

RAGFort: Dual-Path Defense Against Proprietary Knowledge Base Extraction in Retrieval-Augmented Generation

Qinfeng Li^{1*}, Miao Pan^{1*}, Ke Xiong^{1*}, Ge Su¹,
Zhiqiang Shen², Yan Liu², Sun Bing³, Hao Peng⁴, Xuhong Zhang^{1,5†}

¹Zhejiang University

²Ant Group

³Universal Identification Technology (Hangzhou) Co., Ltd

⁴Zhejiang Normal University

⁵Ningbo Global Innovation Center, Zhejiang University

Abstract

Retrieval-Augmented Generation (RAG) systems deployed over proprietary knowledge bases face growing threats from reconstruction attacks that aggregate model responses to replicate knowledge bases. Such attacks exploit both intra-class and inter-class paths—progressively extracting fine-grained knowledge within topics and diffusing it across semantically related ones, thereby enabling comprehensive extraction of the original knowledge base. However, existing defenses target only one path, leaving the other unprotected. We conduct a systematic exploration to assess the impact of protecting each path independently and find that joint protection is essential for effective defense. Based on this, we propose RAGFort, a structure-aware dual-module defense combining *contrastive reindexing* for inter-class isolation and *constrained cascade generation* for intra-class protection. Experiments across security, performance, and robustness confirm that RAGFort significantly reduces reconstruction success while preserving answer quality, offering comprehensive defense against knowledge base extraction attacks.

Code — <https://github.com/happywinder/RAGFort>

Extended version — <https://arxiv.org/abs/2511.10128>

1 Introduction

Large Language Models (LLMs) have demonstrated remarkable capabilities across a wide range of natural language processing tasks, particularly in open-domain question answering and knowledge reasoning (Brown et al. 2020; OpenAI 2023). To enhance factual accuracy and controllability, Retrieval-Augmented Generation (RAG) has emerged as a widely adopted paradigm (Lewis et al. 2020; Izacard and Grave 2021), enabling models to retrieve external knowledge to support answer generation. Beyond public knowledge like Wikipedia, RAG systems are increasingly deployed in high-value domains, such as healthcare (Singhal

et al. 2022), and finance (Wu et al. 2023), where the knowledge base is proprietary, domain-specific, and constitutes a core intellectual asset for commercial applications.

However, the privatized nature of these knowledge bases exposes a new vulnerability (Ni et al. 2025). Specifically, recent studies reveal that even by interacting with the RAG system in a black-box manner, attackers can extract a proxy knowledge base by aggregating responses from the RAG system (Zeng et al. 2024a; Cohen, Bitton, and Nassi 2024). In particular, attackers may utilize the extracted knowledge base to construct functionally equivalent RAG systems (Jiang et al. 2024), thereby effectively replicating and misusing the original RAG system’s capabilities. Such replication enables competitors to offer highly similar services, leading to direct financial losses and substantially eroding the competitive edge of the original provider. We refer to this threat as the *knowledge base extraction attack*.

Specifically, recent studies show that attackers can extract information from a knowledge base by collecting and combining responses generated by the RAG system. For example, document extraction attacks have shown that attackers can craft queries that specifically target certain information. By prompting the model with these carefully designed queries, attackers can induce the RAG system to generate responses that contain or reveal underlying knowledge base content (Zeng et al. 2024a). More alarmingly, some emerging methods pose a more systemic extraction of the entire knowledge base rather than focusing on only isolated pieces of content. To achieve this, these attacks not only iteratively refine queries to extract detailed information within a specific targeted topic (referred to as *intra-class extraction*), but also gradually expand their query scope to cover other semantically related topics. Specifically, attackers leverage previously extracted knowledge to recursively generate queries targeting new semantically related topics. By iteratively expanding the extraction boundary, the attack gradually propagates through the knowledge base, eventually achieving nearly complete coverage (referred to as *inter-class extraction*). For example, RAG-Thief (Jiang et al. 2024) employs an automated agent that starts with

*These authors contributed equally.

† Corresponding to zhangxuhong@zju.edu.cn
Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

an initial adversarial query, stores the extracted chunks in memory, and then analyzes these chunks to generate more targeted queries for uncovering missing or previously unknown information. By iteratively updating its memory, the agent can both fill in missing details within the same topic (i.e., intra-class extraction) and gradually expand extraction to neighboring topics that were previously unknown or unqueried (i.e., inter-class extraction), ultimately reconstructing nearly the entire knowledge base.

However, existing defense strategies are still at a preliminary stage, typically focusing on individual attack paths rather than providing comprehensive protection. For example, intra-class defenses, such as paraphrasing or summarizing retrieved content, reduce leakage within individual chunks but fail to prevent cross-topic aggregation (Zeng et al. 2024a,b). Conversely, inter-class defenses restrict retrieval to closely related content, limiting topic drift but allowing in-depth extraction within a single class (Zeng et al. 2024a), such as limiting retrieval to the few most relevant chunks. Worse still, existing intra-class and inter-class defense schemes are difficult to combine effectively. For example, *Set Distance Threshold* (Zeng et al. 2024a), as an inter-class protection, relies on strictly limiting the retrieval distance so that only highly relevant chunks are returned, which essentially forces the RAG system to **focus on only a few chunks** and prevents extraction from expanding to other topics. In contrast, to achieve intra-class protection, *Summarization with Relevant Query* (Zeng et al. 2024a) emphasizes generating diverse or aggregated content across multiple chunks within the same topic (i.e., through abstractive summarization) to prevent attackers from reconstructing the precise information of any single chunk, and thus inherently requires the RAG system **not to focus on only a few chunks**. As a result, this fundamental conflict poses significant challenges for combining intra-class and inter-class defenses in practice. Nevertheless, no prior work has systematically compared the two defense paths or explored integrated strategies that jointly secure both protection paths.

