

Retrieval as Generation: A Unified Framework with Self-Triggered Information Planning

Bo Li^{1,2}, Mingda Wang³, Gexiang Fang¹, Shikun Zhang¹, Wei Ye^{1*}

¹ National Engineering Research Center for Software Engineering, Peking University

² School of Computer Science, Peking University

³ School of Health Sciences and Biomedical Engineering, Hebei University of Technology
deepblue.lb@gmail.com, wye@pku.edu.cn

 WisdomShell/GRIP  GRIP Project  HuggingFace Model

Abstract

We revisit retrieval-augmented generation (RAG) by embedding retrieval control directly into generation. Instead of treating retrieval as an external intervention, we express retrieval decisions within token-level decoding, enabling end-to-end coordination without additional controllers or classifiers. Under the paradigm of Retrieval as Generation, we propose **GRIP** (Generation-guided Retrieval with Information Planning), a unified framework in which the model regulates retrieval behavior through control-token emission. Central to GRIP is *Self-Triggered Information Planning*, which allows the model to decide when to retrieve, how to reformulate queries, and when to terminate, all within a single autoregressive trajectory. This design tightly couples retrieval and reasoning and supports dynamic multi-step inference with on-the-fly evidence integration. To supervise these behaviors, we construct a structured training set covering answerable, partially answerable, and multi-hop queries, each aligned with specific token patterns. Experiments on five QA benchmarks show that GRIP surpasses strong RAG baselines and is competitive with GPT-4o while using substantially fewer parameters.

1 Introduction

Retrieval-Augmented Generation (RAG) improves LLMs by grounding generation on external evidence, boosting open-domain QA, fact verification, and multi-hop reasoning (Gao et al., 2023; Kandpal et al., 2023; Xu et al., 2024; Luo et al., 2024; Yang et al., 2025b; Zhou et al., 2026; Li et al., 2026b). However, most RAG systems keep retrieval external and one-shot: they retrieve documents from the initial query and then generate from a fixed context (Yu et al., 2023; Shi et al., 2023; Wang et al., 2023), an assumption that can fail when information needs emerge gradually during step-by-step

reasoning or when queries contain ambiguity and hidden dependencies.

To address this limitation, prior work explores retrieval control to improve timing and evidence use. Training-free methods (e.g., DRAGIN (Su et al., 2024) and ETC (Li et al., 2026a)) trigger retrieval via heuristic uncertainty signals, while training-based approaches learn retrieval-related behaviors with supervision, spanning both single-step control (e.g., INFO-RAG (Xu et al., 2024), Self-RAG (Asai et al., 2024), GainRAG (Jiang et al., 2025)) and agentic deep-search systems that iteratively generate queries (e.g., R1-Searcher (Song et al., 2025), ZeroSearch (Sun et al., 2025)). However, many of these methods still rely on auxiliary controllers or multi-stage procedures outside token-level generation, so retrieval timing, query reformulation, and stopping are not represented as explicit, trainable actions within a single decoding trajectory. This separation can also make it harder to attribute errors to specific decisions (e.g., retrieving too early vs. stopping too late) and to learn a consistent policy that couples retrieval with the model’s evolving intermediate states.

Rather than relying on external modules to make discrete retrieval decisions, we embed retrieval behavior directly into the model’s generative policy. Under the paradigm of Retrieval as Generation, we propose **GRIP** (Generation-guided Retrieval with Information Planning), a unified framework where retrieval is controlled by the same token-level decoding process as language generation¹. GRIP regulates retrieval via explicit control tokens, including [RETRIEVE], [ANSWER], [INTERMEDIARY], and [SOLVED]. At its core, self-triggered information planning allows the model

¹Here, ‘Retrieval as Generation’ does not mean that the retriever is internalized into the LLM. The retriever remains external; GRIP generates retrieval-control actions (e.g., retrieval triggering, query reformulation, and termination decisions) as part of the decoding trajectory, and the retrieved documents are then supplied back as context.

*Corresponding author

to decide when to retrieve, formulate follow-up queries conditioned on the evolving reasoning context, and determine when to stop and finalize the answer, all within a single autoregressive trajectory.

To support these behaviors, we introduce a structured training paradigm that supervises four answerability types with distinct control-token patterns. This supervision teaches the model to judge information sufficiency, trigger retrieval when needed, and terminate once the question is resolved. Unlike heuristic triggering, GRIP learns retrieval as a token-driven and context-sensitive capability. By integrating retrieval planning into generation, GRIP treats retrieval as an integral part of reasoning, enabling multi-step inference, adaptive retrieval depth, and robust generalization across tasks.

We conduct experiments on five knowledge-intensive QA benchmarks using multiple metrics. GRIP outperforms strong baselines and is competitive with proprietary API models. We further provide behavioral analyses of retrieval timing, query generation quality, and termination control, complemented by case studies and visualizations. In summary, our main contributions are as follows:

- We propose **GRIP**, a unified *Retrieval as Generation* framework that internalizes retrieval into token-level decoding via explicit control tokens, enabling retrieval planning, query reformulation, and termination within a single autoregressive trajectory.
- We extensively evaluate GRIP on five QA benchmarks and provide a mechanistic analysis of retrieval-as-generation behaviors. Beyond consistent gains over strong baselines, our studies show that GRIP learns task-aware retrieval depth and reliable stopping, yielding more adaptive and controllable retrieval trajectories across datasets.

2 Generation-guided Retrieval with Information Planning

2.1 Token-Level Control for Unified Retrieval and Generation

We propose **GRIP**, a *Retrieval-as-Generation* framework that internalizes retrieval control into token-level decoding. By emitting a small set of control tokens, the model decides when to retrieve, how to reformulate queries, and when to terminate. At the core is *self-triggered information planning*,

which assesses information sufficiency and executes multi-step retrieval. To realize this control in a unified and interpretable manner, GRIP augments the model’s output space with a small set of control tokens, allowing retrieval and termination decisions to be made directly during decoding.

GRIP introduces four tokens: [RETRIEVE] (request external evidence), [INTERMEDIARY] (emit a partial state), [ANSWER] (start the final response), and [SOLVED] (terminate). They form structured patterns that define two main branches: [INTERMEDIARY]→[RETRIEVE] for continuing evidence acquisition, and [ANSWER]→[SOLVED] for completion, enabling multi-step and recursive retrieval when needed.

This minimal interface keeps control within autoregressive generation, enabling end-to-end learning without external controllers.

2.2 Self-Triggered Information Planning

GRIP adopts a language-native mechanism called *self-triggered information planning*, which enables the model to autonomously decide when and how to retrieve information during generation. This process unfolds autoregressively, driven by the emission of control tokens. GRIP organizes its planning behavior into the following loop:

1. Initial Decision: Given an original query, the model first determines whether its internal knowledge suffices.

- If so, it emits [ANSWER] followed by a final prediction, and then terminates with [SOLVED].
- If not, it emits [INTERMEDIARY] to provide partial knowledge or reasoning, followed by [RETRIEVE] and generates a new query for external retrieval.

2. Retrieval and New Query Generation: The model uses the original or generated query to retrieve relevant passages, which are incorporated as augmented input. It then reassesses whether the available information is sufficient:

- If the evidence is sufficient, the model emits [ANSWER] followed by a final prediction and concludes with [SOLVED].
- Otherwise, it generates [INTERMEDIARY] explanation and emits a follow-up [RETRIEVE] token with a new query to continue the loop.

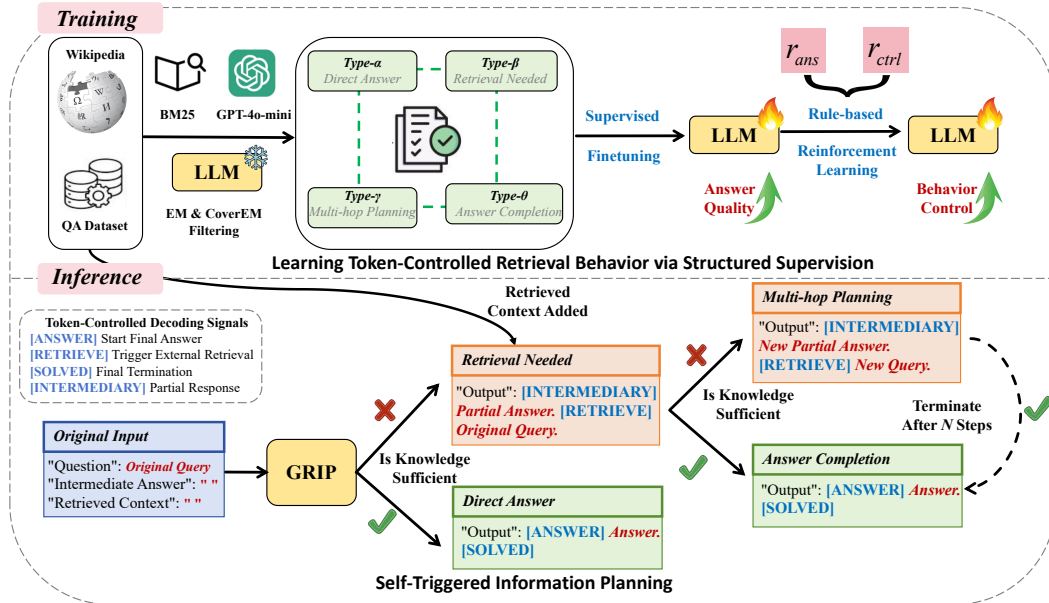


Figure 1: Overview of the GRIP framework. The top half illustrates the training process: four behavior types are constructed for supervised fine-tuning, followed by reinforcement learning to improve answer quality and retrieval control. The bottom half depicts self-triggered information planning, where GRIP dynamically emits control tokens to decide whether to retrieve, generate new queries, or terminate reasoning, all within a unified decoding trajectory.

Notably, GRIP distinguishes between initial and follow-up retrievals: 1) In the initial decision step, the [RETRIEVE] token is followed by the original query, as the model lacks any retrieved context. 2) In subsequent planning steps, the model conditions on its previously generated intermediate response and the original query to synthesize a new follow-up query tailored to unresolved aspects of the task.

3. Multi-hop Planning: This alternation continues recursively: each [INTERMEDIARY] reflects the model’s current understanding, and each [RETRIEVE] introduces a targeted query to acquire additional knowledge.

4. Termination: The retrieval-reasoning loop concludes when the model emits [ANSWER] followed by [SOLVED], indicating that sufficient information has been gathered to produce a final answer. To ensure efficiency and behavioral stability, GRIP uses a decoding-time retrieval budget, which can be adjusted at inference time and is set to three retrieval rounds by default. If this limit is reached, the model is required to finalize its response: it emits [ANSWER] and generates the final answer based on the currently available information.

