

# SecureRAG: End-to-End Secure Retrieval-Augmented Generation

Anonymous ACL submission

## Abstract

Retrieval-augmented generation (RAG) enhances large language models (LLMs) with external knowledge from databases but introduces privacy risks when handling sensitive information. Existing defenses fall short: differential privacy degrades accuracy and remains vulnerable to embedding inversion attacks, while fully homomorphic encryption (FHE) ensures security but lacks access control. We present SecureRAG, an end-to-end secure RAG framework that enforces strict access control while mitigating prompt injection data extraction and embedding inversion attacks. It achieves this by decoupling retrieval into secure search and secure document fetching, using FHE for encrypted search and attribute-based encryption (ABE) for fine-grained access control. SecureRAG supports dynamic database updates, adaptive access policies, and integrates seamlessly with FHE-friendly LLMs, adding only 0.05s of overhead. By providing a fully encrypted, privacy-preserving retrieval framework, SecureRAG enables the secure deployment of domain-specific chatbots in sensitive applications.

## 1 Introduction

Retrieval-augmented generation (RAG)(Lewis et al., 2020) enhances large language models (LLMs)(Meng et al., 2024) by retrieving relevant information from external sources to generate more accurate and context-aware responses without re-training. By bridging information retrieval and text generation, RAG enables cost-effective chatbot customization across the healthcare, finance, and law sectors. For example, Figure 1 illustrates RAG in healthcare, where a hospital corpus aids doctors (*users*) in generating precise diagnostic suggestions.

Despite its advantages, RAG is highly vulnerable to privacy and security risks in sensitive sectors (Zeng et al., 2024; Qi et al., 2024). The primary threat is *sensitive information leakage*, which can

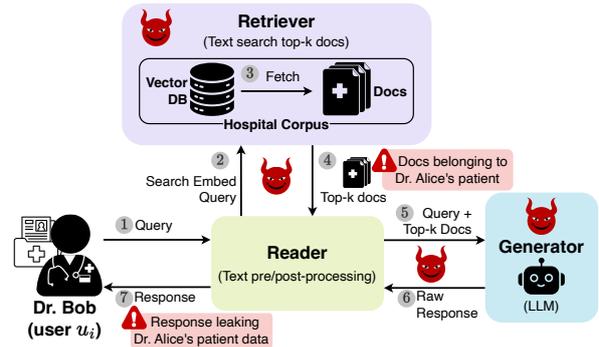


Figure 1: Healthcare-based RAG showing key vulnerabilities: a) retriever exploiting query embeddings and hospital corpus, b) generator misusing queries and retrieved documents, and c) Dr. Bob receiving sensitive data of Dr. Alice’s patient due to lack of access control.

occur through (1) compromised components (e.g., retriever, generator) or (2) adversarial prompt injections that extract restricted data from the database (DB). Beyond direct DB access, exposing cleartext embeddings (queries and documents) risks leaking personally identifiable information (PII). Morris et al. (2023) demonstrated that text embeddings are highly invertible, recovering 89% of PII (e.g., full names) from clinical note embeddings, underscoring the need for equal protection of raw text and embeddings. Additionally, Qi et al. (2024) introduced a prompt injection data extraction attack that targeted RAG’s retrieval DB—rather than the LLM’s training data—successfully extracting 41% of a 77K-word book and 3% of a 1.5M-word corpus using only 100 crafted queries. These findings highlight a critical weakness: RAG lacks effective access control at the retrieval level, leaving it vulnerable to data extraction attacks.

These vulnerabilities are evident in the healthcare scenario shown in Figure 1, where Dr. Bob must uphold patient confidentiality while potentially sharing patients with other doctors. He must also be prevented from using prompt injection attacks to access records of Dr. Alice’s patients, as

they do not collaborate. Deploying separate RAG solutions for each doctor is impractical due to inefficiency, cost, and complexity. Instead, hospitals require a unified RAG system integrated with their database while enforcing strict access controls. The key challenge is ensuring that responses are personalized while *retrieved documents remain both relevant and restricted to authorized patients, preserving security and privacy.*

Most privacy-preserving RAG solutions (Grislain, 2024; Cheng et al., 2024) rely on differential privacy (DP) (Dwork, 2006) as a lightweight defense, adding controllable noise to balance privacy and accuracy. However, DP has critical limitations: (1) it distorts text, reducing retrieval and generation accuracy, (2) its privacy guarantee weakens over repeated queries due to budget exhaustion, and (3) it fails to prevent text embedding inversion (Kim et al., 2022) and prompt injection attacks. These shortcomings make DP unsuitable for an end-to-end, provably secure RAG without accuracy loss.

Fully homomorphic encryption (FHE) (Gentry, 2009) has emerged as a promising approach, offering end-to-end security with strong privacy guarantees and no accuracy loss. Despite its runtime overhead, recent advances demonstrated their effectiveness in preventing information leakage from text embeddings (Kim et al., 2022), enabling secure text classification (Al Badawi et al., 2020; Lee et al., 2022), and supporting secure LLM inference (Zhang et al., 2024; de Castro et al., 2024; Rho et al., 2024). While FHE provides essential building blocks for a secure RAG, existing solutions do not afford access control on retrieved documents.

We propose SecureRAG, an end-to-end secure RAG framework that enforces access control over retrieved documents while preventing prompt injection data extraction and embedding inversion attacks. SecureRAG achieves secure retrieval by splitting the process into two subphases: (1) secure search and (2) secure document fetching. By leveraging FHE’s SIMD property and vertically packing FHE-encrypted embeddings, SecureRAG enables efficient and scalable search over the vector database. To enforce access control, documents are encrypted using an attribute-based encryption (ABE) scheme, ensuring that only authorized users can decrypt retrieved content. SecureRAG supports dynamic management of both the database (e.g., adding/deleting encrypted documents and embeddings) and access rights (e.g., granting/revoking permissions). It seamlessly integrates with FHE-

friendly LLM generators (Rho et al., 2024) without compromising accuracy. We evaluate SecureRAG with the goal of assessing whether encryption can be integrated without compromising accuracy while enabling seamless model updates. Our results show that SecureRAG matches unprotected RAG in accuracy for rank top- $k$  and context precision metrics, the latter assessed by an LLM judge. Performance-wise, SecureRAG retrieves 100 documents from 16,384 under a 2-attribute policy in 0.05s on a single GPU—achieving a 13.6 $\times$  speedup over RemoteRAG (Cheng et al., 2024), which, despite using two GPUs, takes 0.68s to retrieve only 5 documents from 160 with no access control, leaving it vulnerable to prompt injection attacks.

In summary, we introduce SecureRAG, an end-to-end secure RAG framework that enforces access control over retrieved documents while preventing prompt injection data extraction and embedding inversion attacks. SecureRAG achieves this by decoupling retrieval into secure document index search and secure document fetching, utilizing FHE for encrypted embedding search and ABE for fine-grained access control. Our evaluation shows that SecureRAG maintains the accuracy of unprotected RAG while ensuring provable security. With seamless integration into FHE-friendly LLMs and support for dynamic access management and database updates, SecureRAG enables secure, privacy-preserving RAG deployments in sensitive domains such as healthcare.

## 2 Related Work

In RAG, sensitive data leaks through embedding inversion attacks on its components (Morris et al., 2023) or prompt injection queries targeting the extraction of restricted documents (Qi et al., 2024). Existing solutions protect components separately but remain vulnerable to prompt injection attacks.