To bridge this gap, we first pose a key question: *Between intra-class and inter-class protections, which is more critical for safeguarding proprietary knowledge base in RAG systems, and can integrated protection provide stronger security than either path alone?* To explore this, we design a targeted experiment that masks portions of knowledge along each attack path, and assess the resulting impact on answer quality. Our findings reveal a key insight: protecting only one path offers limited resistance, whereas joint protection along both paths reduces the effectiveness of the proxy knowledge base. These results indicate that intra-class robustness and inter-class isolation are complementary, and that effective defense requires their coordination.

Based on this insight, we propose a structure-aware dual-module defense framework, RAGFort, as shown in Figure 2. For inter-class protection, we introduce a structure-aware reindexing strategy that reorganizes the retriever’s dense index to enhance semantic separability between topic classes. This transformation aims to reinforce topic boundaries, making it harder for attackers to expand queries from one class to another, thereby preventing inter-class extrac-

tion attacks. Specifically, the method begins with unsupervised HDBSCAN clustering to identify latent topic structures, followed by supervised contrastive learning to train an embedding encoder that preserves inter-category margins.

For intra-class protection, we propose a cascaded generation framework that combines a lightweight draft model with a robust reference model. In this framework, the draft model first proposes candidate tokens, while the reference model selectively filters or replaces high-risk tokens. Specifically, when the draft model exhibits excessive uncertainty or its output contains sensitive content, the reference model intervenes and generates safer tokens. Essentially, we have designed a near-optimal and efficient rejection strategy that fuses the output distributions of the draft model and the reference model. The new distribution significantly reduces the probability of sensitive content being output and improves the model’s inference performance.

We validate the RAGFort through comprehensive experiments on security, system performance, and ablation studies. For security, we evaluate RAGFort under advanced black-box extraction attacks and show that it reduces the proportion of knowledge base entries recovered by attackers by more than half compared to prior defenses. For system performance, RAGFort consistently preserves high answer quality while also maintaining generation efficiency, thus ensuring practical usability. Through extensive ablation, we show that each module contributes to knowledge protection, and their combination delivers strong, complementary effects. Overall, our contributions are summarized as follows:

- To the best of our knowledge, this is the first systematic study that defends against knowledge base extraction in RAG systems. This overlooked vulnerability poses significant threats in real-world deployments and leads to significant commercial risks for RAG providers.
- We identify and formalize two critical paths to protect the RAG knowledge base: inter-class isolation and intra-class protection. We design a targeted experiment to validate the role of each path and demonstrate that their combined defense is essential for protecting knowledge base extraction attacks.
- We propose a structure-aware dual-module defense that jointly addresses inter-class and intra-class threats. It combines contrastive reindexing to enhance topic-level isolation and a cascade framework to suppress sensitive content generation under adversarial prompts.
- We conduct a comprehensive empirical evaluation of RAGFort. Our evaluation demonstrates that RAGFort significantly enhances knowledge base security while preserving RAG system performance. Ablation studies confirm that each module is effective on its own and further enhances overall protection when combined.

2 Preliminaries

Knowledge Base Extraction Attack. RAG systems generate answers by first retrieving relevant chunks from an external knowledge base via a retriever module and then using a language model to compose the final response. While




Figure 1: Dual-Path Protection Diagram. This diagram illustrates the structure-aware dual-module defense framework (RAG-Fort) for protecting data across vertical and horizontal axes.

RAG systems enhance factual generation by retrieving external knowledge, they also expose proprietary content to extraction attacks. Specifically, recent studies show that adversaries can reconstruct knowledge bases by collecting and combining RAG responses to carefully crafted or paraphrased queries. By systematically piecing together related fragments returned by the system, attackers are able to recover large portions of the original content (Ni et al. 2025; Zeng et al. 2024a; Jiang et al. 2024). Emerging methods further streamline this process by using adaptive querying agents that automatically generate and refine queries to explore both within individual topics (intra-class) and across different topics (inter-class) (Jiang et al. 2024; Cohen, Bitton, and Nassi 2024). This two-path strategy significantly improves the fidelity and coverage of the extraction.

Existing Defense. Existing defense strategies remain at an early stage, limited to a single path, with few truly systematic approaches available. *Intra-class* methods paraphrase or summarize individual chunks to obscure fine-grained content, or replace them altogether with synthetic data (Zeng et al. 2024a,b), while *inter-class* defenses restrict retrieval to highly relevant chunks via distance thresholds, implicitly masking broader topical structures (Zeng et al. 2024a). However, current approaches address only one direction of attack in isolation, leaving RAG systems vulnerable to coordinated extraction strategies.

Cascade Generation. We adopt a two-stage cascade generation strategy (Narasimhan et al. 2024) to regulate knowledge exposure. A lightweight draft model first proposes candidate tokens, which are then validated or replaced by a stronger verifier model. This parallel evaluation ensures both efficiency and quality. Unlike prior work focused on inference speed or quality cascading, we refine the rejection rules to minimize sensitive knowledge retrieval during generation.

Threat Model. In this paper, we consider two parties: the defender and the attacker. Unlike privacy attacks that target individual records, we consider attackers whose objective is to extract the entire underlying knowledge base of the RAG system. The attacker is restricted to a black-box setting and can only access the RAG system via API queries, which aligns with real-world deployments of commercial and enterprise RAG services. Defenders are system providers with full access to the RAG pipeline (including retrievers, databases, and language models). The defense objective is to protect proprietary knowledge bases from reconstruction while ensuring usability for legitimate users.

