed tokens $\mathcal{V}_{\text{allowed}}$ determined by FM-Index constraints, we locate each token's positions \mathcal{P}_t and identify its corresponding future windows \mathcal{W}_t . The adjusted logits are computed as:

$$\tilde{l}(t) = \begin{cases} l(t) + \lambda \cdot \max_{w \in \mathcal{W}_t} \mathcal{S}_w(w), & \text{if } t \in \mathcal{V}_{\text{allowed}} \\ -\infty, & \text{otherwise} \end{cases} \quad (12)$$

where $l(t)$ is the original logits and λ controls the weight of future relevance. With logits adjustment, the token probability in Equation (9) is then computed as $P(e_{i,t}|\dots) = \text{softmax}(\tilde{l}(t))$.

The evidence generation process continues until sufficient information is collected, controlled by special tokens: `<evidence>` begins evidence generation stage, `<sepl>` triggers the next evidence generation under constraints, while `</evidence>` signals completion and transitions to free generation.

3.4 Answer Generation

With the relevant evidences \mathcal{E} generated, the model proceeds to generate the final answer to the original query q , which can be formulated as:

$$P(a|q, \mathcal{C}_{\text{gen}}, \mathcal{E}) = \prod_{t=1}^{T_a} P(a_t|a_{<t}, q, \mathcal{C}_{\text{gen}}, \mathcal{E}), \quad (13)$$

where a is the generated answer sequence of length T_a , a_t is the token at position t in the answer, $a_{<t}$ denotes generated tokens before position t .

3.5 Training of RetroLLM

Since RetroLLM's entire RAG process is one-pass and auto-regressive, we can construct target sequences for supervised fine-tuning to achieve joint learning of retrieval and generation tasks.

Training Data Construction. We simulate the model's inference process to construct training data. For each QA pair (q, a) , we: (1) Use a sparse retriever to obtain clues \mathcal{C}_{aux} and retrieve relevant documents. (2) Locate sentences containing $c \in \mathcal{C}_{\text{aux}}$ within the documents. (3) Apply a reranker to select the top- k_e relevant evidences. (4) Identify evidences that both contain the answer a and are confirmed by an LLM to genuinely answer the query q . (5) Select the top- k evidences up to the first relevant one. (6) For target clues, we utilize an LLM to extract key entities from the query and relevant evidences. An example of the output format is illustrated in Figure 3, and additional details are provided in Appendix D.3.2.

Model Optimization. Since evidence is typically longer compared to clues and answer, we mask out 80% of the tokens in the middle of each target evidence. We employ the standard next token prediction loss as follows:

$$\mathcal{L} = - \sum_{t=1}^{T_c+T_e} \log P(x_t|x_{<t}, q; \theta) - \gamma \sum_{t=1}^{T_a} \log P(y_t|y_{<t}, x, q; \theta), \quad (14)$$

where θ represents the parameters of RetroLLM, x and y represent the target sequence of clues + evidences and answer respectively, and γ is the weight for the answer loss.

4 Experimental Settings

4.1 Datasets and Evaluation Metrics

We conduct experiments on five open-domain QA datasets, including single-hop QA: NQ (Kwiatkowski et al., 2019), TriviaQA (Joshi

Table 1: Overall performance on open-domain QA datasets, including single-hop and multi-hop QA tasks. The best results are in **bold** and the second are underlined. Results from non-proprietary models are in gray color.

Method	In-domain Datasets									Out-of-domain Datasets					
	NQ			TriviaQA			HotpotQA			PopQA			2WIKI		
	Acc	F1	Tok	Acc	F1	Tok	Acc	F1	Tok	Acc	F1	Tok	Acc	F1	Tok
<i>Direct Generation</i>															
Llama3-8B	27.6	30.1	50	56.1	60.2	52	21.2	26.5	56	24.2	26.4	43	20.9	24.3	54
Mistral-7B	30.4	25.2	57	58.8	58.6	57	27.0	23.6	65	25.8	25.2	45	36.5	18.7	58
Qwen2.5-7B	21.8	21.3	52	45.1	48.1	54	21.3	27.5	57	17.1	18.7	45	22.4	28.1	53
ChatGPT	-	-	-	77.0	52.9	-	33.8	24.0	-	26.6	13.2	-	38.0	21.3	-
<i>Retrieval-augmented Generation</i>															
Naive RAG	<u>52.4</u>	41.1	919	69.3	65.9	<u>915</u>	<u>37.8</u>	35.8	<u>960</u>	47.7	38.6	944	<u>38.7</u>	21.7	<u>1000</u>
REPLUG	41.6	41.2	<u>903</u>	65.4	66.5	939	27.8	31.7	965	38.2	37.0	<u>921</u>	24.5	20.8	1007
FLARE	46.8	45.4	1515	67.6	68.9	1418	36.9	34.8	1715	52.4	41.5	1461	34.4	<u>35.2</u>	1795
Self-RAG	41.8	45.2	1203	64.1	53.4	1267	32.1	29.6	1354	39.7	32.7	1236	30.3	25.7	1272
IRCoT	49.6	45.9	1598	66.0	66.1	1715	37.3	<u>41.5</u>	1842	59.8	<u>45.6</u>	1667	29.4	32.4	1707
Iter-RetGen	51.7	<u>48.4</u>	3002	<u>71.0</u>	<u>69.9</u>	2461	37.2	39.0	2545	51.7	47.5	2509	29.2	21.5	2669
Adaptive-RAG	50.5	46.6	946	65.1	65.6	958	37.1	39.1	2080	58.3	40.4	1681	32.1	28.4	2580
<i>Retrieval within Generation</i>															
RetroLLM (Ours)	61.6	49.8	302	74.3	72.8	287	61.9	47.2	607	65.7	43.0	355	48.9	36.2	661

et al., 2017), PopQA (Mallen et al., 2023), and multi-hop QA: HotpotQA (Yang et al., 2018), 2WikiMultiHopQA (2WIKI) (Ho et al., 2020). See Appendix C for detailed statistics. For evaluation metrics, we use Accuracy (Acc), F1 score, and token count (Tok) to assess the quality of generated answers as well as the total number of input and output tokens consumed by LLMs.

4.2 Baselines

The baseline methods include two types: (1) Direct generation: This includes open-source models Llama3-8B (Dubey et al., 2024), Mistral-7B (Jiang et al., 2023a), and Qwen2.5-7B (Yang et al., 2024), and the non-proprietary model ChatGPT (OpenAI, 2022) with results taken from (Zhang et al., 2024b). (2) Retrieval-augmented generation: This includes Naive RAG method and several complex RAG methods, including REPLUG (Shi et al., 2023), Self-RAG (Asai et al., 2024), IRCoT (Trivedi et al., 2023), Iter-RetGen (Shao et al., 2023), and Adaptive-RAG (Jeong et al., 2024). For a fair comparison, all RAG baselines use the E5-base-en (Wang et al., 2022a) retriever, and all LLMs are instruction-tuned with 7B or 8B parameters.

4.3 Implementation Details

Our knowledge source is based on the Wikipedia dump from December 20, 2018, in alignment with DPR (Karpukhin et al., 2020). We use Mistral-7B-Instruct as Backbone LLM. We limit the maximum

number of evidence for single-hop and multi-hop QA to 5 and 10, respectively. We set w_1 and w_2 to 1 and 2, respectively, and λ to 100. For efficient model training, we employ LoRA (Hu et al., 2022), setting training epochs to 3 and γ to 2. We use SPLADE-v3 (Lassance et al., 2024) for clue expansion and use BGE-reranker-base (Xiao et al., 2024) as the scoring model for window sequences. We implement FM-Index based on the sds-lite library (Gog et al., 2014). (Details in Appendix D)

4.4 Experimental Results

Overall Performance We evaluate RetroLLM’s overall downstream performance using NQ, TriviaQA, and HotpotQA for in-domain tasks, and PopQA and 2WIKI for out-of-domain tasks. The results are presented in Table 1. We find that: (1) RAG methods generally outperform direct generation methods (except for ChatGPT), highlighting the knowledge-intensive nature of these tasks that need retrieval augmentation. (2) **RetroLLM outperforms RAG methods across both in-domain and out-of-domain tasks.** This highlights the effectiveness of our unified RAG framework in mastering both evidence retrieval and answer generation, while also demonstrating strong generalization capabilities to unfamiliar domains, which is a crucial and challenging aspect for existing generative retrieval methods. (3) **RetroLLM significantly reduces token consumption (“Tok”) across all datasets compared to RAG meth-**

Table 2: Analysis of retrieval performance of RetroLLM, compared with sparse, dense, and generative retrieval methods. We report average performance on three single-hop and two multi-hop QA datasets.

Method	Single-hop QA			Multi-hop QA		
	R@1	R@5	Num	R@1	R@5	Num
BM25	37.8	56.3	5	26.9	43.1	5
SPLADE-v3	50.6	69.7	5	27.5	42.9	5
E5	54.3	74.3	5	26.9	45.9	5
BGE	53.3	72.8	5	27.4	46.8	5
Naive Constrain	15.7	31.7	5	10.6	20.3	5
RetroLLM	56.6	67.9	3.29	29.3	49.6	4.24

ods. On average, we use approximately 2.1x fewer tokens than Naive RAG and 6x fewer than Iter-RetGen. This efficiency is attributed to RetroLLM’s capability to retrieve fine-grained evidence and dynamically decide the amount of retrieved evidence.

Analysis of Retrieval Performance We also analyze the retrieval performance of RetroLLM compared to sparse and dense retrieval baselines, as shown in Table 2. (1) For single-hop QA tasks, RetroLLM demonstrates superior accuracy on R@1, thanks to the design of clues and future windows, which help precisely locate the relevant evidence. However, the R@5 is lower than strong baselines like E5, as it employs flexible retrieval and uses fewer passages on average (3.29 vs. 5 for baselines). (2) For multi-hop QA tasks, RetroLLM shows superior accuracy compared to all other methods for both R@1 and R@5, while utilizing a smaller average number of 4.24 retrieved passages. (3) Notably, the naive generative retrieval method using constrained beam search performs poorly on all metrics, validating the severity of false pruning.