**Secure text embedding classification.** Many solutions mitigate information leakage by encrypting text embeddings with FHE (Al Badawi et al., 2020; Lee et al., 2022; Kim et al., 2022) for classification tasks. However, they are limited to basic one-to-one similarity comparisons and do not scale to large databases due to FHE’s computational bottlenecks in *search*. SecureRAG addresses this challenge by enabling efficient one-to-many and many-to-many text embedding searches under FHE. It leverages vertical data packing to fully exploit FHE’s SIMD property, significantly improving efficiency and re-

169 ducing encrypted database storage overhead.

170 **Secure inference of LLMs.** LLMs are proprietary,  
171 requiring queries to be sent to external servers for  
172 inference, raising concerns about access and re-  
173 tention of sensitive data. Research on secure in-  
174 ference falls into interactive and non-interactive  
175 models. Interactive models(Akimoto et al., 2023;  
176 Dong et al., 2023; Gupta et al., 2023) rely on mul-  
177 tiparty computation (MPC) but suffer from high  
178 communication overhead, making them imprac-  
179 tical for RAG’s large-scale deployment. Hybrid  
180 approaches(Pang et al., 2024) combine HE and  
181 MPC to mitigate this overhead. FHE-based solu-  
182 tions (de Castro et al., 2024; Zhang et al., 2024;  
183 Rho et al., 2024) offer stronger privacy by trans-  
184 forming LLMs into FHE-friendly architectures, en-  
185 abling non-interactive, end-to-end encrypted infer-  
186 ence. SecureRAG is fully compatible with any  
187 non-interactive FHE-friendly LLM and can seam-  
188 lessly switch between different FHE-friendly LLM  
189 providers without affecting its retriever or access  
190 control policies.

191 **Privacy-preserving RAGs.** While individual RAG  
192 components can be secured, privacy-preserving  
193 RAG solutions aim for end-to-end protection,  
194 primarily against information leakage. Ex-  
195 isting approaches rely on differential privacy  
196 (DP)(Dwork, 2006), but its privacy-accuracy trade-  
197 off degrades LLM performance by altering text  
198 semantics(Mattern et al., 2022). DP also fails to  
199 prevent embedding inversion attacks(Kim et al.,  
200 2022), and its privacy guarantee weakens with re-  
201 peated queries, requiring resets that hinder real-  
202 world deployment(Watson, 2020). Moreover, no  
203 DP-based RAG solutions prevent prompt injection  
204 data extraction attacks. SecureRAG is the first en-  
205 do-end secure RAG framework that enforces access  
206 control over retrieved documents while preventing  
207 information leakage at both the embedding and  
208 generator response levels.

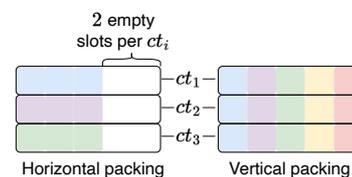
### 209 3 Preliminaries

210 SecureRAG integrates encryption with RAG using  
211 FHE for computations on encrypted text embed-  
212 dings and attribute-based encryption (ABE) for re-  
213 stricting document decryption to authorized users.

#### 214 3.1 Fully Homomorphic Encryption

215 Fully Homomorphic Encryption (FHE) (Gentry,  
216 2009) enables computations on encrypted data  
217 without decryption, providing strong privacy guar-

218 antees. Its IND-CPA security (Cramer et al., 1997)  
219 prevents semi-honest attackers from inferring plain-  
220 texts from ciphertexts. While FHE is computationally  
221 expensive, ongoing optimizations, including  
222 GPU acceleration (Yang et al., 2024), have en-  
223 hanced its practicality. SecureRAG utilizes the  
224 CKKS scheme (Cheon et al., 2017) due to its  
225 support for floating-point operations, enabling en-  
226 crypted computations that closely approximate  
227 cleartext results. CKKS also leverages the single-  
228 instruction multiple-data (SIMD) (Smart and Ver-  
229 cauteren, 2014) property, efficiently packing multi-  
230 ple plaintext values into a single ciphertext. How-  
231 ever, SIMD efficiency depends on the packing  
232 strategy—SecureRAG employs vertical packing,  
233 which is optimized for large-scale vector database  
searches (Figure 2).



234 Figure 2: SecureRAG uses the vertical packing to effi-  
235 ciently store the encrypted text embeddings in DB.

#### 235 3.2 Attribute-Based Encryption

236 Attribute-Based Encryption (ABE) enables fine-  
237 grained access control, allowing decryption only  
238 when user attributes meet a defined policy. In  
239 Ciphertext-Policy ABE (CP-ABE), access policies  
240 are in the ciphertext, giving data owners control,  
241 while in Key-Policy ABE (KP-ABE), policies are  
242 in decryption keys, managed by a central author-  
243 ity. SecureRAG requires KP-ABE, where hospital  
244 authorities control access policies and keys. While  
245 most KP-ABE schemes use pairing-based cryp-  
246 tography, LWE and RLWE-based alternatives (Dai  
247 et al., 2017; Luo et al., 2024) offer stronger security,  
248 and revocability. RLWE-based KP-ABE (Figure 3)  
249 enables *key homomorphism*, allowing homomor-  
250 phic evaluation of public keys over a circuit pol-  
251 icy. Selective security (IND-sCPA) (Goyal et al.,  
252 2006) ensures adversaries without authorized keys  
253 cannot distinguish between encrypted messages.  
254 Like FHE, LWE and RLWE-based KP-ABE rely  
255 on lattice-based hardness, providing post-quantum  
256 security guarantees.

#### 257 4 Threat Model

258 SecureRAG prevents sensitive information leakage  
259 across all RAG stages from malicious components

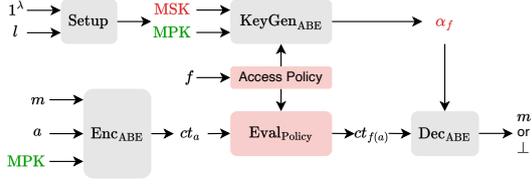


Figure 3: RLWE-based KP-ABE scheme in (Dai et al., 2017). More details are in Appendix A.

(retriever, generator) or users, specifically addressing embedding inversion and prompt injection data extraction attacks. Notably, the latter targets the retrieval database, not the LLM generator’s training data, which SecureRAG does not consider in its threat model. We consider a four-party setting with non-colluding entities: users, a reader, a retriever, and a generator. All parties are semi-honest<sup>1</sup>, except the reader, a trusted third party representing the hospital administrator. The users, representing hospital staff with attribute-based access, query a customized LLM augmented with the hospital corpus for tailored responses. The reader encrypts sensitive data (doctors’ queries, patient records, hospital corpus), sets DB access policies, manages public key infrastructure (PKI), issues user keys, and handles text preprocessing and post-processing. The retriever, a cloud DB provider, stores the encrypted documents and retrieves relevant documents for RAG. The generator, an FHE-friendly proprietary LLM (e.g., OpenAI’s ChatGPT), processes encrypted queries and documents, performs inference under encryption, and returns encrypted responses, protecting its intellectual property. The reader protects hospital data while enabling doctors to use external services. The retriever and generator must not extract meaningful information from queries, corpus, the hospital, or retrieved documents in any form. Users should not intentionally or unintentionally receive responses containing information about other users’ sensitive data or documents they are unauthorized to access.