3 Comparison of Different Defense Paths

To investigate which protection, intra-class or inter-class, is more critical for safeguarding proprietary knowledge in RAG systems, and whether integrated protection offers stronger security than either path alone, we design an experiment that evaluates the impact of different protection paths by selectively masking portions of the knowledge base and analyzing the residual utility of protected knowledge bases.

3.1 Experimental Design

We begin with a complete knowledge base \mathcal{K}_{full} constructed from domain-specific datasets. From this base, we derive three partially leaked variants by systematically removing knowledge portions considered well-protected, which are thus excluded from the proxy knowledge bases. The remaining parts constitute the knowledge bases that the attacker has successfully reconstructed through model stealing.

- **Without Protection:** No entries are removed or masked; the attacker obtains the full original knowledge base without any restrictions.
- **Inter-class Protection:** Removes detailed entries within each category (e.g., specific conditions under "cardiology"), simulating intra-class protection.
- **Intra-class Protection:** Removes entire topic categories (e.g., "dermatology" or "algebra"), simulating inter-class protection.
- **Joint Protection:** Applies both inter-class and intra-class masking to simulate two-path coordinated defense.

Each variant is used to build an independent RAG system, and we evaluate its ability to correctly answer questions. *Higher Answer ACC* indicates that the proxy knowledge base still retains substantial utility, suggesting that the corresponding protection strategy is *less effective*. The detailed experimental setup is described in Section 5.1.

3.2 Results and Analysis

Figure 1 presents the accuracy of attacker-constructed RAG systems under each protection path across three representative domains. The results reveal a clear trend: either inter-class or intra-class protection alone is insufficient to significantly limit knowledge base leakage. By contrast, the joint strategy substantially reduces the utility of the stolen knowledge base. Specifically, the performance of the proxy knowledge base declines as the proportion of protected knowledge increases, and this decline is much more pronounced

with joint protection than with either single-path strategy. This difference grows even further as more data is protected. These findings empirically confirm our central hypothesis: *inter-class and intra-class protection are complementary, and only by combining both protection paths can protection achieve comprehensive and effective defense.*

4 Our Design: RAGFort

In this section, we introduce RAGFort, a structure-aware dual-path defense framework to counter knowledge base reconstruction attacks in RAG systems.

4.1 Overview

RAGFort is achieved by integrating a contrastive reindexing mechanism with a constrained generation cascade, each targeting a different path of the knowledge base, as shown in Figure 2. Specifically, for *inter-class protection*, we restructure the retriever’s index to enforce stronger semantic separation across topic categories. This process begins with unsupervised clustering using HDBSCAN to reveal latent topic structures. Then, we train a supervised contrastive encoder that preserves inter-cluster margins, enabling better class discrimination in the embedding space. As a result, it becomes significantly harder for attackers to retrieve content from unrelated topics.

For *intra-class protection*, we develop a cascade generation mechanism that improves both generation quality and defense robustness. This design combines a fast draft model with a stronger verifier model. During inference, the draft model proposes candidate tokens, which are selectively validated by the reference model based on a rejection rule. To mitigate the generation of sensitive content in cascaded systems, we propose a computationally efficient rejection rule and provide theoretical guarantees that its performance closely approximates that of the optimal rule.

4.2 Inter-Class Protection

Our inter-class protection method proceeds in three stages: unsupervised clustering, supervised contrastive learning, and index replacement.

Pseudo-Labeling via HDBSCAN Clustering. We begin by applying unsupervised clustering to identify latent topic categories within the knowledge base. First, each chunk $k_i \in \mathcal{K} = k_1, \dots, k_N$ is embedded using a pre-trained encoder (e.g., Sentence-BERT) to obtain dense representations $E(k_i)$. Next, we use Hierarchical Density-Based Spatial Clustering of Applications with Noise (HDBSCAN) (Campello, Moulavi, and Sander 2013) to partition the chunks into C clusters $\mathcal{C}_1, \dots, \mathcal{C}_C$, thereby assigning a pseudo-label $y_i \in 1, \dots, C$ to each k_i . Notably, HDBSCAN is chosen due to its capability to automatically infer the number of clusters, handle outliers, and adapt to high-dimensional semantic spaces.

Learning Class-Separable Embeddings via SupCon Loss. Using the pseudo-labels from clustering, we train a new encoder with supervised contrastive learning (SupCon)

to produce class-discriminative embeddings. Given a mini-batch of encoded chunks $\{z_1, \dots, z_B\}$ with corresponding cluster labels $\{y_1, \dots, y_B\}$, the SupCon loss is defined as:

$$\mathcal{L}_{\text{sup}} = \sum_{i=1}^B -\frac{1}{|P(i)|} \sum_{p \in P(i)} \log \frac{\exp(\text{sim}(z_i, z_p)/\tau)}{\sum_{a \in A(i)} \exp(\text{sim}(z_i, z_a)/\tau)}, \quad (1)$$

where $P(i)$ is the set of samples with the same label as i (excluding i), $A(i)$ includes all samples in the batch except i , and $\text{sim}(\cdot, \cdot)$ denotes cosine similarity. τ is the temperature scaling factor. This contrastive objective strengthens the separation between different semantic categories, improving the structural integrity of indexed representations.

Reindexing with Tokenizer-Aligned Contrastive Embeddings. After contrastive training, we obtain a structure-aware encoder $f_{\text{sup}}(\cdot)$ that outputs enhanced chunk embeddings $z'_i = f_{\text{sup}}(k_i)$. We then construct a new dense index for the retriever using these updated representations. Notably, these embeddings are generated based on a custom tokenizer aligned with the SupCon training process.