Ablation Study Table 3 presents the results of the ablation study, evaluating the effectiveness of each component of RetroLLM. It can be observed that: (1) Removing the future window, clue generation, and clue expansion all lead to performance degradation, demonstrating the effectiveness of these specially designed components, as they play an important role in alleviating the false pruning problem in the prefix constraint-based method. (2) Adopting the most basic constrained evidence generation method results in the lowest performance, even lower than without constraints, demonstrating the severity of false pruning. (3) Without constraints, while the model avoids the false pruning

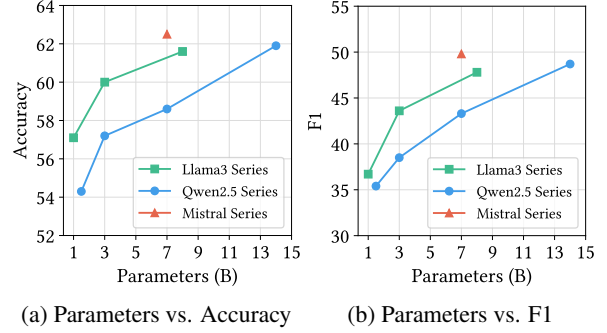


Figure 4: Impact of performance with different base LLMs, reporting average performance on five datasets.

Table 3: Ablation Studies of RetroLLM, considering in-domain and out-of-domain performance.

Method	In-domain		Out-of-domain	
	Acc	F1	Acc	F1
RetroLLM	66.0	56.6	57.3	39.6
w/o Future Window	44.3	43.2	40.9	33.8
w/o Clue Generation	60.6	52.1	56.4	38.1
w/o Clue Expansion	49.6	45.1	44.1	35.4
w/ Naive Constraints	27.2	28.0	21.8	20.7
w/o Constraints	41.6	43.0	31.6	28.1

problem, its performance still notably decreases due to the inability to utilize external knowledge.

Impact of Different Base LLMs To evaluate the performance of RetroLLM using different backbone LLMs with varying parameter sizes, we conducted experiments using the Mistral, Llama3, and Qwen2.5 series, with parameters ranging from 1B to 14B. The results are shown in Figure 4. We observe that: (1) As the parameter size increases, RetroLLM’s performance steadily improves, aligning with the scaling law; (2) There are slight performance differences across the different models (Mistral, Llama3, Qwen2.5). Nonetheless, all models confirm the effectiveness of RetroLLM.

Impact of Generated Evidence Quantity Since RetroLLM can dynamically determine the number of evidence to retrieve, we investigated the effect of different maximum retrieval quantities on performance, with results shown in Figure 5. When retrieving up to 1-5 evidence, performance continues to improve as the number of retrieved pieces increases, suggesting that more evidence contributes to stronger performance in these tasks. However, for multi-hop QA, performance stabilizes around 6 evidence, as more evidence can bring in both useful and distracting information, thereby limiting




Figure 5: Impact on maximum number of generated evidence, reporting average performance on single-hop and multi-hop QA tasks.

further performance gains.

Analysis of Efficiency We also evaluated the efficiency of RetroLLM, considering query latency, token count, and overall performance (see Table 4). We found that: **(1) Latency:** RetroLLM is slightly slower than Naive RAG but significantly faster than other more complex RAG methods. **(2) Token Count:** RetroLLM requires fewer input tokens as it processes only the query, unlike baselines that include retrieved passages. While output tokens are slightly higher due to fine-grained generated evidence. The total token count is significantly reduced due to more precise retrieval granularity. **(3) Performance:** RetroLLM achieves better results, providing an improved cost-efficiency balance.

5 Related Work

Retrieval-augmented Generation Retrieval-augmented generation (RAG) improves generation quality by incorporating relevant context from external knowledge bases, which typically employ a separate dense retriever (Gao et al., 2024; Tan et al., 2024b; Jin et al., 2024b; Dong et al., 2024a; Cheng et al., 2025; Tan et al., 2024a; Jin et al., 2025). Based on training approaches, current RAG systems fall into three categories: (1) Directly prompt of generative models with retrieved context (Press et al., 2023; Trivedi et al., 2023; Li et al., 2025a); (2) Separately training of retriever and/or generator (Karpukhin et al., 2020; Asai et al., 2024; Dong et al., 2025b, 2024b; Zhu et al., 2024; Dong et al., 2024c; Li et al., 2025b; Dong et al., 2025a); and (3) Jointly training of retriever and generator (Lewis et al., 2020; Singh et al., 2021). However, joint training faces challenges due to the architectural differences between retrieval and generation, as well as

Table 4: Efficiency Analysis of RetroLLM, comparing query latency, number of tokens and performance (# P).

Method	Latency (ms)			Token Num			# P
	Retr	Gen	Total	In	Out	Total	F1
Naive RAG	54	528	582	902	17	919	41.1
SelfRAG	89	3180	3269	1096	107	1203	45.2
Iter-RetGen	274	2058	2332	2963	39	3002	48.4
IRCoT	83	1759	1842	1535	63	1598	46.6
RetroLLM	-	-	786	18	297	315	49.8

the need for updating document indices during training. Some approaches aim to unify dense retrieval and generation within a single model, including GritLM (Muennighoff et al., 2024) and OneGen (Zhang et al., 2024a). However, GritLM operates as two distinct models with separate attention mechanisms that share parameters, while OneGen still relies on retrieving passage chunks as input for subsequent generation.

Generative Retrieval Generative retrieval (GR) retrieves by directly generating document identifiers (DocIDs) without the need for traditional document indices (Metzler et al., 2021). Research in this area focuses on: (1) DocID design, including numeric-based DocIDs (Tay et al., 2022; Wang et al., 2022b; Jin et al., 2023; Zeng et al., 2023) and text-based DocIDs (Cao et al., 2021; Bevilacqua et al., 2022; Chen et al., 2022; Li et al., 2023b); (2) DocID memorization strategies, including pseudo-query data augmentation (Zhuang et al., 2022), incorporating ranking feedback (Li et al., 2023a; Tang et al., 2024), and learnable DocIDs (Sun et al., 2023; Yang et al., 2023). However, these methods mainly focus on optimizing retrieval tasks, without considering their connections with downstream tasks. Even though UniGen (Li et al., 2024c) and CorpusLM (Li et al., 2024a) address downstream tasks, they still require mapping the generated DocIDs to the corresponding documents before feeding them into the generator. While RICHES (Jain et al., 2024) attempts to streamline this process, but fails to solve the false pruning issue, which leads to suboptimal downstream task performance.

6 Conclusion

This paper introduces RetroLLM, a unified LLM that seamlessly integrates retrieval and generation by generating evidence from a corpus with FM-Index constraints. This approach eliminates the need for separate retrievers and reduces redundant

input. To improve evidence accuracy, we proposed hierarchical FM-Index constraints and a forward-looking decoding strategy, helping the model focus on relevant information. Experiments show that RetroLLM outperforms existing methods on open-domain QA tasks, marking a step towards a new era of generative retrieval-augmented generation.

7 Limitations

While RetroLLM demonstrates strong performance across various open-domain QA scenarios, it has several limitations that present opportunities for future research:

(1) To improve the robustness of the model-generated clues, we still need to perform clue expansion to ensure the system’s superior performance, as discussed in Section 4.4. This prevents a fully end-to-end optimization of the RAG task. Future work could focus on designing mechanisms that eliminate this need, enabling complete end-to-end RAG optimization.

(2) In terms of efficiency, RetroLLM outperforms most complex RAG methods in query latency. However, it is slightly slower than Naive RAG, as the generated evidence results in more output tokens despite being fine-grained and short. Drawing on the concept of speculative decoding (Leviathan et al., 2023; Xia et al., 2023), future improvements could involve using a smaller language model during the constrained evidence generation phase and switching to a larger model during answer generation. This approach could enhance RetroLLM’s efficiency, comprehensively surpassing existing RAG methods in both performance, latency, and flexibility.

(3) RetroLLM currently only considers the unification of evidence retrieval and final answer generation. It would be worth exploring the incorporation of more model reasoning processes within RetroLLM’s single generation step, such as query intent analysis, question decomposition, retrieval necessity assessment, retrieved evidence relevance judgment, and answer generation with source citations. This integration would contribute to building a more comprehensive, unified RAG system with just a single LLM.

Acknowledgement

This work was supported by Beijing Municipal Science and Technology Project No. Z231100010323009, National Science and Tech-

nology Major Project No. 2022ZD0120103, National Natural Science Foundation of China No. 62272467, Beijing Natural Science Foundation No. L233008, and the fund for building world-class universities (disciplines) of Renmin University of China. The work was partially done at the Engineering Research Center of Next-Generation Intelligent Search and Recommendation, MOE.

References

- Akari Asai, Zeqiu Wu, Yizhong Wang, Avirup Sil, and Hannaneh Hajishirzi. 2024. [Self-rag: Learning to retrieve, generate, and critique through self-reflection](#). In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.
- David M. Beazley. 1996. [SWIG: an easy to use tool for integrating scripting languages with C and C++](#). In *Fourth Annual USENIX Tcl/Tk Workshop 1996, Monterey, California, USA, July 10-13, 1996*. USENIX Association.
- Michele Bevilacqua, Giuseppe Ottaviano, Patrick S. H. Lewis, Scott Yih, Sebastian Riedel, and Fabio Petroni. 2022. [Autoregressive search engines: Generating substrings as document identifiers](#). In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*.
- Nicola De Cao, Gautier Izacard, Sebastian Riedel, and Fabio Petroni. 2021. [Autoregressive entity retrieval](#). In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Jiangui Chen, Ruqing Zhang, Jiafeng Guo, Yiqun Liu, Yixing Fan, and Xueqi Cheng. 2022. [Corpusbrain: Pre-train a generative retrieval model for knowledge-intensive language tasks](#). In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management, Atlanta, GA, USA, October 17-21, 2022*, pages 191–200. ACM.
- Yiruo Cheng, Kelong Mao, Ziliang Zhao, Guanting Dong, Hongjin Qian, Yongkang Wu, Tetsuya Sakai, Ji-Rong Wen, and Zhicheng Dou. 2025. [CORAL: benchmarking multi-turn conversational retrieval-augmented generation](#). In *Findings of the Association for Computational Linguistics: NAACL 2025, Albuquerque, New Mexico, USA, April 29 - May 4, 2025*, pages 1308–1330. Association for Computational Linguistics.
- Guanting Dong, Yifei Chen, Xiaoxi Li, Jiajie Jin, Hongjin Qian, Yutao Zhu, Hangyu Mao, Guorui Zhou, Zhicheng Dou, and Ji-Rong Wen. 2025a. [Tool-star: Empowering llm-brained multi-tool reasoner via reinforcement learning](#). *Preprint*, arXiv:2505.16410.