## 5 SecureRAG

In RAG, the retriever and generator use different text embedding models suitable for text similarity or generation. Thus, SecureRAG extracts the generator’s text embeddings from raw documents on the fly to support any FHE-compatible LLM generator. This enables SecureRAG to integrate with any FHE-friendly LLM generator.

<sup>1</sup>Follow the protocol but try to infer sensitive information.

### 5.1 Key Generation

SecureRAG encrypts text embeddings with FHE, using separate keypairs for the retriever and generator, and encrypts raw documents with KP-ABE to enforce access policies. During setup, the reader, as the system administrator, generates two FHE keypairs:  $(pk_r, sk_r)$  for retriever searches and  $(pk_g, sk_g)$  for generator inference. It shares  $pk_r$  with the retriever and  $pk_g$  with the generator, keeping  $sk_r$  and  $sk_g$  private. The reader also generates a KP-ABE master keypair (MPK, MSK), sharing MPK with users while keeping MSK private.

### 5.2 Encrypted Vector and Document DBs

To build an end-to-end secure RAG while maintaining high efficiency, SecureRAG splits the encrypted vector database (DB) into two parts: an FHE-encrypted chunked vector DB and an ABE-encrypted raw document DB enforcing an access control policy over the retrieved documents.

**FHE-Encrypted chunked vector DB.** The reader encrypts the vector DB following its chunking strategy that optimizes the overall RAG performance. SecureRAG enables the reader to pack  $n$  embeddings  $E = \{e_i\}_{i \in [1, n]}$  of dimension  $d$  using only  $d$  ciphertexts, where  $e_i = (e_{1, i}, \dots, e_{d, i})$  and  $n$  is ciphertext capacity and  $n \gg d$ , which is also the chunk’s size. This is achieved by vertically arranging each chunk’s embeddings and packing them row-wise with one ciphertext  $ct_j = \text{Enc}_{\text{FHE}}^{pk_r}(e_{j, 1}, \dots, e_{j, n})$  per row, resulting in  $d$  ciphertexts per chunk  $ct_E = \{ct_j\}_{j \in [1, d]}$ .

**ABE-Encrypted document DB.** The raw documents  $\{\text{Doc}_t\}_t$  belonging to user  $u_i$  are first ABE-encrypted under user’s attributes  $a_{u_i} \in \{0, 1\}^l$  using the master public key MPK to yield the ciphertext  $ct_{D_t} = \text{Enc}_{\text{ABE}}^{\text{MPK}, a_{u_i}}(\text{Doc}_t)$  over which the reader evaluates the circuit policy  $f$ , resulting in  $ct_{D_t}^f = \text{Enc}_{\text{ABE}}^{\text{MPK}, f(a_{u_i})}(\text{Doc}_t)$ , ciphertexts decryptable with the policy secret key  $\alpha_f$  if and only if  $a_{u_i}$  satisfies the policy  $f$ . To enable a secure direct fetch of those encrypted documents and avoid storing their respective IDs, we use a keyed pseudo-random function (Bellare and Cash, 2010)  $D_t = \text{PRF}_K(d_t)$  that given a secret key  $K$  and an index  $d_t$  it returns the same pseudorandom  $D_t$  completely different from  $d_t$ . The reader generates and sends the ABE-encrypted documents  $\{ct_{D_t}^f\}_t$  along with their pseudorandom identifiers  $\{D_t\}_{t \in [1, k]}$  to the retriever for storage. SecureRAG benefits from this by making fetching secure doc-

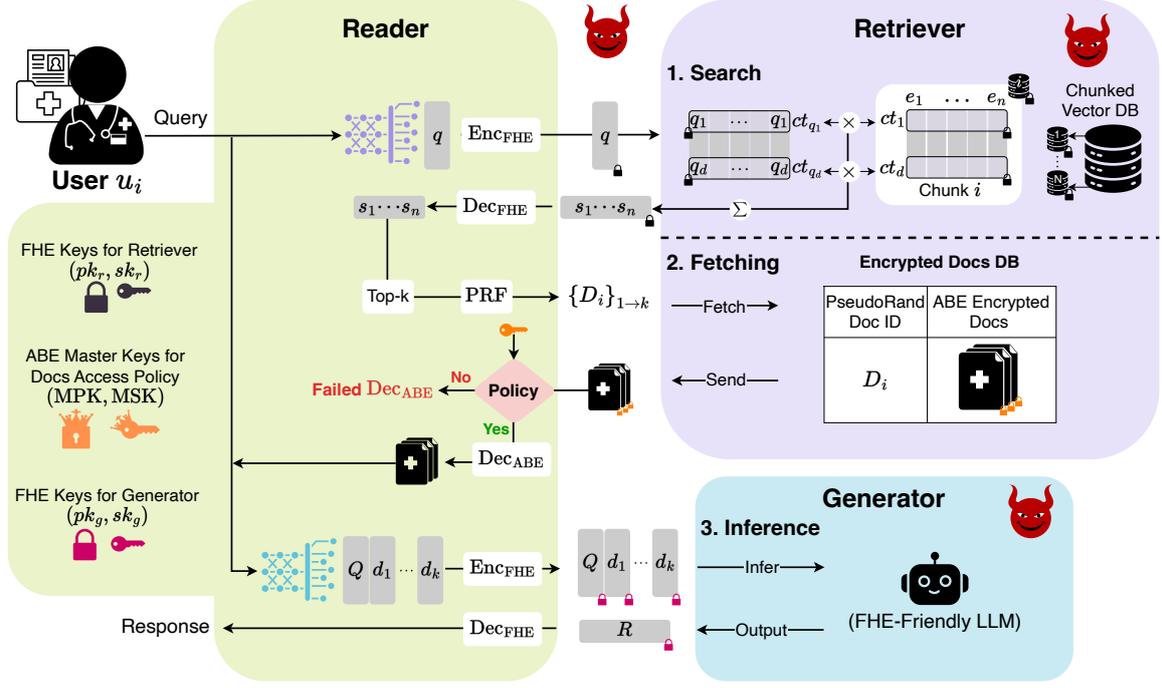


Figure 4: Overview of the SecureRAG framework integrating an FHE-friendly LLM generator with ABE-encrypted documents, ensuring fine-grained access control. 1) The user’s query embedding is FHE-encrypted and used to search for the top-k document indexes in the encrypted chunked vector DB. 2) The ABE-encrypted documents are fetched via pseudo-random IDs and only decryptable if the user  $u_i$  satisfies the access policy. Finally, 3) the embeddings of the query and top-k documents are FHE-encrypted and sent to the LLM generator for secure inference, generating an encrypted response, which is post-processed by the reader before delivering it to the user.

uments efficient without the risk of leaking their actual identifiers, which saves storage space.

### 5.3 Protocol Description

SecureRAG, depicted in Figure 4, consists of three steps: 1) FHE-encrypted search for the top-k document indexes, 2) secure fetching of the ABE-encrypted documents via pseudorandom identifiers, and 3) FHE-encrypted LLM inference. To enhance efficiency, the retrieval part is split into 1) and 2), with SecureRAG filtering documents by user attributes and access policy before step 3).