To preserve the generative performance of the original RAG system, we do not modify the query encoder or the decoder input pipeline. Instead, during inference, the retriever operates on the SupCon-based index, while the generator still consumes the original (non-contrastive) chunk content encoded using the original tokenizer.

4.3 Intra-Class Protection via Cascade

To enhance intra-class protection in RAG systems, we introduce a cascaded generation framework combining a fast draft model with a robust reference model. Let $q(z|x_{<t}, z)$ denote the probability of generating chunk z under the original draft model q . The objective of cascade is to construct a new generation distribution $\pi = (1 - r(x_{<t}, z)) \cdot q_t + r(x_{<t}, z) \cdot p_t$, where $r(x_{<t}, z)$ is a rejection rule that can only take the values 0 or 1. If we want to improve inference efficiency and generation quality and reduce the sampling probability of $\pi(z|x_{<t}, z)$, we need to design an optimal rejection rule r . The optimal rejection rule r should minimize the loss of the large model supervision under the cascade distribution π . To control the cost of fallback to p , we minimize the expected loss under a constraint on the rejection budget B . Meanwhile, to reduce the probability of chunk z being generated, we introduce the threshold C , which requires that the probability of draft model generating sensitive content be significantly lower than that of reference model:

$$\begin{aligned} \min_r \quad & \mathbb{E}_{y \sim \mathbb{P}(\cdot|x_{<t}, z)} [(1 - r(x_{<t}, z)) \cdot \ell(y, q_t) \\ & + r(x_{<t}, z) \cdot \ell(y, p_t)] \\ \text{s.t.} \quad & r(x_{<t}, z) \cdot D_{\text{TV}}(p_t, q_t) \leq B, \\ & (1 - r(x_{<t}, z)) \cdot \frac{q_t(z)}{p_t(z)} \leq C. \end{aligned} \quad (2)$$

Analysis of Rejection Rule. To ease optimization, we transform it into an unconstrained problem using Lagrangian

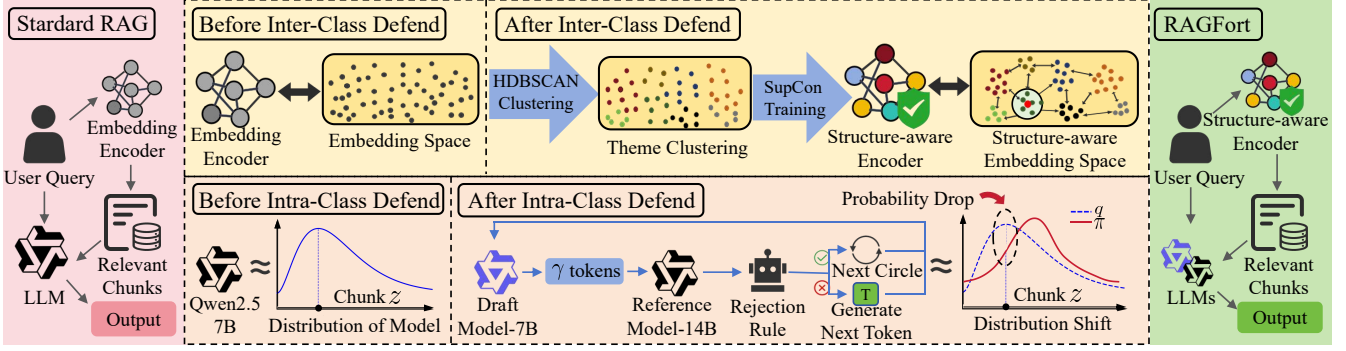


Figure 2: An overview of RAGFort. For inter-class protection, we introduce a structure-aware encoder trained to cluster and separate topics in the embedding space, making it harder for attackers to retrieve content across categories. For intra-class protection, candidate tokens proposed by a draft model are rigorously screened by a verifier model through a tailored rejection rule, which blocks sensitive or risky outputs.

relaxation:

$$\begin{aligned}
& L(r; x_{<t}, z) \\
&= \mathbb{E}_{y \sim \mathbb{P}(\cdot | x_{<t}, z)} [r(x_{<t}, z) \cdot (\ell(y, p_t) + \alpha \cdot D_{TV}(p_t, q_t)) \\
&+ (1 - r(x_{<t}, z)) \cdot (\ell(y, q_t) + \eta \cdot \frac{q_t(z)}{p_t(z)})], \quad (3)
\end{aligned}$$

where α and η are Lagrangian multipliers. The optimal r for minimizing (3) is given by:

Lemma 1 Let $\mathbb{P}_t(y) = \mathbb{P}(y | x_{<t}, z)$ is real distribution. Then the optimal rejection rule is:

$$r^*(x_{<t}, z) = \begin{cases} 1 & \text{if } \mathbb{E}_{y \sim \mathbb{P}_t} [\ell(y, q_t)] > \mathbb{E}_{y \sim \mathbb{P}_t} [\ell(y, p_t)] \\ & + \alpha \cdot D_{TV}(p_t, q_t) - \eta \cdot \frac{q_t(z)}{p_t(z)}, \\ 0 & \text{otherwise.} \end{cases} \quad (4)$$

Intuitively, the cascade system will only choose to generate a verifier model when the expected loss of the verifier model is significantly too high. Since \mathbb{P}_t is unknown, we approximate the expected loss using confidence scores:

$$\begin{aligned}
\hat{r}(x_{<t}, z) &= 1 \quad \text{if and only if} \\
\max_y q_t(y) &< \max_y p_t(y) - \alpha \cdot D_{TV}(p_t, q_t) + \eta \cdot \frac{q_t(z)}{p_t(z)}. \quad (5)
\end{aligned}$$

This rule triggers fallback when p_t is sufficiently more confident than q_t . Next, we analyze the regret (performance gap) between the simple rejection rule \hat{r} and the optimal rejection rule r .