This planning mechanism is fully internal to the model and operates within a unified decoding flow. No confidence thresholds, external classifiers, or prompt chaining are needed. Through supervised

training on structured examples, the model learns to coordinate retrieval and reasoning as a compositional skill. Pseudocode for the GRIP decoding process is provided in Appendix M.

2.3 Learning Token-Controlled Retrieval Behavior via Structured Supervision

To enable GRIP to learn retrieval as a generation-internal process, we formulate a structured supervision paradigm that aligns specific control token patterns with different retrieval behaviors. We organize training examples into four structured supervision scenarios, each reflecting a distinct behavioral trajectory. This token-controlled supervision allows GRIP to acquire compositional retrieval strategies, such as when to retrieve, how to formulate a new query, and when to terminate reasoning. We illustrate the four supervision types and their corresponding token structures in Figure 2.

Type-α samples represent directly answerable queries. These are queries that can be answered correctly using the model’s internal knowledge alone, without any external context. We identify such instances by running the backbone model (e.g., LLaMA3-8B-instruct) and filtering cases that consistently yield exact-match (EM) answers across multiple decoding attempts. The model is trained to generate [ANSWER] followed by the final answer, and conclude with [SOLVED].

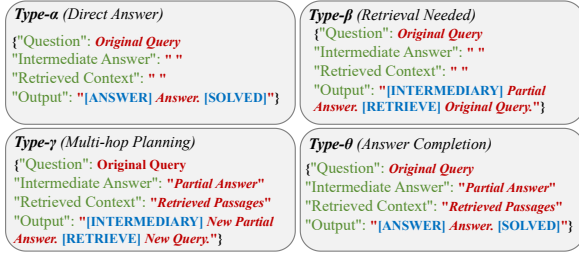


Figure 2: Four types of structured training samples in GRIP, each aligned with specific token patterns for controlling retrieval behavior. This design enables the model to learn when to retrieve, plan new queries, or finalize answers via language-native supervision.

Type- β samples involve cases where the model provides partially correct or noisy answers that contain the gold answer but lack clarity. These cases are detected using a relaxed coverage-based EM metric, defined as responses that contain the target answer but fail to fully resolve the query. Here, the model is expected to emit [INTERMEDIARY] to reflect its current partial knowledge, followed by [RETRIEVE] and a retrieval query. This setup teaches the model to recognize uncertainty and trigger retrieval appropriately.

Type- γ samples simulate multi-hop or complex knowledge cases where neither the model nor basic retrieval suffices. We identify such samples by detecting failures to answer correctly even with BM25-based retrieval, and then rely on GPT-4o-mini to generate an improved follow-up query focused on missing aspects. The model is trained to generate an intermediate answer ([INTERMEDIARY]) and a refined query ([RETRIEVE]), and to iterate through multiple planning steps before concluding with [ANSWER] and [SOLVED].

Type- θ samples are cases where both the model and retriever can find passages that cover the answer superficially, but the correct response requires compressing, selecting, or synthesizing salient details. We use CoverEM to label these samples and train the model to refine retrieved content through reasoning. This teaches the model to emit token transitions like [INTERMEDIARY] \rightarrow [RETRIEVE] \rightarrow [ANSWER], modeling selective information usage in noisy contexts.

Token supervision is applied using standard teacher forcing over these structured sequences. By associating each sample type with a precise token trajectory, the model learns to internalize retrieval control within its generative policy. Representative

examples are provided in Appendix D.

2.4 Training Objective and Optimization

GRIP is trained in two stages: (1) supervised fine-tuning on structured behavioral examples to learn retrieval patterns, and (2) rule-based reinforcement learning (RL) guided by two reward signals. Together, these stages enable the model to acquire retrieval planning skills that are both interpretable and adaptive.

Supervised Fine-Tuning. The initial phase of GRIP training employs supervised learning to expose the model to diverse token-driven retrieval behaviors. Each training instance is labeled with a structured target sequence containing both reasoning content and control tokens (e.g., [ANSWER], [RETRIEVE]), reflecting the semantic structure of the corresponding query type. The objective is to minimize the standard autoregressive cross-entropy loss over the entire target sequence:

$$\mathcal{L}_{SFT} = - \sum_{t=1}^T \log P_{\theta}(y_t | y_{<t}, x)$$

where x is the input query, and y_1, \dots, y_T is the ground-truth sequence containing both natural language and special control tokens.

Reinforcement Fine-Tuning. While SFT provides strong initial behavior modeling, we observe that the model tends to over-trigger retrieval actions, often emitting multiple [RETRIEVE] tokens even when unnecessary. To address this, we apply a RL phase based on rule-based reward optimization framework. We define two complementary reward signals:

1) Answer Fidelity Reward r_{ans} . Measures the similarity between the model’s generated answer and the ground-truth reference using BLEU score (Chang et al., 2025), encouraging accurate and semantically faithful responses.

2) Control Accuracy Reward r_{ctrl} . Rewards the correct emission of control token patterns and penalizes incorrect or missing tokens, thereby sharpening the model’s understanding of retrieval control logic. The final reward is computed as:

$$R = r_{\text{ans}} + r_{\text{ctrl}}$$

The model is optimized using DAPO (Yu et al., 2025), encouraging answer-faithful and token-accurate behaviors that promote a more balanced retrieval strategy.

Overall, GRIP is designed to unify retrieval planning and language generation under a single autoregressive framework. Compared with closely related search-based methods, its advantage comes from two aspects: (1) **retrieval planning is internalized as token-level control within generation itself**, rather than implemented through external search-action loops; and (2) **multi-step retrieval behavior is learned via one-step decision optimization**, rather than long-horizon search-policy optimization. This makes GRIP both conceptually simple and practically stable, while still enabling adaptive retrieval depth and controllable stopping behavior at inference time.

3 Experimental Setup and Main Results

3.1 Datasets and Evaluation Metrics

Datasets. We evaluate GRIP on five widely used benchmarks for open-domain and multi-hop question answering: **HotpotQA** (Ho et al., 2020), **PopQA** (Mallen et al., 2022), **Natural Questions (NQ)** (Kwiatkowski et al., 2019), **WebQuestions (WebQ)** (Berant et al., 2013), and **TriviaQA** (Joshi et al., 2017). These datasets cover diverse query types, including factual recall, contextual understanding, and multi-step reasoning, providing a comprehensive testbed for retrieval-augmented generation. We use the official evaluation splits and verify that there is no overlap between evaluation questions and any training data used in our experiments. See Appendix A for more detail.

Evaluation Metrics. We report three primary QA metrics: **Exact Match (EM)** (exact normalized string match), **ROUGE** (average of ROUGE-1/2/L), and **F1** (token-level F1 between prediction and reference). For compact reporting, we also use **Avg.Score**, defined as the unweighted mean of all reported metric values across all evaluated datasets (i.e., averaging over *datasets* and *metrics* with equal weight). We additionally report **CoverEM** for behavioral analysis; its definition and results are provided in Appendix C.

3.2 Experimental Details

Structured Training Data Construction. We construct structured training data based on the training sets of TriviaQA and NQ, each paired with

three retrieved passages. The same top-3 retrieval setting is applied during inference for consistency. Specifically, we construct 40,000 structured training instances for the supervised fine-tuning stage, covering all four behavior types described in Section 2.3. For the RL stage, we generate an additional 5,000 samples with the same distribution across behavior types. The prompts used to construct structured data are provided in Appendix K. Additional settings for control tokens are described in Appendix E. Unless otherwise specified, we set the maximum retrieval steps to three, and we analyze larger test-time budgets in § 4.4.

Supervised Fine-tuning. We fine-tune GRIP on 40,000 structured samples using LLaMA3-8B (Grattafiori et al., 2024) as the backbone with full-parameter training. The model is trained for 8 epochs with a micro-batch size of 4 per GPU (total batch size 32 across 8 A800 GPUs). We use the AdamW with a learning rate of 1×10^{-6} . A cosine learning rate schedule is applied with 10% warm-up. Input sequences are truncated at 8192 tokens. Due to space limitations, results using the Qwen2.5-7B-Instruct (Yang et al., 2024, 2025a) are provided in the Appendix P.

Reinforcement Learning. We adopt the DAPO (Yu et al., 2025) algorithm, an improved variant of GRPO (Shao et al., 2024), to fine-tune GRIP on 5,000 structured samples. Two task-specific rewards are used: (1) *Answer Fidelity*, computed via BLEU score against the reference answer, more details about the reward selection ablation can be found in the Appendix F; and (2) *Control Accuracy*, which assigns 0.5 points for correctly emitted control tokens. We train the model for only 1 epoch with a learning rate of 1×10^{-7} , which is sufficient to stabilize retrieval behaviors under the DAPO framework. Please refer to Appendix O for more details.

3.3 Comparison Models

We compare GRIP with two groups of baselines. **Training-free** methods include direct prompting (**Instruct**), one-shot retrieval (**Single RAG**), proprietary API models (**GPT-3.5 Turbo**, **GPT-4o**) as references, and dynamic retrieval methods **FLARE** (Jiang et al., 2023), **DRAGIN** (Su et al., 2024) and **ETC** (Li et al., 2026a) (8-shot ICL following their original setups). **Training-based** methods include a matched-data SFT baseline (**SFT-RAG**) and state-of-the-