- Guanting Dong, Keming Lu, Chengpeng Li, Tingyu Xia, Bowen Yu, Chang Zhou, and Jingren Zhou. 2025b. [Self-play with execution feedback: Improving instruction-following capabilities of large language models](#). In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net.
- Guanting Dong, Xiaoshuai Song, Yutao Zhu, Runqi Qiao, Zhicheng Dou, and Ji-Rong Wen. 2024a. Toward general instruction-following alignment for retrieval-augmented generation. *arXiv preprint arXiv:2410.09584*.
- Guanting Dong, Hongyi Yuan, Keming Lu, Chengpeng Li, Mingfeng Xue, Dayiheng Liu, Wei Wang, Zheng Yuan, Chang Zhou, and Jingren Zhou. 2024b. [How abilities in large language models are affected by supervised fine-tuning data composition](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2024, Bangkok, Thailand, August 11-16, 2024*, pages 177–198. Association for Computational Linguistics.
- Guanting Dong, Yutao Zhu, Chenghao Zhang, Zechen Wang, Zhicheng Dou, and Ji-Rong Wen. 2024c. [Understand what LLM needs: Dual preference alignment for retrieval-augmented generation](#). *CoRR*, abs/2406.18676.
- 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 preprint arXiv:2407.21783*.
- Paolo Ferragina and Giovanni Manzini. 2000. [Opportunistic data structures with applications](#). In *41st Annual Symposium on Foundations of Computer Science, FOCS 2000, 12-14 November 2000, Redondo Beach, California, USA*, pages 390–398. IEEE Computer Society.
- Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, Qianyu Guo, Meng Wang, and Haofen Wang. 2024. [Retrieval-augmented generation for large language models: A survey](#). *Preprint*, arXiv:2312.10997.
- Simon Gog, Timo Beller, Alistair Moffat, and Matthias Petri. 2014. From theory to practice: Plug and play with succinct data structures. In *13th International Symposium on Experimental Algorithms, (SEA 2014)*, pages 326–337.
- Xanh Ho, Anh-Khoa Duong Nguyen, Saku Sugawara, and Akiko Aizawa. 2020. [Constructing A multi-hop QA dataset for comprehensive evaluation of reasoning steps](#). In *Proceedings of the 28th International Conference on Computational Linguistics, COLING 2020, Barcelona, Spain (Online), December 8-13, 2020*, pages 6609–6625. International Committee on Computational Linguistics.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. [Lora: Low-rank adaptation of large language models](#). In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.
- Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, and Ting Liu. 2023. [A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions](#). *CoRR*, abs/2311.05232.
- Palak Jain, Livio Baldini Soares, and Tom Kwiatkowski. 2024. [From RAG to RICHES: retrieval interlaced with sequence generation](#). *CoRR*, abs/2407.00361.
- Soyeong Jeong, Jinheon Baek, Sukmin Cho, Sung Ju Hwang, and Jong Park. 2024. [Adaptive-rag: Learning to adapt retrieval-augmented large language models through question complexity](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers), NAACL 2024, Mexico City, Mexico, June 16-21, 2024*, pages 7036–7050. Association for Computational Linguistics.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de Las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, L  lio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timoth  e Lacroix, and William El Sayed. 2023a. [Mistral 7b](#). *CoRR*, abs/2310.06825.
- Huiqiang Jiang, Qianhui Wu, Chin-Yew Lin, Yuqing Yang, and Lili Qiu. 2023b. [LLMLingua: Compressing prompts for accelerated inference of large language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 13358–13376. Association for Computational Linguistics.
- Bowen Jin, Hansi Zeng, Guoyin Wang, Xiusi Chen, Tianxin Wei, Ruirui Li, Zhengyang Wang, Zheng Li, Yang Li, Hanqing Lu, et al. 2023. Language models as semantic indexers. *arXiv preprint arXiv:2310.07815*.
- Jiajie Jin, Xiaoxi Li, Guanting Dong, Yuyao Zhang, Yutao Zhu, Yongkang Wu, Zhonghua Li, Qi Ye, and Zhicheng Dou. 2025. [Hierarchical document refinement for long-context retrieval-augmented generation](#). *Preprint*, arXiv:2505.10413.
- Jiajie Jin, Yutao Zhu, Xinyu Yang, Chenghao Zhang, and Zhicheng Dou. 2024a. [Flashrag: A modular toolkit for efficient retrieval-augmented generation research](#). *CoRR*, abs/2405.13576.
- Jiajie Jin, Yutao Zhu, Yujia Zhou, and Zhicheng Dou. 2024b. [BIDER: bridging knowledge inconsistency for efficient retrieval-augmented llms via key supporting evidence](#). *CoRR*, abs/2402.12174.

- Mandar Joshi, Eunsol Choi, Daniel Weld, and Luke Zettlemoyer. 2017. TriviaQA: A large scale distantly supervised challenge dataset for reading comprehension. In *ACL*, pages 1601–1611, Vancouver, Canada. Association for Computational Linguistics.
- 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 *EMNLP*, 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.
- Carlos Lassance, Hervé Déjean, Thibault Formal, and Stéphane Clinchant. 2024. [Splade-v3: New baselines for SPLADE](#). *CoRR*, abs/2403.06789.
- Yaniv Leviathan, Matan Kalman, and Yossi Matias. 2023. [Fast inference from transformers via speculative decoding](#). In *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, volume 202 of *Proceedings of Machine Learning Research*, pages 19274–19286. PMLR.
- Patrick S. H. Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. [Retrieval-augmented generation for knowledge-intensive NLP tasks](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Xiaoxi Li, Guanting Dong, Jiajie Jin, Yuyao Zhang, Yujia Zhou, Yutao Zhu, Peitian Zhang, and Zhicheng Dou. 2025a. [Search-o1: Agentic search-enhanced large reasoning models](#). *CoRR*, abs/2501.05366.
- Xiaoxi Li, Zhicheng Dou, Yujia Zhou, and Fangchao Liu. 2024a. [Corpuslm: Towards a unified language model on corpus for knowledge-intensive tasks](#). In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2024, Washington DC, USA, July 14-18, 2024*, pages 26–37. ACM.
- Xiaoxi Li, Jiajie Jin, Guanting Dong, Hongjin Qian, Yutao Zhu, Yongkang Wu, Ji-Rong Wen, and Zhicheng Dou. 2025b. [Webthinker: Empowering large reasoning models with deep research capability](#). *CoRR*, abs/2504.21776.
- Xiaoxi Li, Jiajie Jin, Yujia Zhou, Yuyao Zhang, Peitian Zhang, Yutao Zhu, and Zhicheng Dou. 2024b. [From matching to generation: A survey on generative information retrieval](#). *CoRR*, abs/2404.14851.
- Xiaoxi Li, Yujia Zhou, and Zhicheng Dou. 2024c. [Uni-gen: A unified generative framework for retrieval and question answering with large language models](#). In *Thirty-Eighth AAAI Conference on Artificial Intelligence, AAAI 2024, Thirty-Sixth Conference on Innovative Applications of Artificial Intelligence, IAAI 2024, Fourteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2014, February 20-27, 2024, Vancouver, Canada*, pages 8688–8696. AAAI Press.
- Yongqi Li, Nan Yang, Liang Wang, Furu Wei, and Wenjie Li. 2023a. [Learning to rank in generative retrieval](#). *CoRR*, abs/2306.15222.
- Yongqi Li, Nan Yang, Liang Wang, Furu Wei, and Wenjie Li. 2023b. [Multiview identifiers enhanced generative retrieval](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2023, Toronto, Canada, July 9-14, 2023*, pages 6636–6648. Association for Computational Linguistics.
- Alex Mallen, Akari Asai, Victor Zhong, Rajarshi Das, Daniel Khashabi, and Hannaneh Hajishirzi. 2023. [When not to trust language models: Investigating effectiveness of parametric and non-parametric memories](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2023, Toronto, Canada, July 9-14, 2023*, pages 9802–9822. Association for Computational Linguistics.
- Donald Metzler, Yi Tay, Dara Bahri, and Marc Najork. 2021. Rethinking search: making domain experts out of dilettantes. In *ACM SIGIR Forum*, volume 55, pages 1–27. ACM New York, NY, USA.
- Niklas Muennighoff, Hongjin Su, Liang Wang, Nan Yang, Furu Wei, Tao Yu, Amanpreet Singh, and Douwe Kiela. 2024. [Generative representational instruction tuning](#). *CoRR*, abs/2402.09906.
- OpenAI. 2022. [Introducing chatgpt](#). <https://openai.com/blog/chatgpt>.
- Ofir Press, Muru Zhang, Sewon Min, Ludwig Schmidt, Noah Smith, and Mike Lewis. 2023. [Measuring and narrowing the compositionality gap in language models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 5687–5711, Singapore. Association for Computational Linguistics.
- Hongjin Qian, Zheng Liu, Kelong Mao, Yujia Zhou, and Zhicheng Dou. 2024. [Grounding language model with chunking-free in-context retrieval](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2024, Bangkok, Thailand, August 11-16, 2024*, pages 1298–1311. Association for Computational Linguistics.
- Runqi Qiao, Qiuna Tan, Guanting Dong, Minhui Wu, Chong Sun, Xiaoshuai Song, Zhuoma Gongque, Shanglin Lei, Zhe Wei, Miaoxuan Zhang, Runfeng Qiao, Yifan Zhang, Xiao Zong, Yida Xu, Muxi