Lemma 2 Let $\mathbb{P}_t(y) = \mathbb{P}(y | x_{<t}, z)$ is real distribution. Then for fixed $x_{<t}$ and z , we have

$$\begin{aligned}
& L(\hat{r}; x_{<t}, z) - \min_r L(r; x_{<t}, z) \\
&\leq \max_{y \in \mathcal{Y}} |\mathbb{P}_t(y) - q_t(y)| + \max_{y \in \mathcal{Y}} |\mathbb{P}_t(y) - p_t(y)|. \quad (6)
\end{aligned}$$

This result shows that the regret is small when both q_t and p_t are close to the ground-truth distribution \mathbb{P}_t .

The algorithm 1 describes a cascading generation process in which the draft model first proposes a block of γ tokens

and calculates the probability q_t . For each token position, the verifier model concurrently computes the probability p_t and the rejection rule r , and then uses r to generate π . Subsequently, the earliest rejection position j^* is sought. If a rejection occurs, a residual distribution is constructed to optimize token generation at j^* ; otherwise, the draft output is retained, and the verifier model is used to generate the starting token for the next block. Importantly, we utilize the Bernoulli distribution during the rejection process. This ensures that the entire generation process starts from q but ultimately generates tokens whose distribution does not deviate from the verifier model p (Yan, Agarwal, and Venkataraman 2024; Zhou et al. 2023).

5 Experiments

5.1 Evaluation Settings

RAG Architecture and Implementation Details. We implement a standard RAG pipeline using the LangChain framework (Chase 2022). Each data sample (either a QA pair or a document) is treated as a distinct chunk in the knowledge base. At inference time, the built-in LangChain retriever, which operates on dense vector similarity, retrieves the top 5 relevant chunks for each query. The retrieved chunks are then concatenated with the question to form the input to the generator, following the prompt format: “Context: [...] Question: [...] Answer:”. For the generator, we evaluate three large language models: Qwen-14B (Jiang et al. 2023), a high-capacity open-source model; DeepSeek-R1-Distill-Llama-8B (abbreviated as *DeepSeek-R1-8B*) (Guo et al. 2025), a lightweight distilled variant of the LLaMA series optimized for chat-based generation; and Gemma-3-27B (Team et al. 2025), a strong 27B-parameter model featuring enhanced reasoning and factual consistency through large-scale multi-domain instruction tuning. All RAG systems share the same retriever, prompting format, and decoding settings to ensure a fair comparison.

Dataset and Evaluation Protocol. We evaluate on three datasets spanning different domains: (1) **Health-**

Datasets	Models	Without protection		RAGFort (Ours)		Re-ranking		Summarization	
		Worm-attack	RAG-Thief	Worm-attack	RAG-Thief	Worm-attack	RAG-Thief	Worm-attack	RAG-Thief
HealthCareMagic	Qwen-14B	17.60±1.58	57.16±2.13	8.84±0.91	27.96±1.32	16.28±1.17	56.44±2.08	15.44±1.20	47.24±1.81
	DeepSeek-R1-8B	13.92±1.29	51.64±1.92	8.56±0.83	26.72±1.18	12.64±0.94	49.28±1.76	12.88±0.91	42.52±1.54
	Gemma-3-27B	18.36±1.33	59.88±1.71	9.76±0.88	29.44±1.11	18.24±0.96	57.36±1.84	16.12±0.97	47.28±1.62
Enron Email	Qwen-14B	31.60±1.73	52.60±1.95	14.30±1.01	20.60±0.89	28.65±1.58	44.35±1.49	25.35±1.43	46.75±1.66
	DeepSeek-R1-8B	28.75±1.45	47.85±1.88	10.15±0.75	17.65±0.69	23.45±1.12	43.75±1.37	24.75±1.16	40.65±1.31
	Gemma-3-27B	32.35±1.41	54.15±1.99	14.40±0.71	24.55±0.76	28.85±1.28	46.50±1.52	27.60±1.13	47.30±1.43
Math QA	Qwen-14B	47.84±2.24	70.31±2.87	29.06±1.02	38.13±1.41	39.49±1.90	74.69±2.99	35.31±1.78	68.44±2.67
	DeepSeek-R1-8B	43.13±2.11	67.81±2.54	28.44±1.39	37.81±1.87	40.00±1.97	58.75±2.01	41.88±1.91	55.63±1.88
	Gemma-3-27B	48.75±2.26	69.06±2.62	29.69±0.98	38.44±1.38	43.13±1.78	63.44±1.95	43.44±1.51	64.69±1.57
Relative Mean CRR		1×		0.51×		0.91×		0.87×	

Table 1: Security assessment of RAGFort in preventing knowledge base extraction attack. We report the Chunk Recovery Rate (CRR, %) of RAG systems under two black-box attack strategies: *Worm-Attack* and *RAG-Thief*. Each column block corresponds to a specific defense method—*Without protection*, *RAGFort (Ours)*, *Re-ranking*, and *Summarization*—while each sub-column reports CRR under a specific attack. Lower CRR ↓ indicates stronger protection. The last row shows the average attack CRR as a multiple of the baseline without protection (×).