	HotpotQA			PopQA			NQ			WebQ			TriviaQA			Avg. Score
	EM	ROUGE	F1	EM	ROUGE	F1	EM	ROUGE	F1	EM	ROUGE	F1	EM	ROUGE	F1	
<i>Training-free Method</i>																
Instruct	17.2	21.6	25.9	17.4	19.6	23.2	14.4	23.8	20.3	14.8	25.1	29.4	46.2	45.3	55.1	26.6
GPT-3.5 Turbo	26.2	32.4	38.2	29.1	31.1	35.3	20.7	34.3	30.0	15.9	27.4	31.9	55.6	57.8	69.7	35.7
GPT-4o	33.2	40.2	47.0	30.6	32.3	39.9	26.5	42.7	28.3	23.5	32.3	37.0	65.7	64.3	78.2	41.4
Single RAG	26.1	31.6	37.2	22.8	26.6	31.1	19.3	28.4	24.8	14.0	22.7	26.6	46.4	47.0	56.8	30.8
FLARE	23.2	27.9	32.8	14.3	16.0	18.4	14.7	22.4	21.3	24.2	30.3	34.7	48.6	48.5	56.4	28.9
DRAGIN	27.9	32.6	38.7	15.5	16.8	19.8	23.9	32.8	28.5	25.2	31.5	35.7	55.3	53.6	64.6	33.5
ETC	32.5	37.7	44.2	30.5	32.5	37.5	20.9	26.7	30.7	18.9	26.6	30.4	52.9	52.1	63.0	35.8
<i>Training-based Method</i>																
SFT-RAG	20.3	24.1	28.6	29.4	25.2	30.4	20.8	17.9	21.3	18.9	18.3	23.1	50.1	24.8	57.2	27.4
Self-RAG	19.6	23.8	26.7	18.1	22.3	22.8	15.7	22.4	24.0	16.4	26.5	27.4	50.2	47.3	57.5	28.0
INFO-RAG	19.9	23.7	26.9	18.3	22.6	23.0	17.2	22.9	24.9	18.1	27.7	28.9	50.8	47.8	58.1	28.7
RobustRAG	27.6	31.8	37.5	29.7	27.7	32.4	26.4	25.1	29.2	21.5	25.0	29.1	48.8	47.7	57.9	33.2
GainRAG	31.4	35.6	41.8	30.1	33.3	38.1	22.9	27.9	32.2	16.5	24.5	28.9	50.3	49.1	59.2	34.8
R1-Searcher	26.0	29.1	34.9	41.6	35.2	41.3	25.8	24.9	28.7	21.8	26.1	30.6	56.0	53.3	64.9	36.0
RetRobust	29.6	34.9	40.9	34.1	35.1	40.4	24.2	29.0	33.8	21.8	27.4	31.7	53.6	50.9	61.9	36.6
InstructRAG	31.2	36.8	42.3	33.1	35.7	40.3	29.5	29.5	33.6	19.5	26.8	31.4	51.3	52.2	62.5	37.0
GRIP(ours)	33.0	37.6	44.1	38.6	37.5	38.4	32.1	35.8	32.0	31.4	39.3	34.6	57.9	55.9	67.4	41.0 _{44.0}
<i>w/o RL</i>	31.6	36.6	43.0	38.1	37.1	37.6	32.6	36.1	32.7	32.0	39.9	35.1	57.0	55.2	66.8	40.7

Table 1: Main results on five QA benchmarks. We report EM, ROUGE (average of ROUGE-1/2/L), and F1; Avg.Score is the unweighted mean over all datasets and metrics. All *open-source* baselines are reproduced using official implementations and released checkpoints under the same evaluation setting, and results are averaged over three random seeds. Results on a domain-specific dataset are provided in Appendix B.

art frameworks **Self-RAG** (Asai et al., 2024), **INFO-RAG** (Xu et al., 2024), **RobustRAG** (Fang et al., 2024), **RetRobust** (Yoran et al.), **InstructRAG** (Wei et al., 2024), **RobustRAG** (Fang et al., 2024), **GainRAG** (Jiang et al., 2025), and **R1-Searcher** (Song et al., 2025).

To ensure a fair comparison, we reproduce all *open-source* baselines with official code/checkpoints when available and evaluate all models under the same protocol: LLaMA3-8B backbone, Wikipedia corpus, BM25 retriever with top-3 passages per query. For methods without explicit answer spans (e.g., INFO-RAG, RobustRAG, Self-RAG), we apply minimal output standardization for consistent scoring. We ablate retrievers (DPR and BM25+DPR) in Appendix L and use BM25 by default for its performance–efficiency trade-off.

3.4 Main Results

Table 1 presents a comprehensive comparison between GRIP and a broad range of RAG methods. GRIP consistently achieves the best overall performance among open-source systems, outperforming strong training-based baselines such as GainRAG and RobustRAG across all five benchmarks, and also surpassing recent deep-search baselines such as R1-Searcher in Avg.Score. By internalizing

retrieval decisions into the autoregressive generation process via token-level control, GRIP eliminates the need for external heuristics or controllers. Despite using a smaller backbone (LLaMA3-8B), GRIP reaches an Avg.Score of 41.0, close to GPT-4o, demonstrating strong effectiveness under a lightweight setting.

Gains are most pronounced on multi-hop or compositional benchmarks (e.g., HotpotQA and TriviaQA), where dynamic evidence acquisition is crucial. Beyond exact matching, GRIP also yields consistently higher ROUGE and F1, indicating better answer completeness and alignment. Overall, these results support GRIP as an efficient and generalizable framework for controlled retrieval-augmented generation across diverse QA scenarios.

4 Analysis

4.1 Retrieval-Depth Adaptivity

A central benefit of token-level planning is *bidirectional* control over retrieval depth: GRIP suppresses retrieval when the current decoding state is already sufficient, yet deepens retrieval when evidence is missing. To make this behavior concrete, we study two complementary subsets.

First, we consider cases where DRAGIN/GainRAG perform at least one retrieval while GRIP an-

	NQ(54.3%)			WebQ(24.3%)			Avg. Score
	EM	ROUGE	F1	EM	ROUGE	F1	
DRAGIN	30.3	33.9	39.3	42.5	48.4	52.9	41.2
GainRAG	31.6	34.4	39.6	33.0	40.7	45.6	37.4
GRIP	31.1	34.0	39.8	44.9	50.7	55.6	42.7

Table 2: Subset where DRAGIN/GainRAG retrieve at least once while GRIP answers without retrieval. Percentages denote the subset proportion per dataset.

	NQ(23.1%)			WebQ(43.8%)			Avg. Score
	EM	ROUGE	F1	EM	ROUGE	F1	
DRAGIN	11.4	16.7	19.0	18.6	23.8	28.3	19.6
GainRAG	8.8	14.8	17.4	11.1	18.4	22.4	15.5
GRIP	15.4	20.3	23.4	19.9	24.6	28.9	22.1

Table 3: Subset where GRIP performs two retrievals while DRAGIN/GainRAG retrieve at most once. Percentages denote the subset proportion per dataset.

swers directly; Table 2 shows that GRIP achieves higher accuracy on this subset despite using fewer retrieval calls, indicating it can avoid redundant retrieval rather than following a fixed schedule.

Second, we analyze cases where GRIP performs two retrievals while baselines retrieve at most once; Table 3 shows consistent gains on these harder examples, suggesting GRIP can detect evidence insufficiency and extend retrieval accordingly. Together, these results support that GRIP learns context-sensitive retrieval depth control from structured token supervision.

4.2 Adaptive Retrieval Across Tasks

Table 4 reports the mean number of retrieval calls per example. GRIP exhibits clear task-aware adaptivity in retrieval depth: it retrieves more frequently on HotpotQA (1.44) and PopQA (1.58), where multi-hop reasoning or long-tail entities often require external evidence (Mallen et al., 2022; Asai et al., 2024), while triggering substantially fewer retrievals on NQ (0.76), where many questions can be answered from the model’s parametric knowledge. Compared with training-based baselines that follow nearly fixed retrieval schedules (e.g., GainRAG always retrieves once), GRIP adjusts retrieval frequency across datasets without external controllers. Moreover, R1-Searcher performs markedly more retrievals (4.67–5.75 across datasets; 5.12 on average), indicating a substantially higher retrieval cost despite its weaker end performance in our setting. In contrast, GRIP achieves strong accuracy with a

much lower retrieval budget.

	Hotpot	PopQA	NQ	WebQ	Trivia	Avg. Count
DRAGIN	1.13	1.26	1.06	1.02	1.08	1.11
GainRAG	1.0	1.0	1.0	1.0	1.0	1.0
R1-Searcher	5.75	5.34	5.09	4.67	4.76	5.12
GRIP	1.44	1.58	0.76	1.15	1.25	1.24
<i>w/o RL</i>	1.99	1.98	0.77	1.26	1.99	1.60

Table 4: Mean retrieval count per dataset. GRIP adapts retrieval frequency across tasks, and RL further reduces retrieval calls while preserving this adaptivity. Hotpot and Trivia denote HotpotQA and TriviaQA, respectively.

Besides, rule-based RL further refines GRIP’s retrieval policy by consistently reducing redundant calls across *all* datasets: the overall average drops from 1.60 to 1.24 ($\sim 22.5\%$ reduction), while preserving the same task-aware retrieval pattern. A finer-grained breakdown of this behavioral shift is provided in Appendix G.

4.3 Improving Retrieval Quality by Generating New Queries




Figure 3: The percentage of samples where the gold answer span is found in the top-1 or top-3 retrieved passages. GRIP-generated queries significantly improve answer presence in retrieved contexts across both NQ and WebQ, indicating more effective retrieval behavior.

To evaluate the effectiveness of GRIP’s self-triggered information planning, we analyze the retrieval quality of its generated new queries compared to the original ones. Specifically, we measure the proportion of top-1 and top-3 retrieved passages that contain the gold answer. As shown in Figure 3, GRIP significantly improves retrieval coverage on both NQ and WebQ. This improvement does not stem from superficial query rewriting. Instead, it emerges from GRIP’s iterative reasoning process. By leveraging intermediate answers and previously retrieved content, GRIP gradually refines its focus toward unresolved aspects of the question. This progressive formulation yields more targeted new

queries that are more likely to retrieve passages containing the correct answer. Notably, larger gains on top-1 than top-3 suggest that GRIP improves evidence ranking, consistent with query reformulation being coupled to intermediate reasoning states for more efficient retrieval.

4.4 Controllable Retrieval Budget and Depth Extrapolation

We vary the maximum allowed retrieval steps B at inference time and report both the realized retrieval frequency (Avg.Count) and overall Avg.Score (Table 5). As B increases from 3 to 10, Avg.Score rises from 41.0 to 41.8 while Avg.Count grows only from 1.24 to 1.62. Avg.Count stays far below the budget cap across all settings, indicating that GRIP does not blindly consume the available budget but retrieves adaptively when the decoding trajectory suggests an information gap, supporting our token-level planning claim.

Max B	Avg.Count (per dataset)						Avg. Score
	Hotpot	PopQA	NQ	WebQ	Trivia	Avg	
3	1.44	1.58	0.76	1.15	1.25	1.24	41.0
5	1.56	1.73	1.01	1.45	1.31	1.41	41.2
7	1.66	1.79	1.12	1.65	1.35	1.51	41.5
10	1.74	1.86	1.22	1.90	1.39	1.62	41.8

Table 5: Effect of the maximum retrieval budget B at inference time. Hotpot and Trivia denote HotpotQA and TriviaQA, respectively.

This adaptive control also exhibits *extrapolative generalization*: although training trajectories contain at most three retrieval steps, GRIP benefits from larger test-time budgets ($B=5, 7, 10$) with consistent gains, suggesting it learns when additional evidence is needed rather than a fixed-step pattern. Extra retrieval is mostly allocated to evidence-uncertain benchmarks (NQ: 0.76→1.22; WebQ: 1.15→1.90), and improvements show diminishing returns. Overall, B provides a practical control parameter to trade accuracy for cost, while GRIP remains adaptive beyond the retrieval depth seen in training. Appendix N shows the distribution of retrieval counts.