- Diao, Zhimin Bao, Chen Li, and Honggang Zhang. 2024. [We-math: Does your large multimodal model achieve human-like mathematical reasoning?](#) *CoRR*, abs/2407.01284.
- Stephen E. Robertson and Hugo Zaragoza. 2009. [The probabilistic relevance framework: BM25 and beyond](#). *Found. Trends Inf. Retr.*, 3(4):333–389.
- Michael Schindler. 1997. A fast block-sorting algorithm for lossless data compression. In *Proc. Data Compression Conf*, volume 469. Citeseer.
- Zhihong Shao, Yeyun Gong, Yelong Shen, Minlie Huang, Nan Duan, and Weizhu Chen. 2023. [Enhancing retrieval-augmented large language models with iterative retrieval-generation synergy](#). *Preprint*, arXiv:2305.15294.
- Weijia Shi, Sewon Min, Michihiro Yasunaga, Minjoon Seo, Rich James, Mike Lewis, Luke Zettlemoyer, and Wen tau Yih. 2023. [REPLUG: retrieval-augmented black-box language models](#). *CoRR*, abs/2301.12652.
- Devendra Singh, Siva Reddy, Will Hamilton, Chris Dyer, and Dani Yogatama. 2021. End-to-end training of multi-document reader and retriever for open-domain question answering. *Advances in Neural Information Processing Systems*, 34:25968–25981.
- Weiwei Sun, Lingyong Yan, Zheng Chen, Shuaiqiang Wang, Haichao Zhu, Pengjie Ren, Zhumin Chen, Dawei Yin, Maarten de Rijke, and Zhaochun Ren. 2023. [Learning to tokenize for generative retrieval](#). *CoRR*, abs/2304.04171.
- Jiejun Tan, Zhicheng Dou, Wen Wang, Mang Wang, Weipeng Chen, and Ji-Rong Wen. 2024a. [Html-rag: Html is better than plain text for modeling retrieved knowledge in rag systems](#). *arXiv preprint arXiv:2411.02959*.
- Jiejun Tan, Zhicheng Dou, Yutao Zhu, Peidong Guo, Kun Fang, and Ji-Rong Wen. 2024b. [Small models, big insights: Leveraging slim proxy models to decide when and what to retrieve for llms](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, ACL 2024, Bangkok, Thailand, August 11-16, 2024, pages 4420–4436. Association for Computational Linguistics.
- Yubao Tang, Ruqing Zhang, Jiafeng Guo, Maarten de Rijke, Wei Chen, and Xueqi Cheng. 2024. List-wise generative retrieval models via a sequential learning process. *ACM Transactions on Information Systems*.
- Yi Tay, Vinh Tran, Mostafa Dehghani, Jianmo Ni, Dara Bahri, Harsh Mehta, Zhen Qin, Kai Hui, Zhe Zhao, Jai Prakash Gupta, Tal Schuster, William W. Cohen, and Donald Metzler. 2022. [Transformer memory as a differentiable search index](#). In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*.
- Harsh Trivedi, Niranjan Balasubramanian, Tushar Khot, and Ashish Sabharwal. 2023. [Interleaving retrieval with chain-of-thought reasoning for knowledge-intensive multi-step questions](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, ACL 2023, Toronto, Canada, July 9-14, 2023, pages 10014–10037. Association for Computational Linguistics.
- Liang Wang, Nan Yang, Xiaolong Huang, Binxing Jiao, Linjun Yang, Daxin Jiang, Rangan Majumder, and Furu Wei. 2022a. [Text embeddings by weakly-supervised contrastive pre-training](#). *CoRR*, abs/2212.03533.
- Yujing Wang, Yingyan Hou, Haonan Wang, Ziming Miao, Shibin Wu, Qi Chen, Yuqing Xia, Chengmin Chi, Guoshuai Zhao, Zheng Liu, Xing Xie, Hao Sun, Weiwei Deng, Qi Zhang, and Mao Yang. 2022b. [A neural corpus indexer for document retrieval](#). In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*.
- Heming Xia, Tao Ge, Peiyi Wang, Si-Qing Chen, Furu Wei, and Zhifang Sui. 2023. [Speculative decoding: Exploiting speculative execution for accelerating seq2seq generation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023, Singapore, December 6-10, 2023*, pages 3909–3925. Association for Computational Linguistics.
- Shitao Xiao, Zheng Liu, Peitian Zhang, Niklas Muenighoff, Defu Lian, and Jian-Yun Nie. 2024. [C-pack: Packed resources for general chinese embeddings](#). In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2024, Washington DC, USA, July 14-18, 2024*, pages 641–649. ACM.
- An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, Guanting Dong, Haoran Wei, Huan Lin, Jialong Tang, Jialin Wang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Ma, Jianxin Yang, Jin Xu, Jingren Zhou, Jinze Bai, Jinzheng He, Junyang Lin, Kai Dang, Keming Lu, Keqin Chen, Kexin Yang, Mei Li, Mingfeng Xue, Na Ni, Pei Zhang, Peng Wang, Ru Peng, Rui Men, Ruize Gao, Runji Lin, Shijie Wang, Shuai Bai, Sinan Tan, Tianhang Zhu, Tianhao Li, Tianyu Liu, Wenbin Ge, Xiaodong Deng, Xiaohuan Zhou, Xingzhang Ren, Xinyu Zhang, Xipin Wei, Xuancheng Ren, Xuejing Liu, Yang Fan, Yang Yao, Yichang Zhang, Yu Wan, Yunfei Chu, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, Zhifang Guo, and Zhihao Fan. 2024. [Qwen2 technical report](#). *CoRR*, abs/2407.10671.
- Tianchi Yang, Minghui Song, Zihan Zhang, Haizhen Huang, Weiwei Deng, Feng Sun, and Qi Zhang. 2023. [Auto search indexer for end-to-end document](#)

- retrieval. In *Findings of the Association for Computational Linguistics: EMNLP 2023, Singapore, December 6-10, 2023*, pages 6955–6970. Association for Computational Linguistics.
- 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 *EMNLP*, pages 2369–2380, Brussels, Belgium. Association for Computational Linguistics.
- Hansi Zeng, Chen Luo, Bowen Jin, Sheikh Muhammad Sarwar, Tianxin Wei, and Hamed Zamani. 2023. [Scalable and effective generative information retrieval](#). *CoRR*, abs/2311.09134.
- Hansi Zeng, Chen Luo, and Hamed Zamani. 2024. Planning ahead in generative retrieval: Guiding autoregressive generation through simultaneous decoding. *arXiv preprint arXiv:2404.14600*.
- Jintian Zhang, Cheng Peng, Mengshu Sun, Xiang Chen, Lei Liang, Zhiqiang Zhang, Jun Zhou, Huajun Chen, and Ningyu Zhang. 2024a. [Onegen: Efficient one-pass unified generation and retrieval for llms](#). *Preprint*, arXiv:2409.05152.
- Peitian Zhang, Zheng Liu, Yujia Zhou, Zhicheng Dou, and Zhao Cao. 2023. [Term-sets can be strong document identifiers for auto-regressive search engines](#). *CoRR*, abs/2305.13859.
- Xuanwang Zhang, Yunze Song, Yidong Wang, Shuyun Tang, Xinfeng Li, Zhengran Zeng, Zhen Wu, Wei Ye, Wenyuan Xu, Yue Zhang, Xinyu Dai, Shikun Zhang, and Qingsong Wen. 2024b. [RAGLAB: A modular and research-oriented unified framework for retrieval-augmented generation](#). *CoRR*, abs/2408.11381.
- Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, Yifan Du, Chen Yang, Yushuo Chen, Zhipeng Chen, Jinhao Jiang, Ruiyang Ren, Yifan Li, Xinyu Tang, Zikang Liu, Peiyu Liu, Jian-Yun Nie, and Ji-Rong Wen. 2023. [A survey of large language models](#). *CoRR*, abs/2303.18223.
- Yutao Zhu, Zhaoheng Huang, Zhicheng Dou, and Ji-Rong Wen. 2024. [One token can help! learning scalable and pluggable virtual tokens for retrieval-augmented large language models](#). *CoRR*, abs/2405.19670.
- Yutao Zhu, Huaying Yuan, Shuting Wang, Jiongnan Liu, Wenhan Liu, Chenlong Deng, Zhicheng Dou, and Ji-Rong Wen. 2023. [Large language models for information retrieval: A survey](#). *CoRR*, abs/2308.07107.
- Shengyao Zhuang, Houxing Ren, Linjun Shou, Jian Pei, Ming Gong, Guido Zuccon, and Daxin Jiang. 2022. [Bridging the gap between indexing and retrieval for differentiable search index with query generation](#). *CoRR*, abs/2206.10128.

Appendix

A	The FM-Index	15
A.1	Data Structure	15
A.2	Supporting Functions	15
A.3	Examples	16
B	False Pruning in Constrained Decoding	16
B.1	What is False Pruning?	16
B.2	What Causes False Pruning?	17
B.3	How to Mitigate False Pruning?	17
C	Datasets	17
C.1	Details of Datasets	17
C.2	Statistics	17
D	Implementation Details	18
D.1	Implementation Details for Baselines	18
D.2	Implementation Details for Naive Constrained Generation	18
D.3	Implementation Details for RetroLLM	18
D.3.1	FM-Index Building	18
D.3.2	Training	18
D.3.3	Inference	19
E	Additional Experimental Results	19
E.1	Analysis of Retrieval Performance	20
E.2	Impact of Different Base Models	20
E.3	Impact of Generated Evidence Quantity	21
E.4	Analysis of Pre-processing Efficiency	21
E.5	Performance Comparison with Compression-based RAG Methods	21
F	Case Study	21
F.1	Examples from RetroLLM	21
F.2	Comparing RetroLLM with Naive Beam Search Method	22

A The FM-Index

The FM-Index (Ferragina and Manzini, 2000), which stands for Full-text index in Minute space, is a space-efficient data structure designed for indexing large text corpora, combining the Burrows-Wheeler Transform (BWT) and run-length encoding. It enables fast substring searching while providing substantial compression, making it particularly useful in applications such as prefix-constrained decoding.

A.1 Data Structure

The FM-Index is based on the Burrows-Wheeler Transform (BWT) (Schindler, 1997). The BWT of a string S is computed by sorting all cyclic rotations of S lexicographically and then taking the last column of the sorted rotations. This transformation rearranges the characters of the string in a way that enhances its compressibility, which is key to the FM-Index's space efficiency.