**Searching top-k document indexes.** We consider the query embedding  $q$  and document embeddings  $E = \{e_i\}_{i \in [1, n]}$  as normalized  $d$ -dim vectors. Normalized vectors enable efficient encrypted search, as the inner product (IP) is cheaper to compute under encryption than cosine similarity while preserving identical scores. The equation below shows that the inner product avoids computing the embedding norms, which would require expensive computation under FHE.

$$\text{Cosine}(\tilde{q}, \tilde{e}_j) = \frac{\langle \tilde{q}, \tilde{e}_j \rangle}{\|\tilde{q}\| \cdot \|\tilde{e}_j\|} = \langle q, e_j \rangle$$

A naive encrypted search over  $n$  embeddings would

compare the query against one embedding at a time, resulting in  $n$  IPs costing  $n$  homomorphic multiplications. By leveraging the SIMD property and the vertical packing of the document vector DB, SecureRAG computes those  $n$  IPs at once, reducing the computation cost to only  $d$  homomorphic multiplications, where  $d \ll n$ , while the search remains exhaustive. Hence, the reader extracts the user query embedding  $q = (q_1, \dots, q_j, \dots, q_d)$  that is compatible with the retriever’s embedding model. Then, it encrypts each component  $q_j$  as an  $n$ -dim vector of its replica, yielding  $ct_{q_j} = \text{Enc}_{\text{FHE}}^{pk_r}(q_j, \dots, q_j)$ . Next, it sends  $ct_q = \{ct_{q_j}\}_{j \in [1, d]}$  to the retriever for computing IP w.r.t. each chunk as follow:

$$\langle ct_q, ct_E \rangle = \sum_{j \in [1, d]} q_j \times ct_j \quad 388$$

The retriever returns  $ct_S = \text{Enc}_{\text{FHE}}^{pk_r}(s_1, \dots, s_n)$  a ciphertext containing the IP scores  $s_j = \langle q, e_j \rangle$ . The reader decrypts the scores  $\text{Dec}_{\text{FHE}}^{sk_r}(ct_S) = (s_1, \dots, s_n)$ , sorts them, and selects the top-k.

**Fetching relevant documents.** From the indexes  $\{d_t\}_{t \in [1, k]}$  of the top-k scores, the reader recovers the documents’ pseudorandom identi-

fiers  $D_t$ . Next, the reader sends  $\{D_t\}_{t \in [1,k]}$  to the retriever who sends back the ABE-encrypted documents  $\{ct_{D_t}^f\}_{t \in [1,k]}$  where  $ct_{D_t}^f = \text{Enc}_{\text{ABE}}^{\text{MPK},f(a_{u_i})}(\text{Doc}_t)$  with user’s  $u_i$  attributes  $a_{u_i}$ , an  $l$ -dim binary vector with  $l$  being the maximum number of attributes a user can have. Then, the reader decrypts documents using the secret key  $\alpha_f$  for access policy  $f$ . Decryption succeeds only if user  $u_i$ ’s attributes  $a_{u_i}$  satisfy  $f$ ; otherwise, it fails.

**Encrypted inference.** The reader combines the user’s query with the successfully ABE-decrypted documents and extracts their embeddings with an embedding model compatible with the generator. It then sends their FHE-encryption using  $pk_g$  to the FHE-friendly LLM generator, performs the secure inference under encryption, and returns its encrypted response back to the reader. Subsequently, the reader decrypts it using  $sk_g$ , post-processes it, and displays it to the user.

## 5.4 Complexity and Security Analyses

Table 1 presents SecureRAG’s computational complexity and storage requirements for its secure search part, which is the dominant part of the retrieval. Note that FHE schemes can adjust their parameters to expand ciphertext capacity as needed, which would help in improving efficiency. Our security analysis is in Appendix B.

Table 1: SecureRAG’s retriever search storage and computational complexity as  $\mathcal{O}(N \cdot (\#\text{Add}_{\text{HE}} + \#\text{Mult}_{\text{HE}}))$  where  $N$  is the number of chunks.

Dimension		$d$	256	768
Complexity	Add <sub>HE</sub>	$d - 1$	255	767
	Mult <sub>HE</sub> *	$d$	256	768
Storage	Enc Query	$d$	256	768
	Enc DB	$d \cdot N$	$256 \cdot N$	$768 \cdot N$

\* Those are homomorphic multiplications of depth 1.

## 5.5 Dynamic Databases and Access Rights

SecureRAG efficiently and dynamically handles the addition and deletion of documents and their embeddings with on-the-fly user management, including dynamic addition, revocation, and real-time policy updates. **Dynamic Databases.** For the **addition** of a document  $\text{Doc}_t$ , the reader updates the vector DB with the document embedding  $\bar{e}_t = (\bar{e}_{1,t}, \dots, \bar{e}_{d,t})$  by selecting a chunk with an available empty slot  $d_t$ , and sends to the retriever the set of ciphertexts encrypting  $\bar{e}_t$ , that is,  $ct_{\bar{e}_t} = \{ct_{\bar{e}_{j,t}}\}_{j \in [1,d]}$  where  $ct_{\bar{e}_{j,t}} = \text{Enc}_{\text{FHE}}^{pk_r}(\dots, 0, \bar{e}_{j,t}, 0, \dots)$ . Then, the retriever

then updates the encrypted vector DB  $ct_E = \{ct_j\}_{j \in [1,d]}$  using one homomorphic addition per ciphertext  $ct_j := ct_j + ct_{\bar{e}_{j,t}} \forall j \in [1,d]$ . Also, the reader generates the document’s pseudorandom identifier  $D_t = \text{PRF}_K(d_t)$  and ABE-encryption of the document to which it applies the access policy  $f$  resulting in  $ct_{D_t}^f = \text{Enc}_{\text{ABE}}^{\text{MPK},f(a_{u_i})}(\text{Doc}_t)$ . It sends to the retriever  $ct_{D_t}^f$  for storing it under  $D_t$ . Note that our addition can add batch of embeddings at once  $\bar{e}_{t_1}, \dots, \bar{e}_{t_m}$  at once with  $ct_{\bar{e}_{j,t}} = \text{Enc}_{\text{FHE}}^{pk_r}(\dots, 0, \bar{e}_{j,t_1}, \dots, \bar{e}_{j,t_m}, 0, \dots)$ . SecureRAG supports batch **deletion** of embeddings. The reader creates a deletion vector  $v \in \{0, -1\}^n$ , marking deletions with  $-1$ , and encrypts it as  $ct_v = \text{Enc}_{\text{FHE}}^{pk_r}(v)$ . The retriever updates the encrypted DB  $ct_E = \{ct_j\}_{j \in [1,d]}$  using one homomorphic addition and multiplication per ciphertext:  $ct_j := ct_j + ct_v \times ct_j \quad \forall j \in [1,d]$ . The retriever processes encrypted updates blindly, without knowing the modified documents. **Dynamic access control.** The KP-ABE-SW scheme (Luo et al., 2024) supports switchable attributes, enabling dynamic user management and policy updates. SecureRAG leverages these capabilities to manage access rights dynamically, with all modifications handled by the reader, acting as the system administrator.

## 6 Experiments

Our goal is not to enhance state-of-the-art RAG accuracy but to assess whether an encryption layer can be integrated without compromising accuracy or efficiency in sensitive applications. The experiments in Section 6.1 were conducted using Python 3.12 on a NVIDIA RTX A6000 GPU-equipped server. For embedding search, we implemented CUDA C++ experiments using PhantomFHE (Yang et al., 2024) with the CKKS scheme, tested on HPCC with a single NVIDIA A100 core and 16GB of memory. For document decryption, we used PALISADE-abe (PALISADE, 2020), implementing the lattice-based ABE scheme in C++, running on macOS Sequoia on an Apple M3 Pro (12 cores, 36GB RAM). We will publicly release our code<sup>2</sup>.