4.5 Replacing Proprietary Teachers with Open-Source Supervision

To test whether GRIP’s gains rely on proprietary distillation, we conduct a controlled teacher-replacement study on **Type- γ** supervision, which is the only place where our data construction uses an

Teacher	Training	Avg.Score
GPT-4o-mini	GRIP (SFT)	40.7
	GRIP (SFT+RL)	41.0
LLaMA3-8B-Instruct	GRIP (SFT)	38.8
	GRIP (SFT+RL)	39.0
Qwen3-32B	GRIP (SFT)	40.2
	GRIP (SFT+RL)	40.3

Table 6: Teacher replacement results, where only the GPT-labeled subset is re-annotated.

external teacher. Specifically, Type- γ samples use GPT-4o-mini to generate improved follow-up retrieval queries when the backbone model fails. We replace GPT-4o-mini and re-annotate the follow-up queries for *all* Type- γ samples using two open teachers: LLaMA3-8B-Instruct and Qwen3-32B. All other components are kept identical across settings, including the construction and supervision of other samples, the prompts and decoding hyperparameters for annotation.

Table 6 shows that GRIP remains strong under open teachers, indicating that the gains primarily arise from our framework design rather than GPT-specific artifacts. In particular, replacing GPT-4o-mini with Qwen3-32B yields a comparable final performance (40.3 vs. 41.0), suggesting that GRIP can still induce effective dynamic retrieval and query reformulation behaviors without relying on closed-source supervision. Using a backbone-scale teacher (LLaMA3-8B-Instruct) leads to a larger drop, implying that teacher capacity mainly affects the *quality* of trajectory instantiation rather than the *existence* of the learned retrieval policy. Overall, teacher outputs serve as interchangeable demonstrations, while the token-controlled supervision is the key factor that turns retrieval into a learnable generation skill.

4.6 General Capability Preservation

Model	MMLU (Acc)	MBPP (Pass@1)
Instruct	66.56	54.6
SFT-RAG	62.93	47.8
GRIP	65.73	53.8

Table 7: General capability evaluation on closed-form non-RAG benchmarks.

A potential concern is whether fine-tuning GRIP for retrieval planning may degrade the model’s general-purpose abilities on non-RAG tasks. To

Pairwise comparison (GRIP vs.)	Win	Equal	Loss
Instruct	51.0	37.5	11.5
SFT-RAG	90.5	8.5	1.0
R1-Searcher	59.5	32.5	8.0

Table 8: Pairwise GPT-4o evaluation on CNN/Daily-Mail summarization. Results are averaged over both output orders to mitigate position bias.

directly assess this, we conduct additional evaluations on two types of tasks that involve no retrieval: (1) closed-form knowledge and coding benchmarks, including MMLU (Hendrycks et al., 2020) and MBPP (Austin et al., 2021), and (2) an open-ended generation task, CNN/DailyMail summarization (Hermann et al., 2015), evaluated by GPT-4o as a pairwise judge.

Closed-form benchmarks. Table 7 reports results on MMLU and MBPP. Compared with the base INSTRUCT model, GRIP shows only a minor decrease, dropping by 0.83 points on MMLU and 0.8 points on MBPP. This suggests that learning retrieval-planning behaviors does not substantially harm the model’s general knowledge, reasoning, or coding ability. In contrast, the matched fine-tuning baseline SFT-RAG exhibits noticeably larger degradation on both benchmarks, indicating that GRIP better preserves the backbone model’s general capabilities while still introducing retrieval-oriented behaviors.

Open-ended summarization. We further evaluate open-ended generation quality on CNN/Daily-Mail summarization, where no retrieval is involved. Following recent practice, we use GPT-4o as a pairwise judge to compare model outputs. To reduce position bias, each model pair is evaluated in both orders, and we report the order-averaged Win/Equal/Loss rates in Table 8. GRIP remains competitive with the base INSTRUCT model, achieving a 51.0% win rate versus 11.5% loss rate, and is preferred much more often than other fine-tuned or search-oriented baselines such as SFT-RAG and R1-SEARCHER. These results suggest that GRIP preserves general-purpose generation quality while learning retrieval control.

Overall, across both closed-form and open-ended non-RAG evaluations, GRIP shows minimal degradation relative to the base INSTRUCT model, while substantially outperforming retrieval-oriented baselines in preserving general-purpose capabilities. This indicates that integrating retrieval

planning into token-level generation does not come at the cost of broadly useful language modeling ability.

5 Related Work

Most standard RAG pipelines retrieve documents in a static, one-shot fashion based on the initial query (Gao et al., 2023; Fan et al., 2024; Xiong et al., 2024; Ye et al., 2024; Zhang et al., 2025, 2026; Li et al., 2026b). To reduce unnecessary retrieval, several works estimate retrieval necessity at inference time using heuristic uncertainty signals, such as entropy-based filtering or logit-margin thresholds (Jiang et al., 2023; Ram et al., 2023; Su et al., 2024; Li et al., 2026a). While lightweight, these approaches rely on hand-crafted criteria rather than learning retrieval behaviors end-to-end.

Recent studies make retrieval behavior learnable by introducing retrieval-relevant supervision, e.g., self-reflection/critique signals or contrastive/adversarial objectives (Asai et al., 2024; Jiang et al., 2025; Xu et al., 2024; Fang et al., 2024). In parallel, deep search and multi-step retrieval systems further strengthen search-and-reason routines for complex queries (Song et al., 2025; Sun et al., 2025), as well as planning-oriented multi-hop methods (Lee et al., 2024; Jin et al., 2025; Yan and Ling, 2025; Zhu et al., 2026; Xu et al., 2026). However, many of these approaches still rely on external controllers or multi-stage pipelines for retrieval decisions, which can complicate integration and reduce transparency. In contrast, GRIP internalizes retrieval as token-level generative actions, realizing timing, query reformulation, and termination within a single autoregressive trajectory.

6 Conclusion

We present GRIP, a unified retrieval-augmented generation framework that integrates retrieval behavior into the generation process through self-triggered information planning. This design enables dynamic control over retrieval without relying on external modules. Experimental results across five QA benchmarks confirm that GRIP consistently outperforms existing RAG methods and achieves performance comparable to GPT-4o. We hope this work provides a foundation for future efforts in controllable and efficient generation.

Limitations

Although GRIP learns to regulate retrieval through discrete control tokens, its behavior remains sensitive to the design of the retrieval interface, such as the maximum retrieval budget B and the evidence packing strategy (e.g., chunking and top- k). In this work we keep these choices fixed for fair comparison, but more adaptive evidence budgeting and context structuring could further improve robustness across domains.

References

- Akari Asai, Zeqiu Wu, Yizhong Wang, Avirup Sil, and Hannaneh Hajishirzi. 2024. Self-rag: Learning to retrieve, generate, and critique through self-reflection.
- Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, and 1 others. 2021. Program synthesis with large language models. *arXiv preprint arXiv:2108.07732*.
- Jonathan Berant, Andrew K. Chou, Roy Frostig, and Percy Liang. 2013. [Semantic parsing on freebase from question-answer pairs](#). In *Conference on Empirical Methods in Natural Language Processing*.
- Yapei Chang, Yekyung Kim, Michael Krumbick, Amir Zadeh, Chuan Li, Chris Tanner, and Mohit Iyyer. 2025. [Bleuberi: Bleu is a surprisingly effective reward for instruction following](#). *ArXiv*, abs/2505.11080.
- Wenqi Fan, Yujuan Ding, Liangbo Ning, Shijie Wang, Hengyun Li, Dawei Yin, Tat-Seng Chua, and Qing Li. 2024. A survey on rag meeting llms: Towards retrieval-augmented large language models. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 6491–6501.
- Feiteng Fang, Yuelin Bai, Shiwen Ni, Min Yang, Xiaojun Chen, and Ruifeng Xu. 2024. [Enhancing noise robustness of retrieval-augmented language models with adaptive adversarial training](#). In *Annual Meeting of the Association for Computational Linguistics*.
- Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, and Haofen Wang. 2023. Retrieval-augmented generation for large language models: A survey. *arXiv preprint arXiv:2312.10997*.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, and 1 others. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2020. Measuring massive multitask language understanding. *arXiv preprint arXiv:2009.03300*.
- Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. Teaching machines to read and comprehend. *Advances in neural information processing systems*, 28.
- Xanh Ho, A. Nguyen, Saku Sugawara, and Akiko Aizawa. 2020. [Constructing a multi-hop qa dataset for comprehensive evaluation of reasoning steps](#). *ArXiv*, abs/2011.01060.
- Yi Jiang, Sendong Zhao, Jianbo Li, Haochun Wang, and Bing Qin. 2025. Gainrag: Preference alignment in retrieval-augmented generation through gain signal synthesis. *arXiv preprint arXiv:2505.18710*.
- Zhengbao Jiang, Frank F Xu, Luyu Gao, Zhiqing Sun, Qian Liu, Jane Dwivedi-Yu, Yiming Yang, Jamie Callan, and Graham Neubig. 2023. Active retrieval augmented generation. *arXiv preprint arXiv:2305.06983*.
- Bowen Jin, Hansi Zeng, Zhenrui Yue, Dong Wang, Hamed Zamani, and Jiawei Han. 2025. [Search-r1: Training llms to reason and leverage search engines with reinforcement learning](#). *ArXiv*, abs/2503.09516.
- Mandar Joshi, Eunsol Choi, Daniel S. Weld, and Luke Zettlemoyer. 2017. [Triviaqa: A large scale distantly supervised challenge dataset for reading comprehension](#). *ArXiv*, abs/1705.03551.
- Nikhil Kandpal, Haikang Deng, Adam Roberts, Eric Wallace, and Colin Raffel. 2023. Large language models struggle to learn long-tail knowledge. In *International Conference on Machine Learning*, pages 15696–15707. PMLR.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur P. Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, Kristina Toutanova, Llion Jones, Matthew Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob Uszkoreit, Quoc V. Le, and Slav Petrov. 2019. Natural questions: A benchmark for question answering research. *Transactions of the Association for Computational Linguistics*, 7:453–466.
- Myeonghwa Lee, Seonho An, and Min-Soo Kim. 2024. [Planrag: A plan-then-retrieval augmented generation for generative large language models as decision makers](#). In *North American Chapter of the Association for Computational Linguistics*.
- Bo Li, Tian Tian, Zhenghua Xu, Hao Cheng, Shikun Zhang, and Wei Ye. 2026a. Modeling uncertainty trends for timely retrieval in dynamic rag. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pages 31527–31535.