Formally, for a string $S = S_1S_2 \dots S_n$, the BWT, denoted as $\text{BWT}(S)$, is obtained by sorting the cyclic rotations of S lexicographically and taking the last character of each rotation. Let $\mathcal{R}(S)$ denote the set of all cyclic rotations of S , sorted in lexicographical order:

$$\mathcal{R}(S) = \{\sigma_1, \sigma_2, \dots, \sigma_n\} \quad (15)$$

where σ_i denotes the i -th rotation of S . The BWT of S is then the string formed by the last characters of these sorted rotations:

$$\text{BWT}(S) = (\sigma_1[n], \sigma_2[n], \dots, \sigma_n[n]) \quad (16)$$

The FM-Index stores only two columns from the BWT matrix: the first (F) and last (L) columns. These columns capture the relative order of the characters in all cyclic rotations of S . More precisely: - The first column (F) contains the sorted characters of the text S . - The last column (L) contains the last character of each of the cyclic rotations of S .

Additionally, to enable efficient searching, the FM-Index uses additional data structures, such as the Wavelet Tree, to store the L column efficiently. This allows for fast rank and select operations, which are essential for searching.

A.2 Supporting Functions

Backward Search The core feature of the FM-Index is the backward search, which allows for

efficient substring searching. Given a substring $P = p_1p_2 \dots p_k$, the backward search locates all occurrences of P in the original string by iteratively searching through the first (F) and last (L) columns.

The backward search proceeds by iterating through each character of P from right to left. Initially, the search interval spans the entire text S , represented as the BWT matrix. In each step, we examine the current character p_i of the pattern and update the search interval by examining the corresponding entries in the F and L columns. Specifically: (1). We find all occurrences of p_i in the last column (L). (2). We update the search range in the first column (F) to include only the rows corresponding to the occurrences of p_i .

This process is repeated for each character of the pattern, refining the search interval until all occurrences of the substring are found.

The time complexity of the backward search is $O(k \log |V|)$, where k is the length of the pattern P and $|V|$ is the size of the alphabet. This is because each iteration of the search takes $O(\log |V|)$ time, and there are k iterations corresponding to the length of the pattern.

Count The occurrence count function, denoted as $\text{occ}(P)$, counts how many times a pattern P occurs in the original text. The occurrence count is closely related to the backward search. After performing a backward search for a pattern P , the occurrence count is simply the size of the resulting search interval. This interval represents all occurrences of P in the text.

Since computing the occurrence count requires performing a backward search, the time complexity of this operation is also $O(k \log |V|)$, where k is the length of the pattern and $|V|$ is the size of the alphabet.

Locate The locate function, denoted as $\text{locate}(P)$, returns the positions in the original text where the pattern P occurs. This function works by using the results of the backward search to determine the positions of the occurrences. Specifically, once the search interval is determined through backward search, the locate function maps the rows of the interval back to positions in the original text. This mapping is achieved by using the F column, which contains the sorted characters of the text.

The time complexity of the locate function is $O(k \log |V|)$, as it involves performing a backward search for the pattern P .

Extract Given the start and end positions of a substring in the BWT matrix, the extract text function reconstructs the corresponding text substring. The process works by tracing the positions of the characters in the substring through the F and L columns in reverse order. Starting from the end position, the algorithm follows the reverse of the backward search to find the characters in the substring and reconstruct the text.

This operation runs in $O(k)$, where k is the length of the substring to be extracted. The reason for this is that we need to perform k steps to extract a substring of length k , with each step involving simple lookups in the F and L columns.

A.3 Examples

To better understand how the FM-Index works, let us examine a concrete example using the string $S = \text{"banana\$"}$, where $\$$ serves as the end-of-string marker. The first step in constructing the FM-Index is to generate the Burrows-Wheeler Transform. This is accomplished by creating all possible cyclic rotations of the input string and sorting them lexicographically. For our example string, the sorted rotations form a matrix where each row represents one rotation:

Sorted Rotations	F Col.	L Col.
\$banana	\$	a
a\$banan	a	n
ana\$ban	a	n
anana\$b	a	b
banana\$	b	\$
na\$bana	n	a
nana\$ba	n	a

The Burrows-Wheeler Transform of S is then obtained as the last column L: $\text{BWT}(S) = \text{"annb\$aa"}$. The FM-Index maintains this last column L along with the first column F = $\text{"$aaabnn"}$, which contains the lexicographically sorted characters of S . These two columns, combined with auxiliary data structures for efficient rank and select operations, form the core of the FM-Index.

Given a pattern such as $P = \text{"ana"}$, we can perform a backward search from right to left. Starting with the last character 'a', we determine its occurrence range in the F column using the cumulative count $C[a] = 1$ (since only '\$' precedes 'a'). The rank of 'a' up to position 7 in L is 3, updating the search interval to [2, 4]. Next, processing the character 'n', with $C[n] = 5$, we find that the rank of

'n' up to position 4 in L is 2, refining the search interval to [6, 7]. Finally, processing the first character 'a', we use $C[a] = 1$ and find that the rank of 'a' up to position 7 in L is 3, resulting in a final search interval of [3, 4].

Upon identifying the final search interval [3, 4] in the L column, we examine the corresponding characters, which are 'n' and 'b'. These characters represent the ones that precede the pattern "ana" in the original string S . Mapping these to the positions in the F column reveals that the possible characters following "ana" in S are 'n' and '\$'. Specifically, in "banana\$", the substring "ana" is followed by 'n' in the first occurrence and by '\$' in the second occurrence. Therefore, the backward search correctly identifies 'n' and '\$' as the allowable next characters after the prefix "ana", validating the FM-Index derivation process.

B False Pruning in Constrained Decoding

In Section 2.2, we conducted empirical studies revealing that false pruning is a significant issue in constrained decoding. This section delves deeper into understanding this problem.

B.1 What is False Pruning?

False pruning occurs when the search process incorrectly eliminates branches that could contain the optimal solution, preventing the algorithm from identifying the true best outcome. Specifically, in prefix-constrained decoding for language models, false pruning involves incorrectly discarding candidate tokens that meet the prefix constraint but might contribute to the optimal solution (Zhang et al., 2023).

Consider the question: "What is the capital of France?" The correct evidence in the corpus is "The capital city of France is Paris." During decoding with the prefix constraint "The capital", if the model selects "of" instead of "city", a critical issue arises. Suppose the corpus lacks direct statements like "The capital of France" but includes irrelevant examples such as "The capital of America". In this scenario, even though the path starting with "The capital of" could potentially lead to the correct answer about Paris, the model is constrained by the corpus to decode only irrelevant evidence like "The capital of America is Washington D.C."

This situation exemplifies false pruning: the correct solution path is erroneously removed during

beam search or sampling, despite being valid under the prefix constraint. This happens because intermediate token choices steer the model toward contexts where it cannot effectively retrieve the target information about France’s capital. Such failures illustrate how the local, token-by-token nature of decoding can clash with prefix constraints, causing the model to follow suboptimal paths and ultimately fail to generate the correct evidence.

B.2 What Causes False Pruning?

For auto-regressive decoding models, false pruning arises primarily due to two factors:

Excessive Prefix Choices: In large corpora, candidate sequences present a vast number of prefix options initially. The model can generate nearly any short prefix it wants, making it challenging to predict the correct one.

Limited Future Awareness: Even with fewer prefix choices, the model cannot anticipate future content beyond the current token decision. This limitation makes it difficult to select tokens that lead to the correct evidence.

B.3 How to Mitigate False Pruning?

Addressing the root causes of false pruning involves implementing strategies that either narrow the prefix choices or enhance the model’s foresight during decoding.

Reducing Prefix Choices: One effective method is to limit the number of prefix options. Our approach employs clue generation to identify a relevant subset of documents, followed by decoding evidence within this constrained set. This reduction significantly decreases the prefix choices, mitigating the risk of false pruning.

Enhancing Future Relevance Awareness: Another strategy is to provide the model with information about the relevance of future sequences. In our method, we identify the clue’s position within the document and utilize the surrounding text as future windows. By guiding the language model to generate relevant evidence based on these windows and their relevance scores, we improve the model’s ability to connect to the target information.

Set-Based Decoding: Some generative retrieval methods adopt set-based decoding strategies (Zhang et al., 2023; Zeng et al., 2024), which bypass the issues inherent in auto-regressive decoding by directly generating sets of terms. These methods are suitable for retrieval tasks that involve

Table 5: Detailed statistics of datasets and retrieval corpus utilized in our experiments.

Task	Dataset	# Train	# Test
Single-hop QA	NQ	79,168	3,610
Single-hop QA	TriviaQA	78,785	11,313
Single-hop QA	PopQA	/	14,267
Multi-hop QA	HotpotQA	90,447	7,405
Multi-hop QA	2WIKI	/	12,576
Retrieval Corpus		# Passages	# Documents
Wikipedia		21,015,324	3,232,907

decoding document identifiers to fetch corresponding documents. However, they are not applicable to our evidence generation task, where the goal is to generate meaningful evidence directly rather than retrieve identifiers.

C Datasets

C.1 Details of Datasets

In our experiments, we utilize a variety of question answering (QA) datasets to evaluate both single-hop and multi-hop reasoning capabilities. For single-hop QA, we employ the Natural Questions (NQ) (Kwiatkowski et al., 2019) dataset, TriviaQA (Joshi et al., 2017), and PopQA (Mallen et al., 2023), which provide a diverse range of factual questions requiring straightforward retrieval and answer extraction. For multi-hop QA, we use HotpotQA (Yang et al., 2018), which necessitates reasoning across multiple documents, and 2WIKI (Ho et al., 2020), a dataset designed to test more complex multi-hop reasoning scenarios. These datasets are selected to cover a broad spectrum of QA challenges, ensuring a comprehensive evaluation of the model’s retrieval and reasoning capability.

C.2 Statistics

Table 5 presents detailed statistics of the datasets and the retrieval corpus used in our study. For single-hop QA tasks, NQ consists of 79,168 training samples and 3,610 test samples, while TriviaQA has 78,785 training samples and 11,313 test samples. PopQA is used solely for testing, with 14,267 samples. In the multi-hop QA category, HotpotQA includes 90,447 training samples and 7,405 test samples, and 2WIKI provides 12,576 test samples without a training set. The retrieval corpus comprises the Wikipedia dataset, containing 21,015,324 passages and 3,232,907 documents. These statistics highlight the extensive scale of our

experimental setup, facilitating robust training and evaluation of our models.