### 6.1 Performance Evaluation

We evaluate SecureRAG using the standard RAG assessment pipeline (Es et al., 2023; Roucher, 2024), where an LLM judge systematically assesses retriever and generator performance. The judge receives instructions, the query, retrieved

<sup>2</sup>Upon acceptance at <https://github.com/anonymous>

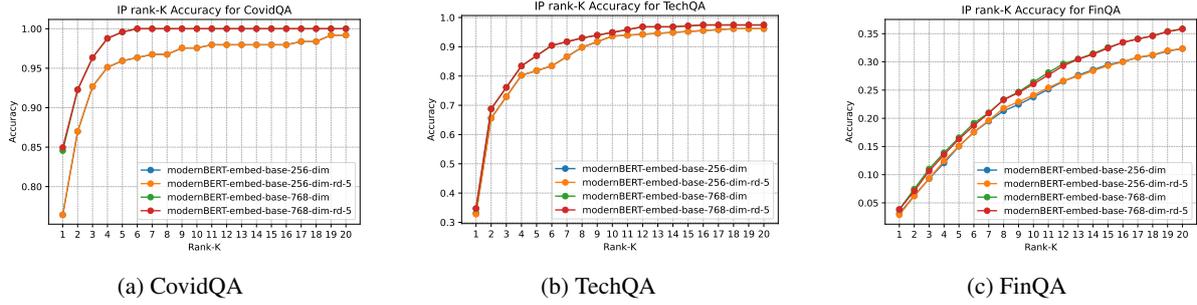


Figure 5: SecureRAG’s retriever performance on CovidQA, TechQA, and FinQA (rounded to 5, orange&red curves) meets unprotected RAG accuracy for rank top- $k$  with  $k \in [1, 20]$  for embeddings of dimensions  $d \in \{256, 768\}$ .

documents (retriever output), and the generated response (LLM output) for structured analysis. To evaluate the impact of encryption on RAG, we apply rounding to 5, reflecting CKKS’s precision limits. Performance in cleartext (no rounding) is compared to the encrypted setting (with rounding) to measure potential losses. **Models.** We use ModernBERT Embed (Nussbaum et al., 2024) for retrieval (supporting 256 and 768 dimensions) and Llama-2-7B (Meta, 2024) as the FHE-friendly generator, optimized for GPUs (Community, 2024; Rho et al., 2024). SecureRAG is tested with ModernBERT Embed as the retriever and Llama-2-7B as the generator, using Llama-3.1-8B as the LLM judge, without fine-tuning any models. **Datasets.** To assess the adaptation of RAG to domain-specific contexts, we consider the following datasets provided in (Friel et al., 2024): the PubMedQA and CovidQA-RAG datasets for the biomedical domain, the TechQA dataset for customer support, and the FinQA and TAT-QA datasets for the financial domain. **Metrics.** We measure the retriever’s effectiveness using the *rank top-k* and the *context precision* metric as defined in (Es et al., 2023), providing the LLM judge with a specific prompt (see Appendix D) instructing it to return a verdict based on the relevance of the retrieved documents w.r.t. the question and the LLM generator’s response. High scores indicate high performance. **Assessment.** For each dataset, the documents column contains documents corresponding to each question. We combined all documents to build the vector DB, extracted their normalized embeddings for  $d \in \{256, 768\}$ , and rounded them to 5. Figure 5 shows that rounding has no impact on retrieval accuracy, as ranking curves with and without rounding overlap almost perfectly, regardless of the embedding dimensionality. Table 2 reports LLM judged context precision for varying retrieved documents ( $k \in \{1, 5, 10\}$ ).

Table 2: Performance of SecureRAG for context precision using a retriever with  $d$ -dim embeddings. Gray (resp. white) cells are with (resp. without) rounding.

Domain	$d$ -dim	256			768		
		K Docs	1	5	10	1	5
Healthcare	PubMedQA	0.889 0.851	0.999 0.999	0.999 0.999	0.839 0.847	0.999 0.999	0.999 0.999
	CovidQA	0.974 0.971	0.999 0.999	0.999 0.999	0.961 0.994	0.999 0.999	0.999 0.999
Finance	TAT-QA	0.929 0.938	0.999 0.999	0.929 0.998	0.911 0.916	0.999 0.999	0.999 0.999
	FinQA	0.955 0.973	0.999 0.999	0.999 0.999	0.970 0.961	0.999 0.999	0.999 0.999
Tech	TechQA	0.999 0.999	0.999 0.999	0.999 0.999	0.999 0.999	0.999 0.999	0.999 0.999

The LLM judge produces nearly identical scores for  $k \geq 5$ , with minor variations at  $k = 1$ , which can be due to its probabilistic nature. SecureRAG maintains RAG accuracy with minimal loss when retrieving one document while preventing prompt injection and data extraction attacks. See Appendix C for PubMedQA and TAT-QA results.

Table 3: SecureRAG’s GPU mean runtime for searching the top-100 embeddings chunk-wise, with  $N$  embeddings per chunk at a 128-bit security level.

Dimension	256	512	768	1024	#Chunks
$N$	16384	16384	16384	16384	–
100 : $N$	18ms	27ms	37ms	47ms	1
100 : $10^6$	1.11s	1.67s	2.29s	2.91s	61
100 : $10^9$	18.31min	27.46min	37.63min	47.81min	61035
$N$	32768	32768	32768	32768	–
100 : $N$	30ms	43ms	58ms	71ms	1
100 : $10^6$	0.93s	1.33s	1.79s	2.2s	30
100 : $10^9$	15.25min	21.87min	29.5min	36.11min	30517

\* Runtime measured over 500 iterations for 100 :  $N$  and extrapolated for 100 :  $10^6$  and 100 :  $10^9$ .

## 6.2 Runtime Evaluation

We measure SecureRAG’s latency for its 3 steps: (1) FHE-encrypted search on GPU, (2) ABE-based secure fetching on CPU, and (3) FHE-encrypted LLM inference, referencing reported GPU benchmarks for state-of-the-art FHE-friendly LLMs.

Table 4: Runtime comparison between SecureRAG and state-of-the-art solutions showing a seamless integration of SecureRAG with the state-of-the-art FHE-friendly LLMs, incurring negligible overhead while effectively preventing prompt injection data extraction attacks.

Solution	Approach	Docs	DB	Dim	Tokens	GPUs	Retriever (s)	Generator (s)	Total (s)	PIDE* Attack	Access Control
(Kim et al., 2022)	FHE	1	1000	768	-	-	0.6	-	-	Vulnerable	✗
RemoteRAG (Cheng et al., 2024)	DP&PHE	5	160	768	-	2	0.68	-	-	Vulnerable	✗
BOLT (Pang et al., 2024)	HE&MPC	-	-	768	128	4	-	185	-	Vulnerable	✗
NEXUS (Zhang et al., 2024)	FHE	-	-	768	128	4	-	37.3	-	Vulnerable	✗
HEaaN (Rho et al., 2024)	FHE	-	-	768	128	1	-	26.5	-	Vulnerable	✗
SecureRAG+NEXUS†	FHE&ABE	100	16384	768	128	4	0.05	37.30	37.35	Prevented	✓
SecureRAG+HEaaN†		100	16384	768	128	1	0.05	26.50	26.55	Prevented	✓

\* Prompt injection data extraction (PIDE) attack (Qi et al., 2024). † Retrieval of 100 documents from a 16384 vector DB using 2 attributes 51.81ms.