Algorithm 1: CASCADES

Input: Draft model q , Reference model p , Rejection rule r , Chunk z , Prefix $x_{<t}$, Block size γ

Output: x_t, \dots, x_{t+j^*}

```

1: for  $j = 0$  to  $\gamma - 1$  do
2:   Compute  $q_{t+j}(\cdot \mid x_{<t+j}, z)$  and  $q_t(z)$ 
3: end for
4: Compute  $p_t(z)$ 
5: for  $j = 1$  to  $\gamma$  do
6:   Compute  $p_{t+j}(\cdot \mid x_{<t+j}, z)$  and  $r_{t+j}(x_{<t+j}, z)$ 
7:    $\pi_{t+j}(\cdot) \leftarrow (1 - r(x_{<t+j}, z)) \cdot q + r(x_{<t+j}, z) \cdot p$ 
8: end for
9:  $j^* = \gamma$ 
10: for  $j = 0$  to  $\gamma - 1$  do
11:   if  $r_{t+j}(x_{<t+j}, z) = 0$  then
12:     continue
13:   else
14:      $a_j \sim \text{Bernoulli} \left( \min \left\{ 1, \frac{p_{t+j}(x_{t+j})}{q_{t+j}(x_{t+j})} \right\} \right)$ 
15:     if  $a_j = 0$  then  $j^* = j$  and break
16:   end if
17: end if
18: end for
19: if  $j^* < \gamma$  then
20:    $p_{\text{cas}}(\cdot) \leftarrow \text{norm}(\max\{0, \pi_{t+j^*}(\cdot) - q_{t+j^*}(\cdot)\})$ 
21: else
22:    $p_{\text{cas}}(\cdot) \leftarrow \pi_{t+\gamma}$ 
23: end if
24: Sample  $x_{t+j^*} \sim p_{\text{ref}}(\cdot)$ 

```

CareMagic (Pampari et al. 2018): medical QA with symptom, diagnosis, and treatment queries; (2) **Enron Email** (Klimt and Yang 2004): corporate email corpus covering varied business contexts; (3) **MathQA** (Amini et al. 2019): math word problems requiring numerical reasoning.

Each RAG system is evaluated using QA pairs constructed from the corresponding dataset. We use GPT-4-turbo (*gpt-4-0125-preview*) (OpenAI 2024) as an automatic evaluator to determine whether each generated answer is semantically consistent and factually aligned with the ground truth. The overall performance is measured by the percent-

age of correct responses, reported as **Accuracy (ACC)**.

Attack Strategies. We evaluate defense robustness under the two most widely recognized attack strategies: (1) **Worm-Attack** (Cohen, Bitton, and Nassi 2024): Utilizes adversarial, self-replicating prompts to trigger cascading extraction and propagation across RAG-based applications, simulating a computer worm in the GenAI ecosystem. Each experiment runs 1,000 attack rounds with 10 collision vector search iterations per round. (2) **RAG-Thief** (Jiang et al. 2024): An agent-based attack that maintains memory of extracted chunks and iteratively generates new adversarial queries through self-improvement, enabling scalable extraction from the knowledge base. Each round issues 10 queries, initialized with 300, 250, and 50 GPT-generated prompts for HealthCareMagic, EnronMail, and MathQA, respectively (final chunk sets: 2,500, 2,000, and 320).

Baseline Defenses. Our proposed method is compared to two recently proposed defense baselines: (1) **Re-ranking Protection** (Zeng et al. 2024a): an inter-class protection that enforces semantic similarity thresholds during retrieval to filter loosely related documents and limit inter-class exposure. (2) **Summarization Protection** (Zeng et al. 2024b): Substitutes retrieved passages with abstract summaries to obscure fine-grained content and reduce intra-class leakage.

Security Metrics. Following prior work (Jiang et al. 2024), we introduce the **Chunk Recovery Rate (CRR)**, which jointly captures both lexical and semantic similarity between generated outputs and original knowledge chunks. A chunk is counted as recovered only if it satisfies both of the following criteria: (1) **ROUGE-L** (Lexical Similarity): The ROUGE-L score between the generated output and the original chunk exceeds a predefined threshold of 0.5. (2) **Semantic Similarity**: The cosine similarity between the embedding vectors of the generated and original text, computed using a pre-trained model such as Sentence-BERT or BERTScore, exceeds a semantic threshold of 0.85.

5.2 Security

To evaluate defense effectiveness, we simulate knowledge extraction attacks and assess how well each method prevents unauthorized replication. The attacker interacts with

Datasets	Models	Without protection		RAGFort (Ours)		RAGFort _{InterOnly}		RAGFort _{IntraOnly}	
		Worm-attack	RAG-Thief	Worm-attack	RAG-Thief	Worm-attack	RAG-Thief	Worm-attack	RAG-Thief
HealthCareMagic	Qwen-14B	17.60±1.58	57.16±2.13	8.84±0.91	27.96±1.32	14.52±0.62	41.16±1.71	13.62±0.83	51.28±2.57
	DeepSeek-R1-8B	13.92±1.29	51.64±1.92	8.36±0.83	26.72±1.18	9.04±0.45	39.56±2.22	11.51±0.68	42.36±2.67
	Gemma-3-27B	18.36±1.33	59.88±1.71	9.76±0.88	29.44±1.11	15.28±0.51	46.32±1.35	14.72±0.73	51.56±2.31
Enron Email	Qwen-14B	31.60±1.73	52.60±1.95	14.30±1.01	20.60±0.89	19.75±0.90	33.25±1.25	26.10±1.75	36.70±1.40
	DeepSeek-R1-8B	28.75±1.45	47.85±1.88	10.15±0.75	17.65±0.69	24.15±0.95	31.40±1.15	22.55±1.65	27.85±2.60
	Gemma-3-27B	32.35±1.41	54.15±1.99	14.40±0.71	24.55±0.76	24.85±1.12	34.85±1.22	27.35±1.30	45.05±1.95
Math QA	Qwen-14B	47.84±2.24	70.31±2.87	29.06±1.02	38.13±1.41	30.94±0.91	54.61±2.19	43.75±1.51	65.63±2.58
	DeepSeek-R1-8B	43.13±2.11	67.81±2.54	28.44±1.39	37.81±1.87	41.25±1.61	52.19±1.68	42.19±2.23	53.78±1.74
	Gemma-3-27B	48.75±2.26	69.06±2.62	29.69±0.98	38.44±1.38	37.81±1.43	54.06±2.02	44.06±1.78	62.19±2.26
Relative Mean CRR		1×		0.51×		0.75×		0.83×	