- Bo Li, Zhenghua Xu, and Rui Xie. 2026b. Language drift in multilingual retrieval-augmented generation: Characterization and decoding-time mitigation. In *Fortieth AAAI Conference on Artificial Intelligence, Thirty-Eighth Conference on Innovative Applications of Artificial Intelligence, Sixteenth Symposium on Educational Advances in Artificial Intelligence, AAAI 2026, Singapore, January 20-27, 2026*, pages 31519–31526. AAAI Press.
- Kun Luo, Zheng Liu, Shitao Xiao, Tong Zhou, Yubo Chen, Jun Zhao, and Kang Liu. 2024. [Landmark embedding: A chunking-free embedding method for retrieval augmented long-context large language models](#). In *Annual Meeting of the Association for Computational Linguistics*.
- Alex Troy Mallen, Akari Asai, Victor Zhong, Rajarshi Das, Hannaneh Hajishirzi, and Daniel Khashabi. 2022. [When not to trust language models: Investigating effectiveness of parametric and non-parametric memories](#). In *Annual Meeting of the Association for Computational Linguistics*.
- Ori Ram, Yoav Levine, Itay Dalmedigos, Dor Muhlgay, Amnon Shashua, Kevin Leyton-Brown, and Yoav Shoham. 2023. In-context retrieval-augmented language models. *Transactions of the Association for Computational Linguistics*, 11:1316–1331.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Jun-Mei Song, Mingchuan Zhang, Y. K. Li, Yu Wu, and Daya Guo. 2024. [Deepseekmath: Pushing the limits of mathematical reasoning in open language models](#). *ArXiv*, abs/2402.03300.
- 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](#). In *North American Chapter of the Association for Computational Linguistics*.
- Huatong Song, Jinhao Jiang, Yingqian Min, Jie Chen, Zhipeng Chen, Wayne Xin Zhao, Lei Fang, and Ji-Rong Wen. 2025. [R1-searcher: Incentivizing the search capability in llms via reinforcement learning](#). *ArXiv*, abs/2503.05592.
- Weihang Su, Yichen Tang, Qingyao Ai, Zhijing Wu, and Yiqun Liu. 2024. [Dragin: Dynamic retrieval augmented generation based on the real-time information needs of large language models](#). *arXiv preprint arXiv:2403.10081*.
- Hao Sun, Zile Qiao, Jiayan Guo, Xuanbo Fan, Yingyan Hou, Yong Jiang, Pengjun Xie, Yan Zhang, Fei Huang, and Jingren Zhou. 2025. [Zerosearch: Incentivize the search capability of llms without searching](#). *ArXiv*, abs/2505.04588.
- Liang Wang, Nan Yang, and Furu Wei. 2023. [Query2doc: Query expansion with large language models](#). In *Conference on Empirical Methods in Natural Language Processing*.
- Zhepei Wei, Wei-Lin Chen, and Yu Meng. 2024. [Instrucrag: Instructing retrieval-augmented generation via self-synthesized rationales](#). *arXiv preprint arXiv:2406.13629*.
- Haoyi Xiong, Jiang Bian, Yuchen Li, Xuhong Li, Mengnan Du, Shuaiqiang Wang, Dawei Yin, and Sumi Helal. 2024. [When search engine services meet large language models: visions and challenges](#). *IEEE Transactions on Services Computing*.
- Shicheng Xu, Liang Pang, Mo Yu, Fandong Meng, Huawei Shen, Xueqi Cheng, and Jie Zhou. 2024. [Unsupervised information refinement training of large language models for retrieval-augmented generation](#). *ArXiv*, abs/2402.18150.
- Shijia Xu, Zhou Wu, Xiaolong Jia, Yu Wang, Kai Liu, and April Xiaowen Dong. 2026. [Self-correcting rag: Enhancing faithfulness via mmkp context selection and nli-guided mcts](#). *Preprint*, arXiv:2604.10734.
- Shi-Qi Yan and Zhen-Hua Ling. 2025. [Rpo: Retrieval preference optimization for robust retrieval-augmented generation](#). *ArXiv*, abs/2501.13726.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, and 41 others. 2025a. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.
- An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jingren Zhou, Junyang Lin, Kai Dang, and 22 others. 2024. Qwen2.5 technical report. *arXiv preprint arXiv:2412.15115*.
- Yuming Yang, Jiang Zhong, Li Jin, Jingwang Huang, Jingpeng Gao, Qing Liu, Yang Bai, Jingyuan Zhang, Rui Jiang, and Kaiwen Wei. 2025b. [Benchmarking multimodal rag through a chart-based document question-answering generation framework](#). *Preprint*, arXiv:2502.14864.
- Fuda Ye, Shuangyin Li, Yongqi Zhang, and Lei Chen. 2024. [R2ag: Incorporating retrieval information into retrieval augmented generation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 11584–11596.
- Ori Yoran, Tomer Wolfson, Ori Ram, and Jonathan Berant. [Making retrieval-augmented language models robust to irrelevant context](#). In *ICLR 2024 Workshop on Large Language Model (LLM) Agents*.
- Qiyang Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Tiantian Fan, Gaohong Liu, Lingjun Liu, Xin Liu, Haibin Lin, Zhiqi Lin, Bole Ma, Guangming Sheng, Yuxuan Tong, Chi Zhang, Mofan Zhang, Wang Zhang, Hang Zhu, and 16 others. 2025. [Dapo: An open-source llm reinforcement learning system at scale](#). *ArXiv*, abs/2503.14476.

W. Yu, Zhihan Zhang, Zhenwen Liang, Meng Jiang, and Ashish Sabharwal. 2023. [Improving language models via plug-and-play retrieval feedback](#). *ArXiv*, abs/2305.14002.

Qianchi Zhang, Hainan Zhang, Liang Pang, Yongxin Tong, Hongwei Zheng, and Zhiming Zheng. 2025. Less is more: Compact clue selection for efficient retrieval-augmented generation reasoning. *arXiv preprint arXiv:2502.11811*.

Qianchi Zhang, Hainan Zhang, Liang Pang, Hongwei Zheng, and Zhiming Zheng. 2026. Stable-rag: Mitigating retrieval-permutation-induced hallucinations in retrieval-augmented generation. *arXiv preprint arXiv:2601.02993*.

Xiaofan Zhou, Huy Nguyen, Bo Yu, Chenxi Liu, and Lu Cheng. 2026. Adaptive stopping for multi-turn llm reasoning. *arXiv preprint arXiv:2604.01413*.

Wei Zhu, Zhiwen Tang, and Kun Yue. 2026. Symphony: Synergistic multi-agent planning with heterogeneous language model assembly. *arXiv preprint arXiv:2601.22623*.

Appendix

A Test Dataset Statistics

A.1 Dataset.

We evaluate GRIP on **five** QA benchmarks as the main evaluation suite in the main paper, and **additionally** report results on **BioASQ** in this appendix as a domain-specific benchmark. HotpotQA focuses on multi-hop reasoning. PopQA, NQ, and WebQ are open-domain QA datasets, targeting factual questions with varying entity popularity and coverage. TriviaQA emphasizes reading comprehension from web and Wikipedia sources, featuring complex, compositional questions authored by trivia enthusiasts. Finally, BioASQ contains biomedical questions that require specialized knowledge and is *not* included in the main-suite average unless explicitly stated. These datasets collectively test the model’s adaptability to different retrieval and reasoning demands.

	#Samples	Type
HotpotQA	7,405	Multi-hop
PopQA	14,267	Open-domain
NQ	3,610	Open-domain
WebQ	2,032	Open-domain
TriviaQA	11,313	Reading Comprehension
BioASQ	885	Domain-specific

Table 9: Statistics and types of the test datasets used in our evaluation. The datasets vary in size and reasoning focus, covering multi-hop, open-domain, reading comprehension, and biomedical QA tasks.

A.2 Evaluation Metrics.

We report three automatic metrics on all QA benchmarks, where higher is better. For each test instance, we compute the metric against the corresponding reference answer(s), and then average scores over the evaluation set.

Exact Match (EM). EM is a binary metric that equals 1 if the predicted answer exactly matches any reference answer after standard normalization (e.g., lowercasing and trimming extra whitespaces), and 0 otherwise. The final EM is the average over all instances.

Token-level F1. F1 measures the token overlap between the prediction and a reference answer. Following the standard QA evaluation protocol, we compute precision and recall based on overlapping tokens, and the final F1 is averaged across instances.

ROUGE. We compute ROUGE- n for $n \in \{1, 2, L\}$ between the generated output and the reference text(s). We then report $\text{ROUGE} = \frac{1}{3}(\text{ROUGE-1} + \text{ROUGE-2} + \text{ROUGE-L})$.

B Performance On The Domain-specific Dataset

	ROUGE	F1	Avg.Score
FLARE	30.2	69.4	49.8
DRAGIN	47.1	81.2	64.2
SFT-RAG	39.1	58.6	48.9
Self-RAG	45.9	68.8	57.4
INFO-RAG	39.4	59.2	49.3
RobustRAG	53.9	80.9	67.4
GainRAG	49.3	74.0	61.7
GRIP	54.8	84.4	69.6

Table 10: Performance on the domain-specific BioASQ benchmark (higher is better). Avg.Score is the un-weighted mean of ROUGE and F1.

We further evaluate all models on BioASQ, a domain-specific QA benchmark where questions often require specialized biomedical knowledge beyond the parametric capacity of general-purpose LLMs. As shown in Table 10, GRIP achieves the best overall performance, obtaining the highest ROUGE (54.8) and F1 (84.4), and thus the highest Avg.Score (69.6). Compared with strong training-based baselines such as RobustRAG and GainRAG, GRIP yields more accurate and better-supported answers, suggesting that its self-triggered information planning can more reliably decide when to retrieve and how to formulate domain-relevant queries under high knowledge uncertainty.

C Cover Match Analysis for Answer Stability

To further assess semantic correctness beyond exact matching, we report CoverEM scores across all datasets. CoverEM is designed to measure whether the model output covers the reference answer string. After applying the same text normalization as EM, CoverEM equals 1 if the normalized prediction contains the normalized reference answer string as a contiguous substring (for any reference in the reference set), and 0 otherwise. The final CoverEM is averaged across instances.

As shown in Table 11, GRIP significantly outperforms all baselines, including powerful models like GPT-4o and GainRAG. This suggests that even when predictions diverge from the gold answer in

surface form, GRIP still generates semantically aligned spans with correct content. The consistent gains across datasets suggest that GRIP’s outputs are not only accurate but also robust to surface-level variations.

	Hotpot	PopQA	NQ	WebQ	Trivia	Avg. CoverEM
GPT-4o	38.6	45.2	48.5	43.6	82.1	51.6
DRAGIN	31.6	23.0	38.4	41.5	68.0	40.5
GainRAG	38.0	44.4	39.1	36.9	64.2	44.5
GRIP	45.0	49.1	41.2	52.4	74.9	52.5

Table 11: CoverEM scores across five QA datasets. GRIP achieves the best average CoverEM. Hotpot and Trivia denote HotpotQA and TriviaQA, respectively.