**Search Performance.** Table 3 reports SecureRAG’s mean runtime for top-100 document searches using embeddings of dimensions  $d \in \{256, 512, 768, 1024\}$ , evaluated at two chunk capacities (16384 and 32768 embeddings per chunk). For single-chunk searches, runtime ranges from 18ms to 47ms at 16384 capacity and 30ms to 71ms at 32768, showing efficiency gains with larger chunks. For large-scale DBs, SecureRAG scales effectively: retrieving from 1M embeddings requires 30 chunks, adding  $< 2.2$ s in runtime, while retrieval from 1B embeddings spans  $\sim 30$ K chunks in 36.11 minutes. SecureRAG achieves a  $13.6\times$  speedup over RemoteRAG (Cheng et al., 2024), which takes 0.68s on two GPUs to retrieve just 5 documents from 160, lacking access control and remaining vulnerable to prompt injection data extraction attacks.

**Access Control Overhead.** Table 5 measures ABE-decryption runtime for 1 to 100 encrypted documents ( $\sim 350$ – $700$  words) with 2 to 10 attributes. Decryption time scales with attributes but remains minimal compared to search, peaking at 92.43ms, demonstrating efficient enforcement of access control with low overhead.

**End-to-End Runtime.** Table 4 compares SecureRAG with existing approaches that protect only specific RAG components. SecureRAG adds just 0.05s overhead to FHE-friendly non-interactive LLMs, leading to a total runtime of 26.55s to 37.35s. Thus, SecureRAG effectively mitigates prompt injection data extraction attacks while seamlessly integrating with state-of-the-art FHE-friendly LLMs, ensuring strong privacy protection at minimal computational cost.

## 7 Conclusion

Privacy and security are critical for responsible RAG deployment, especially in healthcare, where

Table 5: SecureRAG’s CPU mean runtime for ABE-decrypting  $K$  documents assuming one document per ciphertext using a 128-bit security level.

# Attributes	2	4	6	8	10
16384 bits documents $\sim 350$ words					
1 doc	0.47ms	1.04ms	1.23ms	2.03ms	2.25ms
10 docs	1.20ms	2.14ms	3.15ms	4.09ms	5.08ms
100 docs	13.81ms	20.53ms	28.04ms	35.29ms	41.44ms
32768 bits documents $\sim 700$ words					
1 doc	1.45ms	3.03ms	4.08ms	5.12ms	6.08ms
10 docs	3.13ms	5.30ms	7.37ms	9.31ms	11.20ms
100 docs	27.69ms	45.49ms	61.55ms	76.62ms	92.43ms

\* Runtime measured over 500 iterations.

unauthorized access can lead to severe violations. This paper presents SecureRAG, an end-to-end secure RAG framework that integrates FHE and ABE to prevent information leakage, enforce access control, and defend against prompt injection and embedding inversion attacks. SecureRAG splits retrieval into secure search and secure document fetching, ensuring only authorized users access relevant documents without compromising accuracy. Our evaluation shows SecureRAG matches unprotected RAG in rank top- $k$  and context precision metrics. With a single GPU, it retrieves 100 documents from 16K under a 2-attribute policy in 51.81ms, achieving a  $13\times$  speedup over existing solutions, which retrieve only 5 documents from 160 with no access control. SecureRAG supports dynamic database updates and adaptive access control while seamlessly integrating with FHE-friendly LLMs, adding only 0.05s of overhead. By effectively preventing prompt injection data extraction attacks, SecureRAG provides a scalable, practical solution for privacy-preserving RAG deployments. It addresses core security challenges, laying the foundation for future research on secure and privacy-aware chatbots.

## 8 Limitations

SecureRAG has the following limitations. Its overall runtime is heavily influenced by the efficiency of the FHE-friendly LLM generator it integrates with. Additionally, it operates in three rounds, two of which involve the retriever; reducing these interactions could improve efficiency. Another limitation is that the number of supported attributes is fixed during setup, requiring careful estimation of expected attributes. Increasing this number would slightly impact the retriever’s runtime. Another limitation of SecureRAG is the heavy key management burden on the reader’s side. As a trusted third party system administrator, the reader is responsible for handling cryptographic keys for ABE and FHE operations, including key distribution, updates, and revocations. This overhead can increase storage complexity and require efficient key management strategies to maintain scalability. Also, a trusted third party introduces a single point of failure, as if it is compromised, could undermine the entire system; however, in practice, many real-world deployments rely on a centralized authority for efficiency and trust management. This risk can be mitigated through decentralization schemes such as multi-authority ABE (MA-ABE) or threshold cryptography to ensure no single entity has absolute control.

## Acknowledgments

## References

Yoshimasa Akimoto, Kazuto Fukuchi, Youhei Akimoto, and Jun Sakuma. 2023. Privformer: Privacy-preserving transformer with mpc. In *2023 IEEE 8th European Symposium on Security and Privacy (EuroS&P)*, pages 392–410. IEEE. <https://ieeexplore.ieee.org/document/10190506>.

Ahmad Al Badawi, Louie Hoang, Chan Fook Mun, Kim Laine, and Khin Mi Mi Aung. 2020. Privft: Private and fast text classification with homomorphic encryption. *IEEE Access*. <https://ieeexplore.ieee.org/abstract/document/9296754>.

Mihir Bellare and David Cash. 2010. Pseudorandom functions and permutations provably secure against related-key attacks. In *Annual Cryptology Conference*. Springer.

Yihang Cheng, Lan Zhang, Junyang Wang, Mu Yuan, and Yunhao Yao. 2024. Remoterag: A privacy-preserving llm cloud rag service. *arXiv preprint arXiv:2412.12775*.

Jung Hee Cheon, Andrey Kim, Miran Kim, and Yongsoo Song. 2017. Homomorphic encryption

for arithmetic of approximate numbers. In *International Conference on the Theory and Applications of Cryptology and Information Security (ASIACRYPT)*. Springer. [https://link.springer.com/chapter/10.1007/978-3-319-70694-8\\_15](https://link.springer.com/chapter/10.1007/978-3-319-70694-8_15).

CKKS Community. 2024. <https://iheaan.com/?target=heaan-llm>.

Ronald Cramer, Rosario Gennaro, and Berry Schoenmakers. 1997. A secure and optimally efficient multi-authority election scheme. *European transactions on Telecommunications*.

Wei Dai, Yarkın Doröz, Yuriy Polyakov, Kurt Rohloff, Hadi Sajjadpour, Erkey Savaş, and Berk Sunar. 2017. Implementation and evaluation of a lattice-based key-policy abe scheme. *IEEE Transactions on Information Forensics and Security*, 13(5):1169–1184.

Leo de Castro, Antigoni Polychroniadou, and Daniel Escudero. 2024. Privacy-preserving large language model inference via gpu-accelerated fully homomorphic encryption. In *Neurips Safe Generative AI Workshop 2024*.

Ye Dong, Wen-jie Lu, Yancheng Zheng, Haoqi Wu, Derun Zhao, Jin Tan, Zhicong Huang, Cheng Hong, Tao Wei, and Wenguang Chen. 2023. Puma: Secure inference of llama-7b in five minutes. *arXiv preprint arXiv:2307.12533*. <https://arxiv.org/pdf/2307.12533>.

Cynthia Dwork. 2006. Differential privacy. In *International colloquium on automata, languages, and programming*. Springer.