Table 2: Ablation study of RAGFort defense modules against knowledge base extraction attacks. We report the Chunk Recovery Rate (CRR, %) for the full *RAGFort* and its two path variants under different attack strategies, where a lower CRR indicates stronger protection. *RAGFort*_{InterOnly} includes only the inter-class protection module (contrastive reindexing), while *RAGFort*_{IntraOnly} includes only the intra-class protection module (cascade).

Datasets	Models	Before protection		After protection	
		ACC	FLOPs (T)	ACC	FLOPs (T)
HealthCareMagic	Qwen-14B	68.16±1.45	38.03±0.27	66.57±1.22	19.40±0.35
	DeepSeek-R1-8B	61.36±0.61	20.37±0.23	61.12±0.44	20.83±0.24
	Gemma-3-27B	69.64±0.75	84.83±0.58	68.88±0.36	34.33±0.21
Enron Email	Qwen-14B	99.10±0.11	106.14±0.30	98.85±0.25	56.84±0.40
	DeepSeek-R1-8B	97.50±0.15	56.90±0.15	97.60±0.17	57.43±0.48
	Gemma-3-27B	99.25±0.21	236.47±0.86	99.15±0.38	91.40±0.33
Math QA	Qwen-14B	85.31±1.63	47.09±0.35	82.81±1.51	19.30±0.37
	DeepSeek-R1-8B	78.44±1.49	25.22±0.15	77.19±2.47	25.37±0.36
	Gemma-3-27B	87.50±0.41	105.21±0.22	85.94±1.26	31.95±0.31

Table 3: System performance before and after applying RAGFort protection. We report the Answer Accuracy (ACC, %) and computational cost (FLOPs (T)) of RAG systems with and without RAGFort.

the RAG system to reconstruct a proxy knowledge base, and we measure defense performance by the number of recovered chunks. Lower values indicate stronger protection.

As shown in Table 1, RAGFort achieves the lowest Chunk Recovery Rate (CRR) across both Worm-Attack and RAG-Thief, demonstrating significantly stronger protection than existing defenses. Specifically, RAGFort reduces the Relative Mean CRR to just 0.51× of the original. By comparison, re-ranking and summarization achieve only 0.91× and 0.87×, respectively, indicating substantially weaker protection. These results validate the effectiveness of our dual-path design: contrastive reindexing restricts inter-class exposure, while cascade suppresses intra-class leakage, jointly blocking both attack paths.

5.3 System Performance

To assess the impact on system performance of RAGFort, we compare QA accuracy and FLOPs before and after protection (as shown in Table 3). Results show that RAGFort almost does not affect system performance: for all models and datasets, the ACC drops by less than 2 percentage points (e.g., Qwen-14B on HealthCareMagic: 68.16 → 66.57; DeepSeek-R1-8B: 61.36 → 61.12), and FLOPs remain nearly unchanged or even decrease.

We believe the small accuracy drop mainly arises from the gap between the draft and original models, not the defense itself. Specifically, Qwen-14B, which uses a smaller draft model, sees a slightly greater drop than DeepSeek-R1-8B, which uses the same model for both stages. Thus, we rec-

ommend stronger draft models when accuracy is a priority. Regarding efficiency, in the Qwen-14B group, the decrease in FLOPs results from replacing large-model inference with small-model inference, while in DeepSeek-R1-8B, which uses the same model as the draft model, FLOPs are nearly unchanged. This indicates that even when no smaller draft model is available, RAGFort introduces negligible overhead.

5.4 Ablation Study

To assess the individual contributions of each defense module, we perform an ablation study by isolating the inter-class and intra-class protection components. Similar to the security evaluation, we apply attacks on RAG systems equipped with only one defense module and measure the resulting chunk recovery rate and QA accuracy. This allows us to quantify the security gains provided by each component and validate their complementary roles. As shown in Table 2, neither RAGFort_{InterOnly} nor RAGFort_{IntraOnly} alone provides sufficient defense. For instance, under RAG-Thief, the full RAGFort reduces CRR from 57.16% to 27.96% (HealthCareMagic, Qwen-14B), whereas the single-module variants still expose over 40% of content. Averaged across all settings, the full method lowers the relative mean CRR to 0.51×, significantly outperforming RAGFort_{InterOnly} (0.75×) and RAGFort_{IntraOnly} (0.83×). These results demonstrate that inter-class and intra-class protection are complementary, only their joint deployment can provide robust and comprehensive defense against knowledge extraction attacks.