D Structured Dataset Samples

We now provide several dataset samples to help readers understand the structured training data easily.

```

Type-α Example 1
1 {
2   "Question": "Who is the actor that plays lip on shameless",
3   "Output": "[ANSWER] Jeremy Allen White [SOLVED]",
4   "Intermediate Answer": "",
5   "Retrieved Context": ""
6 }

Type-α Example 2
1 {
2   "Question": "Where is the national institute of health located",
3   "Output": "[ANSWER] Bethesda, Maryland [SOLVED]",
4   "Intermediate Answer": "",
5   "Retrieved Context": ""
6 }

```

Figure 4: **Type-α**: Samples that are directly answerable using the model’s internal knowledge alone, where the model emits [ANSWER] followed by the answer and concludes with [SOLVED] without any retrieval.

```

Type-β Example 1
1 {
2   "Question": "Who owns the westgate hotel in las vegas",
3   "Output": "[INTERMEDIARY] Westgate Las Vegas Resort & Casino is owned by Westgate Resorts. [R",
4   "Intermediate Answer": "",
5   "Retrieved Context": ""
6 }

Type-β Example 2
1 {
2   "Question": "Where did the battle of midway take place",
3   "Output": "[INTERMEDIARY] The Battle of Midway took place in the Pacific Ocean, near the Mid",
4   "Intermediate Answer": "",
5   "Retrieved Context": ""
6 }

```

Figure 5: **Type-β**: Samples where the model produces partial or noisy answers containing the gold answer but lacking clarity, encouraging it to recognize uncertainty and issue a retrieval request through [INTERMEDIARY] followed by [RETRIEVE].

```

Type-γ Example 1
1 {
2   "Question": "What was the character name of the Girl from Uncle?",
3   "Output": "[Intermediary] The character name of the Girl from Uncle is \April Dancer,\ who is featured in the television series \The Girl from U.N.C.L.E.\ [RETRIEVE] What is the name of the main character in the television series \The Girl from U.N.C.L.E.\?",
4   "Intermediate Answer": "Cathy Gale",
5   "Retrieved Context": "That Uncle Sam, as they call him, would certainly forever have lost the services of at least one soldier. The precise origin of the Uncle Sam character is unclear, but a popular legend is that the name "Uncle Sam" was derived from Samuel Wilson, a meatpacker from Troy, New York who supplied rations for American soldiers during the War of 1812. There was a requirement at the time for contractors to stamp their name and where the rations came from onto the food they were sending..."
6 }

Type-γ Example 2
1 {
2   "Question": "Big Ben is in which London building?",
3   "Output": "[Intermediary] Big Ben is located in the Elizabeth Tower at the north end of the Palace of Westminster in London. [RETRIEVE] What is the historical significance of the Elizabeth Tower in relation to Big Ben?",
4   "Intermediate Answer": "The Palace of Westminster",
5   "Retrieved Context": "Big Ben Big Ben is the nickname for the great Bell of the clock at the north end of the Palace of Westminster in London and is usually extended to refer to both the clock and the clock tower. The official name of the tower in which Big Ben is located was originally the Clock Tower, but it was renamed Elizabeth Tower in 2012 to mark the Diamond Jubilee of Elizabeth II. The tower was designed by Augustus Pugin in a neo-gothic style. When completed in 1859, ..."
6 }

```

Figure 6: **Type-γ**: Multi-hop or complex knowledge samples where both parametric memory and basic retrieval fail, requiring the model to iteratively construct better sub-questions and reason across multiple planning steps before concluding with [ANSWER] and [SOLVED].

E Various Control Tokens Settings

In the preliminary stage of GRIP, we explored multiple strategies for defining control tokens used to guide the decoding process. Specifically, we experimented with the following approaches:

- Using natural-language-style tokens such as [ANSWER] or [RETRIEVE] directly as special control tokens;
- Replacing control tokens with reserved but unused tokens from the pretrained model vocabulary, such as [ANSWER] replaced by <|reserved_special_token_51|>;
- Introducing four newly defined tokens into the tokenizer vocabulary and initializing them with embeddings of semantically related words. For example, [RETRIEVE] is initialized with the embedding of the word retrieve.

Empirical results showed that these strategies resulted in only marginal differences in overall model performance. To reduce implementation complexity and ensure compatibility with standard decoding processes, we selected the first approach for all subsequent experiments.

F Ablation on the Reward Metric

We replace the answer-fidelity reward r_{ans} with BLEU, ROUGE, token-F1, or EM while keeping

```

Type-θ Example 1
1 {
2   "Question": "Which opera by Rossini has the alternative title 'The Useless Precaution'?",
3   "Output": "[ANSWER] THE BARBER OF SEVILLE [SOLVED]",
4   "Intermediate Answer": "L'occasione fa il ladro",
5   "Retrieved Context": "natural process of comparing oneself with others. The early experimental work of the Gestaltists in Germany placed the beginning of problem solving study (e.g., Karl Duncker in 1935 with his book 'The psychology of productive thinking'). Later this experimental work continued through the 1960s and early 1970s with research conducted on relatively simple (but novel for participants) laboratory tasks of problem solving. The use of simple, novel tasks was due to the clearly defined..."
6 }

Type-θ Example 2
1 {
2   "Question": "Which duo had a Number Two hit in 1984 with 'Agadoo'",
3   "Output": "[ANSWER] BLACK LACE [SOLVED]",
4   "Intermediate Answer": "Johnny Logan and Foster & Allen",
5   "Retrieved Context": "the top 100 at number 64 on 16 August 1988, but only stayed in the chart for two weeks. The song has found success in other countries as well, such as Ireland, New Zealand and France. In the Irish Singles Chart, the song peaked at number 5, but only spent 5 weeks in the charts overall. 'Agadoo' spent a longer 11 weeks in the New Zealand Singles Chart, debuting at number 37 on 21 October 1984 and hitting a peak of number 9 in its sixth week. The song fell out the chart quite abruptly, ..."
6 }

```

Figure 7: **Type-θ**: Cases where the retrieved content covers the answer superficially but demands further reasoning and content compression, training the model to selectively synthesize key information via [INTERMEDIARY] → [RETRIEVE] → [ANSWER] transitions.

Reward metric for r_{ans}	Avg.Score	Avg.Count
BLEU (default)	41.0	1.24
ROUGE	40.5	1.24
F1	40.8	1.25
EM	40.6	1.24

Table 12: Reward-metric ablation for RL: we only change the overlap metric used in the answer-fidelity reward r_{ans} ; all other settings are fixed.

the RL pipeline unchanged. As shown in Table 12, performance is highly consistent across metrics (within 0.5 Avg.Score), and the average number of retrieval calls is nearly unchanged. This suggests that the RL stage is not sensitive to a specific overlap definition and primarily provides fine-grained shaping of answer quality. The near-identical Avg.Count indicates that the control reward dominates the retrieval policy, while r_{ans} mainly affects answer wording rather than when to retrieve. We use BLEU by default, as it offers a smoother graded signal than strict EM while achieving the best Avg.Score under essentially the same retrieval budget.

G Behavioral Shift after Rule-based RL

We analyze GRIP’s retrieval behavior before and after reinforcement learning (GRIP (w/o RL) vs. GRIP) on two representative datasets: NQ and WebQ. As shown in Figure 8, Rule-based RL significantly reduces unnecessary second retrieval steps, e.g., the proportion of Retrieve=2 on WebQ




Figure 8: Comparison of retrieval behavior distributions before and after reinforcement learning on NQ and WebQ.

drops from 59.6% to 32.5%, while Retrieve=1 rises markedly from 6.6% to 37.5%. This shift indicates that GRIP learns to make earlier stopping decisions once sufficient context is available, thereby avoiding over-retrieval. On NQ, where many questions can be answered directly, GRIP maintains a high Retrieve=0 rate (from 57.8% to 60.9%), confirming its ability to preserve efficient strategies on easier tasks. Overall, these results highlight GRIP’s capacity to adapt its retrieval depth based on the complexity of individual questions, demonstrating more interpretable and context-aware control over retrieval behaviors.

Overall, RL mainly converts many “Retrieve=2” cases into “Retrieve=1” rather than increasing “Retrieve=0”, suggesting it improves early stopping once sufficient evidence is obtained.

H Case Studies of GRIP’s Retrieval Planning Behavior

To better illustrate the internal behavior of GRIP during retrieval-augmented generation, we present a series of qualitative case studies that highlight how the model dynamically plans, regulates, and adapts its retrieval strategy. These examples cover representative scenarios including retrieval under information insufficiency, multi-step query refinement, and generation under explicit control signals. Each case demonstrates a different aspect of GRIP’s token-level decision-making process, providing deeper insights into its reasoning trajectory, interpretability, and interaction between generation and retrieval.

Case 1: Fallback to Internal Knowledge When GRIP Cannot Answer

To understand how GRIP behaves when external information is insufficient or retrieval is delayed, we examine a case where the model fails to answer the original question. Instead of hallucinating a con-

fidient response, GRIP generates partial reasoning steps based on its internal knowledge, revealing its awareness of information gaps. This demonstrates GRIP’s ability to initiate cautious reasoning before retrieval is triggered, contrasting with methods like DRAGIN that may prematurely commit to incomplete answers.

In both examples shown in Figure 9, GRIP is presented with questions for which no retrieved evidence is initially available. Rather than producing a low-confidence final answer, GRIP emits an [INTERMEDIARY] token followed by an internally composed partial response grounded in its parametric knowledge. For instance, when asked “what country was Slovakia?” or “who is Nebula on Guardians of the Galaxy?”, GRIP provides factual context (e.g., “Slovakia is a country in Central Europe” or “Nebula is a character in the Guardians franchise...”), before emitting [RETRIEVE] along with a reformulated sub-question. This pattern reflects GRIP’s ability to recognize uncertainty, enabling it to avoid hallucination while still making forward progress in reasoning. Such fallback behavior illustrates a core strength of GRIP: its capacity to delay final commitment until sufficient evidence is gathered.

```

Case 1: Fallback to Internal Knowledge When GRIP Cannot Answer
1 {
2   "input": [
3     "Question": "what country was slovakia?",
4     "Intermediate Answer": "",
5     "Retrieved Context": ""
6   ],
7   "GRIP's Output": [
8     "[INTERMEDIARY] Slovakia is a country in Central Europe. [RETRIEVE] what country was slovakia?",
9   ]
10 }

Case 1: Fallback to Internal Knowledge When GRIP Cannot Answer
1 {
2   "input": [
3     "Question": "who is nebula on guardians of the galaxy",
4     "Intermediate Answer": "",
5     "Retrieved Context": ""
6   ],
7   "GRIP's Output": [
8     "[INTERMEDIARY] Nebula is a character in the Guardians of the Galaxy franchise who is a former gangster and the adoptive daughter of Thanos [RETRIEVE] who is nebula on guardians of the galaxy",
9   ]
10 }

```

Figure 9: Example of GRIP falling back to internal knowledge and deferring retrieval when no context is available.