Shahul Es, Jithin James, Luis Espinosa-Anke, and Steven Schockaert. 2023. Ragas: Automated evaluation of retrieval augmented generation. *arXiv preprint arXiv:2309.15217*.

Robert Friel, Masha Belyi, and Atindriyo Sanyal. 2024. Ragbench: Explainable benchmark for retrieval-augmented generation systems. *arXiv preprint arXiv:2407.11005*.

Craig Gentry. 2009. Fully homomorphic encryption using ideal lattices. In *Proceedings of the forty-first annual ACM symposium on Theory of computing*.

Vipul Goyal, Omkant Pandey, Amit Sahai, and Brent Waters. 2006. Attribute-based encryption for fine-grained access control of encrypted data. In *Proceedings of the 13th ACM conference on Computer and communications security*.

Nicolas Grislain. 2024. Rag with differential privacy. *arXiv preprint arXiv:2412.19291*.

Kanav Gupta, Neha Jawalkar, Ananta Mukherjee, Nishanth Chandran, Divya Gupta, Ashish Panwar, and Rahul Sharma. 2023. Sigma: Secure GPT inference with function secret sharing. *Cryptology ePrint Archive*. <https://eprint.iacr.org/2023/1269.pdf>.

704	Donggyu Kim, Garam Lee, and Sungwoo Oh. 2022.	Zhenting Qi, Hanlin Zhang, Eric Xing, Sham Kakade,	759
705	Toward privacy-preserving text embedding similarity	and Himabindu Lakkaraju. 2024. Follow my instruc-	760
706	with homomorphic encryption. In <i>Proceedings of</i>	tion and spill the beans: Scalable data extraction	761
707	<i>the Fourth Workshop on Financial Technology and</i>	from retrieval-augmented generation systems. <i>arXiv</i>	762
708	<i>Natural Language Processing (FinNLP)</i> . <a href="https://aclanthology.org/2022.finnlp-1.4.pdf">https://</a>	<i>preprint arXiv:2402.17840</i> .	763
709	<a href="https://aclanthology.org/2022.finnlp-1.4.pdf">aclanthology.org/2022.finnlp-1.4.pdf</a> .		
710	Garam Lee, Minsoo Kim, Jai Hyun Park, Seung-	Donghwan Rho, Taeseong Kim, Minje Park, Jung Woo	764
711	won Hwang, and Jung Hee Cheon. 2022. Privacy-	Kim, Hyunsik Chae, Jung Hee Cheon, and Ernest K	765
712	Preserving Text Classification on BERT Embeddings	Ryu. 2024. Encryption-friendly llm architecture.	766
713	with Homomorphic Encryption. In <i>Proceedings of</i>	<i>arXiv preprint arXiv:2410.02486</i> . <a href="https://arxiv.org/pdf/2410.02486">https://arxiv.</a>	767
714	<i>the 2022 Conference of the North American Chap-</i>	<a href="https://arxiv.org/pdf/2410.02486">org/pdf/2410.02486</a> .	768
715	<i>ter of the Association for Computational Linguis-</i>	Aymeric Roucher. 2024. RAG Evaluation -	769
716	<i>tics: Human Language Technologies</i> . <a href="https://aclanthology.org/2022.naacl-main.231.pdf">https://</a>	Hugging Face Open-Source AI Cookbook.	770
717	<a href="https://aclanthology.org/2022.naacl-main.231.pdf">aclanthology.org/2022.naacl-main.231.pdf</a> .	<a href="https://huggingface.co/learn/cookbook/en/rag_evaluation">https://huggingface.co/learn/cookbook/en/</a>	771
718	Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio	<a href="https://huggingface.co/learn/cookbook/en/rag_evaluation">rag_evaluation</a> .	772
719	Petroni, Vladimir Karpukhin, Naman Goyal, Hein-	Nigel P Smart and Frederik Vercauteren. 2014. Fully	773
720	rich Küttler, Mike Lewis, Wen-tau Yih, Tim Rock-	homomorphic simd operations. <i>Designs, codes and</i>	774
721	täschel, et al. 2020. Retrieval-augmented generation	<i>cryptography</i> .	775
722	for knowledge-intensive nlp tasks. <i>Advances in Neu-</i>	Lauren Watson. 2020. An introduction to differen-	776
723	<i>ral Information Processing Systems</i> .	tial privacy. <a href="https://laurenwatson.github.io/blogposts/2020-10-25-dp/">https://laurenwatson.github.io/</a>	777
724	Fucaai Luo, Haiyan Wang, Xingfu Yan, and Jiahui Wu.	<a href="https://laurenwatson.github.io/blogposts/2020-10-25-dp/">blogposts/2020-10-25-dp/</a> .	778
725	2024. Key-policy attribute-based encryption with	Hao Yang, Shiyu Shen, Wangchen Dai, Lu Zhou, Zhe	779
726	switchable attributes for fine-grained access control	Liu, and Yunlei Zhao. 2024. Phantom: A cuda-	780
727	of encrypted data. <i>IEEE Transactions on Information</i>	accelerated word-wise homomorphic encryption li-	781
728	<i>Forensics and Security</i> .	brary. <i>IEEE Transactions on Dependable and Se-</i>	782
729	Justus Mattern, Benjamin Weggenmann, and Florian	<i>ecure Computing</i> . <a href="https://ieeexplore.ieee.org/abstract/document/10428046">https://ieeexplore.ieee.org/</a>	783
730	Kerschbaum. 2022. The limits of word level differ-	<a href="https://ieeexplore.ieee.org/abstract/document/10428046">abstract/document/10428046</a> .	784
731	ential privacy. In <i>Findings of the Association for</i>	Shenglai Zeng, Jiankun Zhang, Pengfei He, Yue Xing,	785
732	<i>Computational Linguistics: NAACL 2022</i> .	Yiding Liu, Han Xu, Jie Ren, Shuaiqiang Wang,	786
733	Xiangbin Meng, Xiangyu Yan, Kuo Zhang, Da Liu, Xi-	Dawei Yin, Yi Chang, et al. 2024. The good and the	787
734	aojuan Cui, Yaodong Yang, Muhan Zhang, Chunxia	bad: Exploring privacy issues in retrieval-augmented	788
735	Cao, Jingjia Wang, Xuliang Wang, et al. 2024. The	generation (rag). <i>arXiv preprint arXiv:2402.16893</i> .	789
736	application of large language models in medicine: A	Jiawen Zhang, Xinpeng Yang, Lipeng He, Kejia Chen,	790
737	scoping review. <i>Iscience</i> , 27(5).	Wen-jie Lu, Yinghao Wang, Xiaoyang Hou, Jian Liu,	791
738	Meta. 2024. Meta llama 2-7b. <a href="https://huggingface.co/meta-llama/Llama-2-7b">https://huggingface.</a>	Kui Ren, and Xiaohu Yang. 2024. Secure trans-	792
739	<a href="https://huggingface.co/meta-llama/Llama-2-7b">co/meta-llama/Llama-2-7b</a> .	former inference made non-interactive. <i>Cryptology</i>	793
740	John Xavier Morris, Volodymyr Kuleshov, Vitaly	<i>ePrint Archive</i> . <a href="https://eprint.iacr.org/2024/136.pdf">https://eprint.iacr.org/2024/</a>	794
741	Shmatikov, and Alexander M Rush. 2023. Text em-	<a href="https://eprint.iacr.org/2024/136.pdf">136.pdf</a> .	795
742	beddings reveal (almost) as much as text. In <i>The</i>	<b>A RLWE-based KP-ABE Scheme</b>	796
743	<i>2023 Conference on Empirical Methods in Natural</i>	We recall the RLWE-based KP-ABE scheme pro-	797
744	<i>Language Processing</i> . <a href="https://openreview.net/pdf?id=EDuKP7DqCk">https://openreview.net/</a>	posed in (Dai et al., 2017) and enhanced with at-	798
745	<a href="https://openreview.net/pdf?id=EDuKP7DqCk">pdf?id=EDuKP7DqCk</a> .	tribute revocability in (Luo et al., 2024).	799
746	Zach Nussbaum, John X. Morris, Brandon Duderstadt,	• $\text{Setup}(1^\lambda, l) \rightarrow \{\text{MPK}, \text{MSK}\}$ where $\lambda$ is the	800
747	and Andriy Mulyar. 2024. Nomic embed: Training a	security parameter, $l$ is the number of user	801
748	reproducible long context text embedder. <i>Preprint</i> ,	maximum attributes, MPK is public master key,	802
749	<a href="https://arxiv.org/abs/2402.01613">arXiv:2402.01613</a> .	and MSK is secret master key.	803
750	PALISADE. 2020. Palisade abe - lattice cryptography	• $\text{Enc}_{\text{ABE}}(m, a, \text{MPK}) \rightarrow ct_a$ where $m$ is the mes-	804
751	library experimental repository. <a href="https://gitlab.com/palisade/palisade-abe">https://gitlab.</a>	age, $a$ the user’s attributes, and $ct_a$ the ouput	805
752	<a href="https://gitlab.com/palisade/palisade-abe">com/palisade/palisade-abe</a> .	ciphertext linked to the user’s attributes.	806
753	Qi Pang, Jinhao Zhu, Helen Möllering, Wenting Zheng,	• $\text{KeyGen}_{\text{ABE}}(\text{MSK}, \text{MPK}, f) \rightarrow \alpha_f$ where $f$ is the	807
754	and Thomas Schneider. 2024. BOLT: Privacy-	circuit policy, which is a boolean circuit, and	808
755	Preserving, Accurate and Efficient Inference for	$\alpha_f$ is the policy decryption key.	809
756	Transformers. In <i>Symposium on Security and Pri-</i>		
757	<i>vacacy (SP)</i> . IEEE. <a href="https://ieeexplore.ieee.org/abstract/document/10646705">https://ieeexplore.ieee.org/</a>		
758	<a href="https://ieeexplore.ieee.org/abstract/document/10646705">abstract/document/10646705</a> .		