D Implementation Details

D.1 Implementation Details for Baselines

All RAG baselines are implemented based on the FlashRAG framework, which is an open-source retrieval-augmented generation toolkit (Jin et al., 2024a). For Self-RAG (Asai et al., 2024), we use the trained selfrag-llama2-7B checkpoint. For all other baselines, we use Mistral-7B-Instruct as the backbone model, aligning with our RetroLLM. All hyper-parameter configurations are set to default in FlashRAG.

D.2 Implementation Details for Naive Constrained Generation

For the naive approach to constrained beam evidence generation, we set `num_beams` = 5, `num_beam_groups` = 5, and `diversity_penalty` = 1.0 for constrained beam search. The `num_beam_groups` and `diversity_penalty` parameters are crucial; without setting these two parameters, the sequences generated by each beam would be highly similar, leading to a significant decrease in evidence accuracy. These parameters ensure diversity among the multiple generated sequences and sort the `beam_size` generated evidences from high to low according to the generation probability of the language model, so that more relevant evidence can be ranked ahead.

For cases where an answer needs to be generated, we continue to freely generate the answer without constraint after each beam, and the final answer given is the answer generated after the top-1 sequence. The specific input and output formats are shown in F.

D.3 Implementation Details for RetroLLM

The implementation of RetroLLM mainly includes FM-Index building, training, and inference. All experiments are conducted on 8 NVIDIA A800-80GB GPUs and an Intel(R) Xeon(R) Platinum 8358 CPU @ 2.60GHz with 64 cores.

D.3.1 FM-Index Building

We implement the FM-Index data structure based on the SDSL-lite (Succinct Data Structure Library) framework (Gog et al., 2014), which is an efficient C++ template library specifically designed for implementing compressed data structures. We then

implemented the functionalities used in this paper at the C++ level, including prefix locating, finding allowed next tokens, counting occurrences, extracting sequences, etc. We also built and stored an FM-Index Manager on the C++ side to map given DocIDs to their corresponding document FM-Indexes. To allow Python code to call these C++ implementations, we used the SWIG (Simplified Wrapper and Interface Generator) tool (Beazley, 1996).

D.3.2 Training

Training Data Construction The data construction approach simulates the model’s inference process. For each labeled QA pair (q, a) , we first utilize a sparse retriever SPLADE-v3 (Lassance et al., 2024) to obtain top-8 clues \mathcal{C}_{exp} and retrieve top-20 documents. We then locate the sentences containing these clues within the documents, followed by employing a reranker to obtain the top- k_e relevant evidences \mathcal{E}_{rel} , where k_e is set to 5 for single-hop QA and 10 for multi-hop QA tasks. Next, we examine whether the labeled answer a is contained within each evidence e to determine if the evidence can address the original query q . To further ensure labeling accuracy, we employ a Llama3.1-70B-Instruct (Dubey et al., 2024) model to judge whether each $e \in \mathcal{E}_{\text{rel}}$ can genuinely answer the query q . We consider an evidence e relevant only if it both contains a and is labeled as relevant by the LLM. Subsequently, we select the top- $k \leq k_e$ evidences where the k -th evidence is the first relevant e . For target clues, we utilize Llama3.1-70B-Instruct to extract key entities from the query and relevant evidences to construct target clues \mathcal{C}_{gen} . This process yields the training pair $(q, \mathcal{C}_{\text{gen}}, \mathcal{E}, a)$, with the target format illustrated in Figure 3 and Appendix F.

Model Optimization As described in Section 3.5, we use a standard sequence-to-sequence loss to train the model. For efficient model training, we employ LoRA (Hu et al., 2022), setting `lora_r` to 16 and `lora_alpha` to 64. We set training epochs to 3 and γ to 2. Since evidence generation is performed under constraints, most of the middle tokens in evidence generation have limited choices under the constraints of the FM-Index; the crucial parts are the first few tokens at the beginning of the evidence and the tokens at the end that decide to finish the evidence. Therefore, we set the middle 80% tokens of each evidence not to participate in training, so that the model training focuses

Table 6: Detailed retrieval performance on five open-domain QA datasets, comparing sparse, dense, and generative approaches. The best results are highlighted in **Bold**.

Method	In-domain Datasets									Out-of-domain Datasets					
	NQ			TriviaQA			HotpotQA			PopQA			2WIKI		
	R@1	R@5	Num	R@1	R@5	Num	R@1	R@5	Num	R@1	R@5	Num	R@1	R@5	Num
<i>Sparse Retrieval</i>															
BM25	24.1	46.2	5	49.6	68.5	5	31.2	48.7	5	39.6	54.3	5	22.6	37.5	5
SPLADE-v3	45.4	68.0	5	58.8	75.9	5	32.9	45.3	5	47.6	65.2	5	22.2	40.6	5
<i>Dense Retrieval</i>															
E5	55.7	77.3	5	61.6	77.8	5	32.3	52.0	5	51.7	70.9	5	21.6	39.8	5
BGE	50.3	73.6	5	58.7	75.1	5	33.7	54.7	5	50.8	69.6	5	21.1	38.9	5
<i>Generative Retrieval</i>															
Naive Constrain	13.1	26.9	5	23.0	46.9	5	11.8	21.6	5	10.9	21.2	5	9.4	19.0	5
RetroLLM	51.6	62.5	3.20	61.1	71.0	2.80	35.6	57.3	3.86	57.0	70.1	4.07	23.0	41.8	4.40

Table 7: Detailed performance comparison of RetroLLM using various base models, including the Llama3 series, Qwen-2.5 series, and Mistral series, with parameter sizes ranging from 1B to 14B. All base models we used are the instruction-tuned versions. The best results are highlighted in **Bold**.

Base Model	In-domain Datasets									Out-of-domain Datasets					
	NQ			TriviaQA			HotpotQA			PopQA			2WIKI		
	Acc	F1	Tok	Acc	F1	Tok	Acc	F1	Tok	Acc	F1	Tok	Acc	F1	Tok
<i>Llama3 Series</i>															
Llama3.2-1B	54.4	35.8	260	64.4	52.9	288	58.8	33.5	573	63.3	32.9	344	44.5	28.5	583
Llama3.2-3B	58.9	45.4	278	67.8	62.1	267	61.3	37.8	609	64.7	40.4	338	47.3	32.2	632
Llama3-8B	59.2	46.4	306	72.7	69.3	256	62.2	47.4	575	65.2	41.4	338	48.7	36.1	668
<i>Qwen2.5 Series</i>															
Qwen2.5-1.5B	50.1	34.3	200	57.2	51.2	170	57.0	32.6	539	59.5	32.6	286	47.5	26.3	650
Qwen2.5-3B	52.1	36.8	236	61.4	56.3	212	60.6	34.1	628	64.0	34.8	336	48.1	30.6	694
Qwen2.5-7B	54.9	42.3	230	64.5	62.4	196	61.9	42.0	549	62.8	37.1	313	48.7	32.5	634
Qwen2.5-14B	58.6	50.6	225	72.8	69.5	186	62.6	45.9	568	64.3	40.8	343	51.3	36.9	687
<i>Mistral Series</i>															
Mistral-7B	61.6	49.8	302	74.3	72.8	287	61.9	47.2	607	65.7	43.0	355	48.9	36.2	661

more on the key parts. Since we added special tokens to represent the start, separation, and end operations of clue and evidence generation, in the model parameters trained, besides the parameters trained by LoRA, we also added the embeddings corresponding to the new tokens to effectively learn the generation of new tokens.

D.3.3 Inference

As illustrated in Figure 3, RetroLLM includes the following three stages. In the clue generation stage, RetroLLM first generates clues with corpus-level FM-Index constraints. The format of this part is “<|clue|> c_1 <|sep|> c_2 <|sep|> ... <|/clue|>”. During clue generation, we simultaneously expand clues with the sparse lexical and expansion model SPLADE-v3 (Lassance et al., 2024). We set the number of expanded clues to 8 and the maximum

number of generated clues to 5.

In the evidence generation stage, evidence is generated based on document-level FM-Index constraints and future window relevance. The format of this part is “<|evidence|> e_1 <|sep|> e_2 <|sep|> ... <|/evidence|>”. We limit the maximum number of generated evidence for single-hop and multi-hop QA to 5 and 10, respectively. We set w_1 and w_2 to 1 and 2, respectively, and λ to 100.

In the answer generation stage, no constraints are added during the final answer generation.

E Additional Experimental Results

This section presents detailed experimental results and analysis, including retrieval performance, the impact of RetroLLM performance with different base models and generated evidence quantity.

Table 8: Detailed performance with different numbers of generated evidence.

# Num	In-domain Datasets				Out-of-domain Datasets					
	NQ		TriviaQA		HotpotQA		PopQA		2WIKI	
	Acc	F1	Acc	F1	Acc	F1	Acc	F1	Acc	F1
1	42.2	40.5	59.3	61.6	50.6	44.2	43.9	40.9	35.1	31.3
2	50.6	42.3	66.3	65.9	59.8	43.8	52.8	45.9	39.8	34.6
3	54.4	42.5	69.3	67.2	61.9	43.0	55.7	45.5	42.1	34.5
4	56.7	43.1	70.9	67.6	64.6	41.0	57.7	45.7	43.9	34.8
5	61.5	49.4	74.6	72.9	66.8	43.0	59.4	46.8	45.9	36.6
6	61.7	49.5	74.6	73.0	67.4	42.8	60.1	47.1	47.9	37.1
7	61.7	49.5	74.6	72.9	67.6	42.5	60.8	47.0	48.4	36.5
8	61.7	49.5	74.6	72.9	68.0	42.7	61.2	46.9	48.6	37.2
9	61.7	49.5	74.6	72.9	68.0	42.7	61.6	47.1	48.7	37.0
10	61.7	49.5	74.6	72.9	68.5	42.7	61.9	47.1	48.9	36.2

E.1 Analysis of Retrieval Performance

We analyze the retrieval performance of RetroLLM compared to sparse and dense retrieval baselines, as discussed in Section 4.4. The results are shown in Table 6

(1) For single-hop QA tasks, RetroLLM demonstrates superior accuracy on R@1, thanks to the design of clues and future windows, which help precisely locate the relevant evidence. For instance, on the PopQA dataset, RetroLLM achieves an R@1 of **57.0%**, surpassing the best dense retriever E5, which attains **51.7%**. Additionally, RetroLLM uses fewer passages on average (**2.80** for TriviaQA and **3.20** for NQ) compared to the fixed number of 5 in baseline methods, indicating more efficient retrieval.