Case 2: GRIP Refines Queries Through Iterative Retrieval Planning

We present a multi-step example in which GRIP issues several [RETRIEVE] signals, each followed by a newly generated query. Initially, the retrieved content is insufficient to support a final answer. However, GRIP leverages its intermediate outputs to

progressively formulate more focused and precise new queries. This behavior illustrates the model’s capacity for retrieval planning: dynamically adjusting its information-seeking behavior based on context evolution and accumulated evidence.

In this case (Figure 10), we observe that the intermediate answer produced in the first and second rounds is nearly identical, indicating that the initial retrieval step failed to surface useful content for answering the original question. Despite this, GRIP does not prematurely conclude. Instead, it formulates a more specific and context-aware query in the second round: “What is the name of the artist performing in the halftime show for the upcoming Super Bowl?” This refined question enables the retriever to access a higher-quality passage containing the correct entity. In the third round, GRIP successfully emits [ANSWER] and [SOLVED], demonstrating that effective retrieval often hinges not on initial coverage but on the model’s ability to iteratively reformulate sub-questions until the right evidence is found. This case showcases GRIP’s retrieval planning mechanism in action, where small linguistic changes in the generated query lead to significant downstream improvements in answerability.

```

Case 2: GRIP Refines Queries Through Iterative Retrieval Planning
1 First-round answer
2 { "input": [
3   "Question": "who is playing the halftime show for the superbowl?",
4   "Intermediate Answer": "",
5   "Retrieved Context": ""
6 ],
7 "GRIP's Output": [
8   "[INTERMEDIARY] The performer for the halftime show of the Super Bowl is typically a well-known musical artist, but the specific artist for each year's show can vary. [RETRIEVE] who is playing the halftime show for the superbowl?"]
9 ]
10 }
11 Second-round answer
12 { "input": [
13   "Question": "who is playing the halftime show for the superbowl?",
14   "Intermediate Answer": "The performer for the halftime show of the Super Bowl is typically a well-known musical artist, but the specific artist for each year's show can vary.",
15   "Retrieved Context": "The Super Bowl Halftime Show is one of the most-watched musical performances of the year, featuring major artists from across genres..."
16 ],
17 "GRIP's Output": "[INTERMEDIARY] The performer for the halftime show of the Super Bowl is a well-known musical artist, but the specific artist varies each year. [RETRIEVE] What is the name of the artist performing in the halftime show for the upcoming Super Bowl?"]
18 }
19 Third-round answer
20 { "input": [
21   "Question": "who is playing the halftime show for the superbowl?",
22   "Intermediate Answer": "The performer for the halftime show of the Super Bowl is a well-known musical artist, but the specific artist varies each year.",
23   "Retrieved Context": "The halftime show performer for the upcoming Super Bowl has been officially announced by the NFL. This year, the headline act will be Usher, a Grammy-winning R&B singer known for his chart-topping hits and dynamic stage presence..."
24 ],
25 "GRIP's Output": "[ANSWER]Justin Timberlake [SOLVED]"
26 }

```

Figure 10: Example of GRIP refining its retrieval query over multiple rounds to obtain a better answer.

Case 3: GRIP’s Summarization Behavior by Answering with [ANSWER]

In this case, we investigate how GRIP responds when explicitly instructed to finalize an answer via the [ANSWER] control token. Despite not having retrieved all necessary evidence, the model inte-

grates its current reasoning trajectory and partial observations to generate a coherent summary response. This highlights GRIP’s ability to adapt to control signals and perform best-effort answering under constrained or premature termination, showcasing a key advantage of token-level behavioral regulation.

Figure 11 illustrates a chained reasoning trajectory where GRIP first generates an incorrect intermediate answer based on parametric memory (1880), then revises it after retrieving a more accurate year (1871) from external content. Rather than concluding immediately, the model goes one step further: it reformulates the question to explore the historical significance of the event, triggering another retrieval round. Finally, GRIP compresses the relevant passage into a concise factual response (1871 A.D.) using the [ANSWER] and [SOLVED] tokens. This case demonstrates GRIP’s ability to integrate evidence across turns, refine its beliefs based on new context, and summarize multi-step reasoning into a final answer—without hallucinating or skipping steps. It also showcases how retrieval can reshape the reasoning trajectory itself, not just inform it.

```

Case 3: Case 3: GRIP's Summarization Behavior with [ANSWER]
1 First-round answer
2 { "Input": [
3   "Question": "when was the debating club established in almora?",
4   "Intermediate Answer": "",
5   "Retrieved Context": "",
6   "GRIP's Output": [
7     "[[INTERMEDIARY] The debating club in Almora was established in the year 1888 [RETRIEVE]
8     when was the debating club established in almora?]]"
9 Second-round answer
10 { "Input": [
11   "Question": "when was the debating club established in almora?",
12   "Intermediate Answer": "The debating club in Almora was established in the year 1888.",
13   "Retrieved Context": "The Debating Club in Almora has long been recognized as a prominent
14   intellectual forum in the region, fostering public speaking, critical thinking,..."
15   "GRIP's Output": ["[[INTERMEDIARY] The debating club in Almora was established in 1871
16   A.D. by Pt. Buddhiballav Pant [RETRIEVE] what is the significance of the debating club
17   established in Almora in 1871?]]"]
18 Third-round answer
19 { "Input": [
20   "Question": "What is the significance of the debating club established in Almora in
21   1871?",
22   "Intermediate Answer": "The performer for the halftime show of the Super Bowl is a well-
23   known musical artist, but the specific artist varies each year.",
24   "Retrieved Context": "The Almora Debating Club, founded in the late 19th century, quickly
25   became a central institution for intellectual and political discourse in the Kumaon region.
26   It served as a platform where students, professionals,...",
27   "GRIP's Output": ["[ANSWER]1871 A.D. [SOLVED]"]

```

Figure 11: Example of GRIP refining, reformulating, and compressing retrieved knowledge into a final answer using [ANSWER].

I Failure Analysis

This section presents selected failure cases of GRIP to better understand its limitations in retrieval-augmented generation. While GRIP exhibits strong performance across benchmarks, we observe occasional failure patterns that reveal important areas

for improvement. Specifically, we analyze representative cases involving incorrect final answers and redundant retrieval behaviors.

Case 1: Incorrect Final Answer

In some cases, GRIP successfully retrieves relevant evidence but fails to generate the correct final answer. This often occurs when the retrieved passages contain implicit or multi-faceted information that requires deeper logical inference. These failures highlight limitations in GRIP’s reasoning capability under answer ambiguity or insufficient summarization, even when the retrieval component functions as intended.

Case 2: Redundant Retrieval Steps

GRIP’s token-level control enables flexible retrieval depth, but in certain examples, the model triggers unnecessary retrievals that do not improve or even dilute answer quality. Such cases typically occur when the model fails to recognize that sufficient information is already available, resulting in redundant sub-queries and noisy context accumulation. These examples illustrate the challenges of balancing retrieval sensitivity and reasoning sufficiency, especially in high-recall settings.

J Latent State Visualization of Control Decisions



Figure 12: PCA visualization of the model’s hidden states before emitting [ANSWER] (blue) and [RETRIEVE] (orange) on the NQ dataset. Each point corresponds to the hidden representation of the last token in the input prompt. The distinct clusters indicate that the model encodes separable internal states for different control decisions.

To better understand the model’s internal decision dynamics, we visualize the hidden representations associated with the two control tokens: [ANSWER] and [RETRIEVE]. For each sample, we

Method	BM25	DPR	Hybrid
Single-RAG (training-free)	30.8	31.9	33.7
RobustRAG (training-based)	33.2	32.1	33.8
GRIP	41.0	39.2	41.0

Table 13: Retriever ablation on the main evaluation suite. We replace the default BM25 retriever with a dense DPR retriever and a BM25+DPR hybrid retriever (denoted **Hybrid**), while keeping the backbone (LLaMA3-8B), corpus (Wikipedia), and top- k passages ($k=3$) fixed. Single-RAG and RobustRAG are representative training-free and training-based baselines from the main table, respectively.

extract the hidden state of the last token in the input prompt, immediately before the model emits a control token. These hidden states represent the model’s contextual understanding just prior to triggering a retrieval or producing an answer. We then apply principal component analysis (PCA) to reduce the hidden states to two dimensions and plot the results in Figure 12.

The visualization reveals a clear separation between the two control behaviors in the latent space. Samples that lead to [RETRIEVE] tend to cluster in a distinct region from those that lead to [ANSWER]. This suggests that the model has learned to encode different latent states for retrieval versus answer emission, reinforcing the effectiveness of control-token-based behavioral modulation.

K Prompts Used in Structured Data Generation

As part of our data generation pipeline, we employ GPT-based prompting to simulate how a model might revise its internal reasoning and generate improved follow-up queries. Given the original question, a partial (intermediary) answer, and a set of retrieved documents, the prompt instructs the model to refine the known fact and formulate a more targeted query that increases the likelihood of successful retrieval. This process is used to construct structured training data that supports GRIP’s ability to perform iterative retrieval planning in a self-supervised fashion.

L Retriever Ablation

To assess how sensitive our conclusions are to the choice of retriever, we evaluate GRIP and two representative baselines under three retrieval backends: sparse BM25, dense DPR, and a BM25+DPR hy-

System:

Given:

- question: the original user question
- retrieval: 3 context list of retrieved documents
- reference: the model’s first attempt at a partial answer

Task:

1. Re-summarize the intermediary fact that is most useful for answering the question, by combining the question, the retrieved documents, and initial_intermediary.
 - If initial_intermediary was incorrect or incomplete, discard or correct it.
 - Verify whether any retrieved document contains incorrect or misleading information based on your own knowledge, and correct the intermediary fact accordingly.
 - If it was correct, you may refine it.
2. Based on that finalized intermediary fact, plus the question and retrieved documents, generate a new search query that is highly likely to return evidence needed for the full answer.

Output:

Produce exactly one line in the following format, with no extra text: [Intermediary] <your refined known fact> [RETRIEVE] <your new query>

Question:

{question}

Retrieved docs:

{docs_str}

Initial_intermediary:

{initial_intermediary}

Your output:

Figure 13: Prompt used for data generation with GPT-4o-mini.

brid retriever. We select Single-RAG as a representative training-free method and RobustRAG as a representative training-based method. For each retriever, we keep all other factors unchanged, including the LLaMA3-8B backbone, the Wikipedia corpus, and the number of retrieved passages ($k=3$).