- $\text{Eval}_{\text{ABE}}(ct_a, f) \rightarrow ct_{f(a)}$  where  $ct_{f(a)}$  is a ciphertext linked to the policy  $f$ .
- $\text{Dec}_{\text{ABE}}(ct_{f(a)}, \alpha_f, \tilde{a}) \rightarrow \tilde{m}$  or  $\perp$  where  $\tilde{m}$  the recovered message if the attribute  $\tilde{a}$  satisfies the policy  $f$  otherwise the decryption fails  $\perp$ .

## B SecureRAG Security Analysis

Our security analysis follows our threat model discussed in Section 4, where the parties are assumed semi-honest and non-colluding, except for the reader, who is a trusted third party. We recall that semi-honest parties adhere to the protocol as specified but aim to infer sensitive information about other participants solely through their interactions.

**Compromised User.** A semi-honest user, who interacts with the system only by sending queries and receiving responses, may attempt to extract information about other users’ documents. They could do this by crafting malicious queries targeting unauthorized documents in the database. However, SecureRAG prevents such attacks through ABE-encryption of documents. Even if a query matches an unauthorized document, decryption will fail because the document’s ciphertext is bound to a policy  $f$  that excludes the user’s attributes.

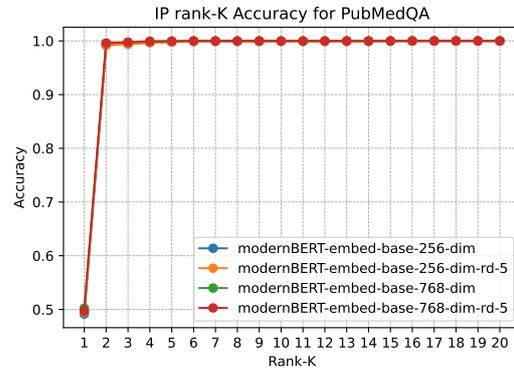
**Compromised Retriever.** A semi-honest retriever that stores the ABE-encrypted documents and their FHE-encrypted embeddings can try to learn the user query, which is FHE-encrypted, the retrieved documents under both forms. For the FHE-encrypted query and embeddings, a compromised retriever cannot infer meaningful information thanks to the IND-CPA security property of FHE, which ensures that FHE ciphertexts remain indistinguishable, even when their underlying plaintexts are identical. The ABE-encrypted documents are fetched using pseudorandom identifiers that look like random values to the retriever. Thus, it cannot learn which documents are fetched. Moreover, the RLWE-based KP ABE scheme that encrypts the documents satisfies the *selective security* (*IND-sCPA*) property (Goyal et al., 2006) that prevents an attacker, who claimed to possess certain attributes from between ciphertexts of two chosen plaintexts as long as they do not satisfy the access policy  $f$ .

**Compromised Generator.** A semi-honest generator that receives encrypted and top-k documents and returns its response encrypted can try to in-

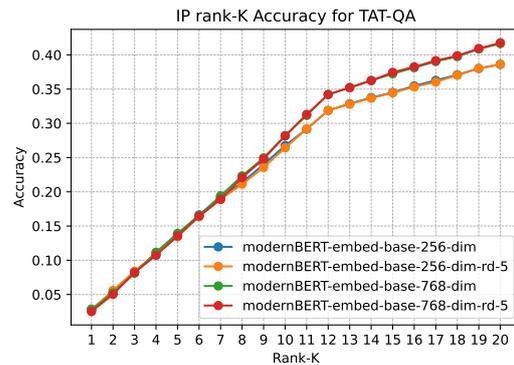
fer information about the query and the selected documents. However, this is prevented by the IND-CPA property of FHE, which entails that the FHE ciphertexts cannot be distinguished even if their underlying plaintexts are identical. Given that, such a compromised generator performs the inference on encrypted data protected by the IND-CPA property, it will be incapable of learning any meaningful information.

## C SecureRAG performance on other datasets

Similarly to Figure 5, Figure 6 shows that rounding does not affect retrieval accuracy, as the ranking curves for the PubMedQA and TAT-QA datasets with and without rounding align almost perfectly, irrespective of embedding dimensionality.



(a) PubMedQA



(b) TAT-QA

Figure 6: SecureRAG’s retriever performance on the PubMedQA and TAT-QA datasets (rounded to 5, orange&red curves) matches unprotected RAG accuracy for rank top-k with  $k \in [1, 20]$ .

## D Prompts for the LLM judge

We used the following prompt from RAGAS (Es et al., 2023) for our evaluation of the context precision metric.

880

### **Context Precision Prompt**

881

882

883

884

- Instruction: Given question, answer and context verify if the context was useful in arriving at the given answer. Give verdict as 1 if useful and 0 if not.

885

886

- Prompt: 'instruction question: question context: context answer: answer verdict: '