(2) For multi-hop QA tasks, RetroLLM shows superior accuracy compared to all other methods for both R@1 and R@5, while utilizing a smaller average number of retrieved passages. Specifically, on HotpotQA, RetroLLM achieves an R@1 of **35.6%**, outperforming E5’s **32.3%** and SPLADE-v3’s **32.9%**. On the 2WIKI dataset, RetroLLM attains an R@1 of **23.0%**, higher than E5’s **21.6%**, demonstrating its effectiveness in multi-hop retrieval scenarios while using only **4.40** passages on average versus the baseline’s 5.

(3) Notably, the naive generative retrieval method using constrained beam search performs poorly on all metrics, further validating the severity of false pruning, as discussed in Section 2.2. For example, on the NQ dataset, the naive method achieves an R@1 of only **13.1%**, significantly lower than RetroLLM’s **51.6%**. Similarly, on TriviaQA, it attains an R@1 of **23.0%** compared to RetroLLM’s **61.1%**, highlighting the substantial

performance gap and the advantages of our approach.

E.2 Impact of Different Base Models

To evaluate the performance of RetroLLM using different backbone LLMs with varying parameter sizes, we conducted experiments using the Mistral, Llama3, and Qwen2.5 series, with parameters ranging from 1B to 14B, as discussed in Section 4.4. The results are shown in Figure 7. We observe that:

(1) **As the parameter size increases, RetroLLM’s performance steadily improves, aligning with the scaling law.** For example, within the Llama3 series, the accuracy on the NQ dataset rises from 54.4% for Llama3.2-1B to 59.2% for Llama3-8B. Similarly, the F1 score on TriviaQA improves from 52.9% to 69.3%. In the Qwen2.5 series, the accuracy on NQ increases from 50.1% for Qwen2.5-1.5B to 58.6% for Qwen2.5-14B, and the F1 score climbs from 34.3% to 50.6%. This consistent enhancement across different model sizes indicates that larger base models contribute to better retrieval performance in RetroLLM.

(2) **There are slight performance differences across the different models (Mistral, Llama3, Qwen2.5), with Mistral generally outperforming Llama3, which in turn outperforms Qwen2.5.** Specifically, Mistral-7B achieves the highest accuracy on several datasets, such as **61.6%** on NQ and **74.3%** on TriviaQA, surpassing both Llama3-8B and Qwen2.5-14B. On the PopQA dataset, Mistral-7B attains an accuracy of **65.7%**, compared to 65.2% for Llama3-8B and 64.3% for Qwen2.5-14B. Despite these variations, all models confirm the effectiveness of RetroLLM, as even smaller

Table 9: Comparison of pre-processing efficiency between RetroLLM and dense retriever E5-Base.

Method	Index Method	Total Storage	Building Time
Plain Text	-	12.7 GB	-
E5-Base	Dense Emb.	60.1 GB	7.57 h
RetroLLM	FM-Index	20.7 GB	5.39 h

models like Qwen2.5-1.5B achieve reasonable performance (e.g., 50.1% accuracy on NQ and 57.2% on TriviaQA), demonstrating that RetroLLM is robust across different base models and parameter sizes.

E.3 Impact of Generated Evidence Quantity

Since RetroLLM can dynamically determine the number of evidence to retrieve, we investigated the effect of different maximum retrieval quantities on performance, as discussed in Section 4.4. The results are shown in Table 8. We observe that:

(1) **When retrieving up to 1-5 evidence, performance continues to improve as the number of retrieved pieces increases, suggesting that more evidence contributes to stronger performance in these tasks.** For instance, on the NQ dataset, the accuracy improves from 42.2% when retrieving only one piece of evidence to 61.5% with five pieces. Similarly, the accuracy on TriviaQA rises from 59.3% to 74.6% as the number increases from one to five. This trend indicates that accessing more evidence enables RetroLLM to retrieve relevant information more effectively, enhancing answer accuracy.

(2) **However, for multi-hop QA, performance stabilizes around 6 evidence, as more evidence can bring in both useful and distracting information, thereby limiting further performance gains.** Specifically, on the HotpotQA dataset, the accuracy increases from 50.6% with one piece of evidence to 67.4% with six pieces, but additional evidence beyond this point yields diminishing returns (e.g., 68.5% accuracy at ten pieces). This suggests that while some additional evidence is beneficial, too much can introduce noise that counteracts the benefits, highlighting the importance of a balanced retrieval strategy.

E.4 Analysis of Pre-processing Efficiency

We conducted experiments to compare the pre-processing efficiency between RetroLLM and E5-Base, including index building time and total stor-

Table 10: Performance comparison with compression-based RAG methods, reporting Accuracy metric.

Method	NQ	TQA	Hotpot	PopQA	2WIKI
LongLLMLingua	50.8	71.1	38.5	57.7	31.9
RECOMP-abs	51.8	68.2	37.8	64.2	32.2
Selective-Context	49.0	67.4	34.6	51.0	25.8
RetroLLM	61.6	74.3	61.9	65.7	48.9

age. For E5-Base, we used an NVIDIA A800-80GB GPU, while for RetroLLM, we used an Intel(R) Xeon(R) Platinum 1304 8358 CPU. The experimental results are as follows:

As shown in the results in Table 9, **RetroLLM’s pre-processing is more efficient than E5-Base in both index building time and total storage**, demonstrating RetroLLM’s superior pre-processing efficiency. Moreover, building FM-Index does not require GPU resources and can be efficiently constructed using CPUs only.

E.5 Performance Comparison with Compression-based RAG Methods

We compared the performance of RetroLLM with several common compression methods, including LongLLMLingua, RECOMP-abstractive, and Selective-Context. Our implementation is based on the FlashRAG toolkit, with consistent parameter settings, and uses Wikipedia as the corpus, the same as RetroLLM. The experimental results are as follows:

As shown in Table 10, **RetroLLM significantly outperforms other compression methods, indicating that RetroLLM is more effective than methods that first retrieve long contexts before compression**, thanks to the hierarchical FM-Index, clue prediction, and forward-looking constrained evidence generation mechanisms.

F Case Study

This section presents examples from RetroLLM and compares them with outputs from a naive constrained beam search method. These examples illustrate the detailed workings of our method and highlight the shortcomings of the naive approach.

F.1 Examples from RetroLLM

Tables 11 and 12 display examples from single-hop and multi-hop question-answering (QA) datasets, respectively. The overall process of RetroLLM consists of two main stages: *clue generation* and *evidence generation*.

In the **clue generation** stage, RetroLLM identifies key terms or phrases related to the question, encapsulated within the `<clue>` and `</clue>` tokens, separated by `<sepl>` token. These clues serve to focus the model’s attention on pertinent concepts or entities that are crucial for answering the question. For instance, in Example #1 from the NQ Dataset (Table 11), the clue generated is “The Star,” which directly relates to the movie mentioned in the question.

Following clue generation, the **evidence generation** stage involves retrieving relevant information from the knowledge corpus, guided by the previously identified clues. The retrieved evidence is enclosed within the `<evidence>` and `</evidence>` tokens. This evidence often comprises multiple snippets of information, separated by `<sepl>` token, which collectively support the final answer. For example, in Example #1 from the HotpotQA Dataset (Table 12), the model generates evidence about “Old and in the Way” and “Owsley Stanley,” providing a coherent chain of information that leads to the correct answer.

F.2 Comparing RetroLLM with Naive Beam Search Method

Tables 13 and 14 compare the outputs of the naive constrained beam search method with those of RetroLLM for a question from the NQ Dataset. The naive method attempts to generate evidence under corpus-level FM-Index constraints, but this approach leads to several issues.

The beams generated by the naive method contain evidence that is largely irrelevant or incoherent. Although some initial phrases may appear related to the question, the continuation often deviates significantly, producing sentences that do not contribute to answering the question correctly. For instance, the naive method incorrectly identifies “Roger Maris,” “1903,” and “Jonathan Elliot” as answers to the question “who got the first nobel prize in physics?” These incorrect answers result from the model’s inability to maintain topic coherence under the stringent corpus-level constraints, a phenomenon known as the *false pruning problem* discussed in Section 2.2 and Appendix B.

In contrast, RetroLLM addresses this issue by first generating clues to narrow down the relevant document subset and then performing forward-looking constrained evidence generation within this subset. This method reduces false pruning and enhances the accuracy of evidence retrieval.

Table 11: Examples from RetroLLM on single-hop QA datasets, with special tokens used in the clue and evidence generation stages highlighted in orange box and blue box, respectively. All generated evidence is from Wikipedia.

Example #1 from NQ Dataset
<p>Question: when does the movie the star come out?</p> <p>Labeled Answer: ["November 17, 2017"]</p>
Results by RetroLLM
<p>Model Input: Question: when does the movie the star come out? Your Response:</p>
<p>Model Output: <clue> The Star </clue> <evidence> the film, then titled "The Lamb", was given an official release date of December 8, 2017. On June 20, 2016, it was announced that Brian Henson and Lisa Henson from the Jim Henson Company would be executive producers for the film, now titled "The Star", making it the first Henson-produced film for Sony since "The Adventures of Elmo in Grouchland" in 1999. <sepl> The film's title was changed from the working title "Dog Years", in a mutual decision between the filmmakers and A24 (who made the suggestion). It was released on February 22, 2018, through DirecTV Cinema before being released in a limited release and through video on demand on March 30, 2018. <sepl> Chenoweth, Tracy Morgan, Tyler Perry, and Oprah Winfrey. The film was released by Columbia Pictures through its parental label Sony Pictures Entertainment on November 17, 2017 in the United States. The film received mixed reviews and grossed over \$62 million worldwide. </evidence> The answer is November 17, 2017.</p>
Example #2 from TriviaQA Dataset
<p>Question: Who was the man behind The Chipmunks?</p> <p>Labeled Answer: ["David Seville"]</p>
Results by RetroLLM
<p>Model Input: Question: Who was the man behind The Chipmunks? Your Response:</p>
<p>Model Output: <clue> The Chipmunks </clue> <evidence> 1972), better known by the stage name David Seville, was an American singer-songwriter and actor, best known for creating the novelty musical group The Chipmunks (now better known as "Alvin and the Chipmunks"). Initially a stage and film actor, he rose to prominence in 1958 with the songs "Witch Doctor" and "The Chipmunk Song (Christmas Don't Be Late)", which both became "Billboard" number-one singles. </evidence> The answer is David Seville.</p>
Example #3 from PopQA Dataset
<p>Question: What is Carsten Carlsen's occupation?</p> <p>Labeled Answer: ["pianist", "composer"]</p>
Results by RetroLLM
<p>Model Input: Question: What is Carsten Carlsen's occupation? Your Response:</p>
<p>Model Output: <clue> Carlsen </clue> <evidence> Carsten Carlsen (5 June 1892 – 28 August 1961) was a Norwegian pianist and composer. Carsten Marensius Carlsen was born in Oslo, Norway. His parents were Anton Carlsen (1862–1943) and Louise Larsen (1876–1957). <sepl> Carlsen has directed more than 25 films, 19 for which he wrote the screenplay. In 2006, he received the Golden Swan Lifetime Achievement Award at the Copenhagen International Film Festival. Carlsen was born on 4 June 1927 in Aalborg, Denmark. </evidence> The answer is pianist and composer.</p>

Table 12: Examples from RetroLLM on multi-hop QA datasets, with special tokens used in the clue and evidence generation stages highlighted in orange box and blue box, respectively. All generated evidence is from Wikipedia.