6 Conclusion

In this paper, we address the problem of protecting proprietary knowledge bases in retrieval-augmented generation systems. Through systematic experiments, we demonstrate that effective protection requires jointly securing both intra-class and inter-class paths. Based on these insights, we propose RAGFort, a dual-path defense combining contrastive reindexing for inter-class isolation and cascade generation for intra-class robustness. Experimental results show that RAGFort significantly reduces knowledge base reconstruction while maintaining high answer quality and efficiency. To the best of our knowledge, RAGFort is the first systematic framework addressing dual-path protection in RAG systems, offering a practical solution for secure RAG deployments.

Acknowledgements

This work was supported by the Key Project of the National Natural Science Foundation of China under Grant no. 62536007, the Zhejiang Province Science Foundation under Grant no. LD24F020002 and the Zhejiang Province's 2025 "Leading Goose + X" Science and Technology Plan under Grant no. 2025C02034.

References

- Amini, A.; Gabriel, S.; Lin, P.; Koncel-Kedziorski, R.; and Hajishirzi, H. 2019. MathQA: Towards interpretable math word problem solving with operation-based formalisms. In *NAACL*.
- Brown, T. B.; Mann, B.; Ryder, N.; et al. 2020. Language models are few-shot learners. *NeurIPS*.
- Campello, R. J.; Moulavi, D.; and Sander, J. 2013. Density-based clustering based on hierarchical density estimates. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, 160–172. Springer.
- Chase, H. 2022. LangChain. <https://github.com/hwchase17/langchain>. Accessed: 2025-06-23.
- Cohen, S.; Bitton, R.; and Nassi, B. 2024. Unleashing worms and extracting data: Escalating the outcome of attacks against rag-based inference in scale and severity using jailbreaking. *arXiv preprint arXiv:2409.08045*.
- Guo, D.; Yang, D.; Zhang, H.; Song, J.; Zhang, R.; Xu, R.; Zhu, Q.; Ma, S.; Wang, P.; Bi, X.; et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Izacard, G.; and Grave, E. 2021. Towards faithful retrieval-augmented generation: A benchmark, evaluation and model. In *NeurIPS*.
- Jiang, C.; Pan, X.; Hong, G.; Bao, C.; and Yang, M. 2024. Rag-thief: Scalable extraction of private data from retrieval-augmented generation applications with agent-based attacks. *arXiv preprint arXiv:2411.14110*.
- Jiang, Y.; Hua, W.; Qin, Y.; and et al. 2023. Qwen: Scaling Up Language Models with Decoupled Parameter and Data Growth.
- Klimt, B.; and Yang, Y. 2004. The Enron corpus: A new dataset for email classification research. In *European Conference on Machine Learning*, 217–226. Springer.
- Lewis, P.; Perez, E.; Piktus, A.; et al. 2020. Retrieval-augmented generation for knowledge-intensive NLP tasks. In *NeurIPS*.
- Narasimhan, H.; Jitkrittum, W.; Rawat, A. S.; Kim, S.; Gupta, N.; Menon, A. K.; and Kumar, S. 2024. Faster Cascades via Speculative Decoding. *arXiv preprint arXiv:2405.19261*.
- Ni, B.; Liu, Z.; Wang, L.; Lei, Y.; Zhao, Y.; Cheng, X.; Zeng, Q.; Dong, L.; Xia, Y.; Kenthapadi, K.; et al. 2025. Towards Trustworthy Retrieval Augmented Generation for Large Language Models: A Survey. *arXiv preprint arXiv:2502.06872*.
- OpenAI. 2023. GPT-4 Technical Report. <https://openai.com/research/gpt-4>.
- OpenAI. 2024. GPT-4 Turbo Technical Report. <https://openai.com/gpt-4-turbo>. Accessed: July 2025.
- Pampari, A.; Raghavan, P.; Liang, J.; and Yu, R. J. M. 2018. emrQA: A large corpus for question answering on electronic medical records. In *EMNLP*.
- Singhal, K.; Azizi, S.; Tu, T.; et al. 2022. Large language models encode clinical knowledge. *arXiv preprint arXiv:2212.13138*. Demonstrates retrieval-augmented prompting on medical QA tasks.
- Team, G.; Kamath, A.; Ferret, J.; Pathak, S.; Vieillard, N.; Merhej, R.; Perrin, S.; Matejovicova, T.; Ramé, A.; Rivière, M.; et al. 2025. Gemma 3 technical report. *arXiv preprint arXiv:2503.19786*.
- Wu, S.; et al. 2023. BloombergGPT: A large language model for finance. *arXiv preprint arXiv:2303.17564*. Includes experiments combining BloombergGPT with retrieval-augmented prompting.
- Yan, M.; Agarwal, S.; and Venkataraman, S. 2024. Decoding speculative decoding. *arXiv preprint arXiv:2402.01528*.
- Zeng, S.; Zhang, J.; He, P.; Liu, Y.; Xing, Y.; Xu, H.; Ren, J.; Chang, Y.; Wang, S.; Yin, D.; et al. 2024a. The Good and The Bad: Exploring Privacy Issues in Retrieval-Augmented Generation (RAG). In *Findings of the Association for Computational Linguistics ACL 2024*, 4505–4524.
- Zeng, S.; Zhang, J.; He, P.; Ren, J.; Zheng, T.; Lu, H.; Xu, H.; Liu, H.; Xing, Y.; and Tang, J. 2024b. Mitigating the privacy issues in retrieval-augmented generation (rag) via pure synthetic data. *arXiv preprint arXiv:2406.14773*.
- Zhou, Y.; Lyu, K.; Rawat, A. S.; Menon, A. K.; Rostamizadeh, A.; Kumar, S.; Kagy, J.-F.; and Agarwal, R. 2023. Distillspec: Improving speculative decoding via knowledge distillation. *arXiv preprint arXiv:2310.08461*.