Table 13 shows that hybrid retrieval consistently improves performance for both Single-RAG and RobustRAG, indicating that combining sparse and dense signals can provide higher-quality evidence than either retriever alone. For GRIP, DPR alone is slightly worse than BM25, while the hybrid retriever matches the best result. Overall, GRIP remains substantially stronger than both baseline families across all retriever choices, suggesting that its gains stem from the proposed token-level planning and training recipe rather than a particular retriever implementation. In the main paper we adopt BM25 as the default retriever for efficiency and controlled comparison, and report these additional retriever variants here for completeness.

M Pseudocode of GRIP Decoding with Self-Triggered Information Planning

The following pseudocode outlines the decoding procedure of GRIP, which integrates retrieval control into autoregressive generation through self-triggered information planning. It illustrates how the model emits control tokens to decide when to retrieve, generate sub-questions, and finalize answers during the decoding process.

Algorithm 1 GRIP Decoding with Self-Triggered Information Planning

Input: query x , model π_θ , retriever \mathcal{R} , max steps B
Output: final answer y

- 1: Initialize $y \leftarrow \emptyset$, $m \leftarrow \emptyset$, $b \leftarrow 0$
- 2: **while** $b < B$ **do**
- 3: Sample $t \sim \pi_\theta(x, m, y)$
- 4: Append t to y
- 5: **if** $t = [\text{ANSWER}]$ **then**
- 6: Generate final answer until [SOLVED]
- 7: **return** y
- 8: **else if** $t = [\text{INTERMEDIARY}]$ **then**
- 9: Generate partial answer a , then generate next token t'
- 10: **if** $t' = [\text{RETRIEVE}]$ **then**
- 11: Generate new query q
- 12: Retrieve $c \leftarrow \mathcal{R}(q)$
- 13: $m \leftarrow m \cup \{c\}$; $b \leftarrow b + 1$
- 14: **else**
- 15: Append fallback message
- 16: **end if**
- 17: **else if** $t = [\text{SOLVED}]$ **then**
- 18: **return** y
- 19: **end if**
- 20: **end while**
- 21: **return** y

N Retrieval-Count Distributions under Budget Control

We vary the inference-time maximum retrieval budget $B \in \{3, 5, 7, 10\}$ and record the realized number of retrieval calls r per example. Figure 14 reports the empirical distribution of r on each benchmark. The dashed vertical line highlights $r=3$, which matches the maximum retrieval depth used in our constructed training trajectories. Across all datasets, GRIP does not saturate the budget: even when B is large, most probability mass remains on small r (typically $r \leq 2$), indicating that retrieval

is triggered selectively when the decoding trajectory suggests an information gap, instead of being executed as a fixed schedule.

Importantly, the figure also demonstrates *extrapolative generalization* beyond the training depth. Our training trajectories are constructed with at most three retrieval steps, yet increasing B at test time yields non-trivial mass at $r > 3$. At $B=10$, the fraction of examples with $r > 3$ reaches 17.0% on WebQ, 11.7% on NQ, 7.8% on PopQA, 7.3% on HotpotQA, and 3.4% on TriviaQA. This shows that GRIP learns a transferable decision principle of *when additional evidence is needed*, rather than overfitting to a fixed maximum-step pattern seen during training.

Finally, the extent of extrapolation is task-dependent. Datasets that are more evidence-uncertain exhibit heavier tails beyond $r=3$ (e.g., WebQ and NQ), whereas factoid-style benchmarks concentrate more mass at $r=1$ or $r=2$ and rarely exceed the training depth (e.g., TriviaQA). Overall, budget control provides a practical mechanism to trade retrieval cost for accuracy while preserving GRIP’s adaptive behavior beyond the retrieval depth seen in training.

O DAPO Configuration for the RL Stage

We adopt the official DAPO recipe from the VERL codebase and keep its core algorithmic configuration unchanged; our modifications are limited to the data configuration and our task-specific reward definitions described in §2.4.² In particular, we enable CLIP-HIGHER, token-level policy gradient loss, dynamic sampling, and soft overlong reward shaping, and disable KL regularization (both KL-in-reward and KL loss).

Soft overlong punishment is enabled and activates when the response length exceeds `overlong_buffer.len`. Dynamic sampling is enabled as in the default recipe to encourage advantage diversity during RL training.

P Results with Qwen2.5-7B

To verify that our findings are not specific to a single backbone, we additionally evaluate GRIP and the reproduced baselines using **Qwen2.5-7B** under the same evaluation protocol. In particular, we keep the external corpus (Wikipedia), retriever (BM25), and top- k passages ($k=3$) unchanged, and

²https://github.com/volcengine/verl/blob/main/recipe/dapo/run_dapo_qwen2.5_32b.sh

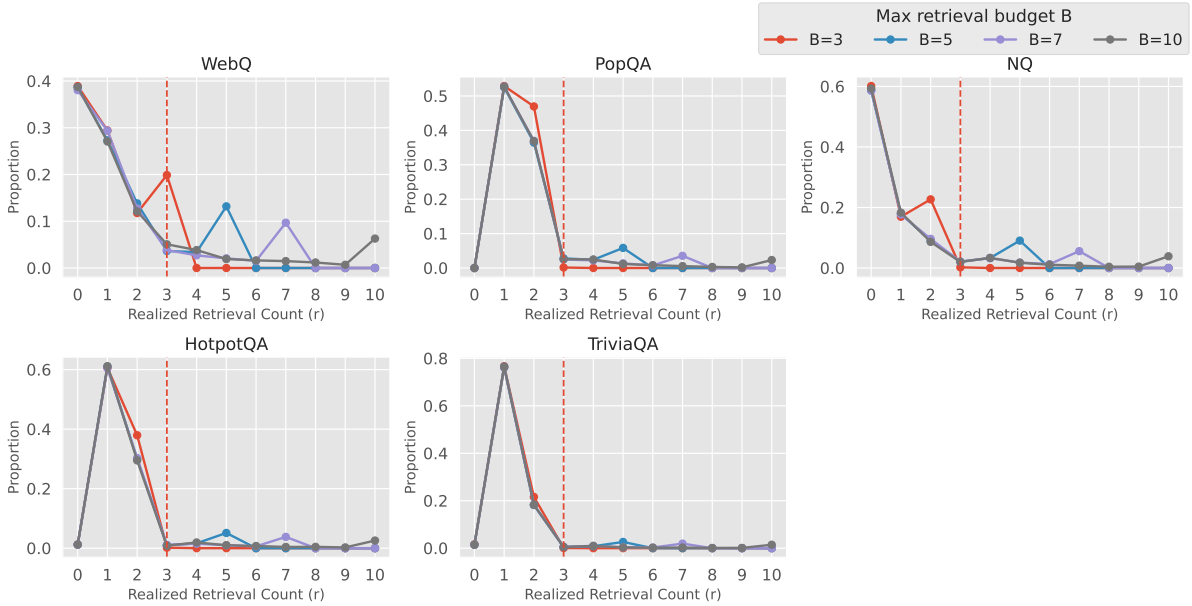


Figure 14: Realized retrieval-count distributions of GRIP under different maximum retrieval budgets B . Each subplot corresponds to one benchmark; curves show the fraction of examples that trigger r retrieval calls at inference time. The dashed vertical line marks $r=3$, the maximum retrieval depth observed in our training trajectories. As B increases, GRIP places more probability mass beyond $r>3$, indicating extrapolation beyond the training depth, while the distribution remains concentrated at small r , suggesting selective and task-aware retrieval rather than mechanically consuming the available budget.

Component	Status	Key settings
CLIP-HIGHER	on	clip_ratio_low=0.2, clip_ratio_high=0.28, clip_ratio_c=10.0
Token-level PG loss	on	loss_agg_mode=token-mean
Dynamic sampling	on	enabled (default DAPO setting)
Soft overlong shaping	on	overlong_buffer.len=8192, penalty_factor=1.0
KL in reward	off	use_kl_in_reward=False, kl_coef=0.0
KL loss	off	use_kl_loss=False, kl_loss_coef=0.0

Table 14: Key DAPO components used in our RL stage. All unspecified hyperparameters follow the referenced DAPO recipe.

apply the same decoding and scoring procedure as in the main experiments.

The results on Qwen2.5-7B exhibit the *same trend* as those on LLaMA3-8B: GRIP consistently outperforms both training-free and training-based RAG baselines across datasets, while maintaining strong retrieval efficiency through token-level planning. The relative ranking among baseline methods is also stable, indicating that the gains of GRIP stem from the proposed retrieval-as-generation formulation and scenario-typed supervision rather than backbone-specific effects. We further observe that the RL stage provides an additional but modest improvement on Qwen2.5-7B, consistent with our main findings that RL primarily refines retrieval behavior and stabilizes termination decisions.

	HotpotQA			PopQA			NQ			WebQ			TriviaQA			Avg. Score
	EM	ROUGE	F1	EM	ROUGE	F1	EM	ROUGE	F1	EM	ROUGE	F1	EM	ROUGE	F1	
Qwen-3-4B-Instruct																
Instruct	8.0	13.6	16.5	5.3	8.0	10.3	6.2	11.3	13.7	7.7	16.3	20.0	22.1	24.1	30.1	14.2
Single-RAG	18.4	24.6	28.8	18.3	21.7	25.5	12.9	18.9	21.9	8.7	16.7	20.2	36.0	38.0	46.0	23.8
Robust-RAG	27.3	30.8	36.7	27.3	28.2	33.2	20.9	23.8	28.1	19.8	24.5	28.7	47.7	45.2	55.1	31.8
R1-Searcher	18.2	21.2	25.4	12.5	14.4	18.0	11.8	16.8	20.1	12.9	21.6	25.4	36.7	35.1	43.5	22.3
GRIP	31.9	39.7	46.2	27.8	30.9	36.6	21.1	28.6	33.2	16.8	27.5	32.3	49.9	48.2	59.1	35.3
Qwen2.5-7B-Instruct																
Instruct	18.3	21.5	26.0	11.5	13.0	16.3	11.8	16.7	20.0	14.5	22.6	26.4	37.1	36.0	44.4	22.4
Single-RAG	30.6	34.9	40.9	23.6	26.6	31.1	18.9	24.6	28.3	15.3	22.6	26.3	45.8	46.1	55.7	31.4
Robust-RAG	26.1	30.8	35.5	27.5	28.6	33.5	20.5	23.8	27.8	21.1	25.8	30.0	48.7	46.7	56.7	32.2
R1-Searcher	20.2	24.3	28.8	11.6	13.8	17.0	14.3	19.2	23.0	19.3	26.7	30.7	40.1	39.4	48.3	25.1
GRIP	32.0	37.6	44.6	28.7	30.9	36.3	21.2	27.7	32.5	19.3	28.3	33.2	51.9	49.6	60.5	35.6

Table 15: Main results on five QA benchmarks under two backbones. We report EM, ROUGE, F1, and Avg. Score.