Example #1 from HotpotQA Dataset
<p>Question: Which American audio engineer and clandestine chemist, who was a key figure in the San Francisco Bay Area hippie movement during the 1960s recorded the album "Old and in the Way?"</p> <p>Labeled Answer: ["Owsley Stanley"]</p>
Results by RetroLLM
<p>Model Input: Question: Which American audio engineer and clandestine chemist, who was a key figure in the San Francisco Bay Area hippie movement during the 1960s recorded the album "Old and in the Way?" Your Response:</p>
<p>Model Output: <clue> Old and in the Way <sepl> Owsley Stanley </clue> <evidence> Owsley Stanley III (January 19, 1935 – March 12, 2011) was an American audio engineer and clandestine chemist. He was a key figure in the San Francisco Bay Area hippie movement during the 1960s and played a pivotal role in the decade's counterculture. <sepl> Old & In the Way (album) Old & In the Way is the self-titled first album by the bluegrass band Old & In the Way. It was recorded 8 October 1973 at The Boarding House in San Francisco by <sepl> Owsley Stanley and Vickie Babcock. <sepl> The 1960s, elements of the expanding Beat movement were incorporated into the hippie and larger counterculture movements. Neal Cassady, as the driver for Ken Kesey's bus Furthur, was the primary bridge between these two generations. Ginsberg's work also became an integral element of early 1960s hippie culture. <sepl> In San Francisco, including Quicksilver Messenger Service, Jefferson Airplane, early Jefferson Starship, Old & In the Way, Janis Joplin, Big Brother and the Holding Company, Taj Mahal, Santana, Miles Davis, The Flying Burrito Brothers, Jimi Hendrix, Johnny Cash, and Blue Cheer. </evidence> The answer is <sepl> Owsley Stanley.</p>
Example #2 from 2WikiMultiHopQA Dataset
<p>Question: Where was the director of film Ronnie Rocket born?</p> <p>Labeled Answer: ["Missoula", "Missoula, Montana"]</p>
Results by RetroLLM
<p>Model Input: Question: Where was the director of film Ronnie Rocket born? Your Response:</p>
<p>Model Output: <clue> Ronnie Rocket </clue> <evidence> Ronnie Rocket is an unfinished film project written by David Lynch, who also intended to direct it. Begun after the success of Lynch's 1977 film "Eraserhead", "Ronnie Rocket" was shelved after Lynch felt he would be unable to find financial backing for the project. <sepl> The script to the AFI, but they rejected it as they felt the planned 45-minute runtime was too long for such a figurative, nonlinear script. After the success of "Eraserhead", Lynch decided to focus on the screenplay of his second film, entitled "Ronnie Rocket", which he also wanted to direct. <sepl> He can use to produce music or cause destruction. The boy names himself Ronnie Rocket and becomes a rock star, befriending a tap-dancer named Electra-Cute. The film was to make use of several themes that have since become recurring elements within David Lynch's works, with a write-up for "The A.V. <sepl> Dave Dennison, Max and Dani Dennison's father, in Disney's "Hocus Pocus". Rocket was born in Bangor, Maine, the son of Mary Aurelia (Fogler) and Sumner Abbott "Ham" Claverie. <sepl> In 2008, Refn co-founded the Copenhagen-based production company Space Rocket Nation. Refn was born in Copenhagen, Denmark, and raised partly in New York, United States. Refn's parents are Danish film director and editor Anders Refn and cinematographer Vibeke Winding. <sepl> The two realized that "Ronnie Rocket" was unlikely to find sufficient financing to be produced, Lynch asked to see some already-written scripts to work from for his next film instead. Cornfeld found four scripts he felt Lynch would be interested in, but on hearing the name of the first, the director decided his next project would be "The Elephant Man". <sepl> His next film, settling on what would become 1980's "The Elephant Man". "Ronnie Rocket" was to feature many of the elements which have since come to be seen as Lynch's hallmarks; including industrial art direction, 1950s popular culture and physical deformity. <sepl> He shared with Billy Williams. Taylor was born in Hampstead, London and entered the film industry in his late teens working for Gainsborough Pictures at Lime Grove in Shepherds Bush. Taylor's first film was as a clapper boy on "The Young Mr Pitt" (1942). <sepl> In his films have earned them a reputation as works that "disturb, offend or mystify" general audiences. Although born in <sepl> Missoula, Montana, Lynch spent his youth traveling across the United States due to his father Donald's job for the Department of Agriculture; as a result, Lynch attended school across several states. </evidence> The answer is <sepl> Missoula, Montana.</p>

Table 13: An example comparing outputs from RetroLLM and the naive constrained beam search method. Special tokens used during the clue and evidence generation stages are highlighted in orange boxes and blue boxes, respectively. Content colored in green indicates correct (or partially correct) answers, whereas content colored in red indicates incorrect answers. All generated evidence is from Wikipedia.

Example #1 from NQ Dataset
Question: who got the first nobel prize in physics? Labeled Answer: ["Wilhelm Conrad Röntgen"]
Results by RetroLLM
Model Input: Question: who got the first nobel prize in physics? Your Response:
Model Output: <clue> Nobel Prize <sepl> Physics </clue> <evidence> the others being the Nobel Prize in Chemistry, Nobel Prize in Literature, Nobel Peace Prize, and Nobel Prize in Physiology or Medicine. The first Nobel Prize in Physics was awarded to physicist Wilhelm Röntgen in recognition of the extraordinary services he rendered by the discovery of the remarkable rays (or X-rays). </evidence> The answer is Wilhelm Röntgen .
Results by Naive Constrained Beam Search
Model Input: Question: who got the first nobel prize in physics? Your Response:
Model Output (beam_size = 5): Beam 1: <evidence> in 1901, the first year of play for the American League. The modern-era record for lowest batting average for a player that qualified for the batting title is .197 by roger maris in 1962. </evidence> The answer is Roger Maris . Beam 2: <evidence> the first physics and chemistry laboratory in a Danish school in 1903, the inclusion of housekeeping in the curriculum, and the first girls' school to include carpentry in the curriculum in 1904. </evidence> The answer is 1903 . Beam 3: <evidence> william johnson Jonathan Elliot (historian) Jonathan Elliot (178420131846) was a 19th-century </evidence> The answer is Jonathan Elliot . Beam 4: <evidence> the first physics Nobel prize winner born in an Arab country. In 2015, Cohen-Tannoudji signed the Mainau Declaration 2015 which calls for a boycott of israel. </evidence> The answer is Cohen - Tannoudji . Beam 5: <evidence> the first prize was awarded in 2004. He was the sixth of the seven children in the Shaw family, and was nicknamed Uncle Six (Luk Sook). </evidence> The answer is He was the sixth of the seven children in the Shaw family .

Table 14: An example comparing outputs from RetroLLM and the naive constrained beam search method. Special tokens used during the clue and evidence generation stages are highlighted in orange boxes and blue boxes, respectively. Content colored in green indicates correct answers, whereas content colored in red indicates incorrect answers. All generated evidence is from Wikipedia.

Example #2 from TriviaQA Dataset
<p>Question: Who was the man behind The Chipmunks?</p> <p>Labeled Answer: ["David Seville"]</p>
Results by RetroLLM
<p>Model Input: Question: Who was the man behind The Chipmunks? Your Response:</p>
<p>Model Output: <clue> The Chipmunks </clue> <evidence> known by the stage name David Seville, was an American singer-songwriter and actor, best known for creating the novelty musical group The Chipmunks (now better known as "Alvin and the Chipmunks"). Initially a stage and film actor, he rose to prominence in 1958 with the songs "Witch Doctor" and "The Chipmunk Song (Christmas Don't Be Late)," which both became "Billboard" number-one singles. </evidence> The answer is David Seville. </p>
Results by Naive Constrained Beam Search
<p>Model Input: Question: Who was the man behind The Chipmunks? Your Response:</p>
<p>Model Output (beam_size = 5): Beam 1: <evidence> david was optioned to the Triple-A Richmond Braves to start the 1992 season. He made his major league debut for the Braves on </evidence> The answer is David. Beam 2: <evidence> the chipmunks' speech was created by speeding up sound clips of normal speech. In a number of the shorts that followed, many of these same sound clips were reused, but the chipmunks' voices were created by using a higher pitch on the original recordings. </evidence> The answer is Ross Bagdasarian. Beam 3: <evidence> alvinophilia is a partialism in which a man or woman is attracted to the human navel. According to a study, it is a moderately prevalent fetish among men, and less so among women. </evidence> The answer is Alvin. Beam 4: <evidence> the chipmunks are also mostly supplemental prey but are considered more easily caught than tree squirrels, considering that they are more habitual terrestrial foragers. </evidence> The answer is Alvin. Beam 5: <evidence> the chipmunks are also mostly supplemental prey but are considered more easily caught than tree squirrels, considering that they are more habitual terrestrial foragers. </evidence> The answer is Tree Squirrels. </p>