# $p^2$RAG: Privacy-Preserving RAG Service Supporting Arbitrary Top-$k$ Retrieval

Yulong Ming [1]  Mingyue Wang [2]  Jijia Yang [1]  Cong Wang [1]  Xiaohua Jia [1]

## Abstract

Retrieval-Augmented Generation (RAG) enables large language models to use external knowledge, but outsourcing the RAG service raises privacy concerns for both data owners and users. Privacy-preserving RAG systems address these concerns by performing secure top-$k$ retrieval, which typically is secure sorting to identify relevant documents. However, existing systems face challenges supporting arbitrary $k$ due to their inability to change $k$, new security issues, or efficiency degradation with large $k$. This is a significant limitation because modern long-context models generally achieve higher accuracy with larger retrieval sets. We propose $p^2$RAG, a privacy-preserving RAG service that supports arbitrary top-$k$ retrieval. Unlike existing systems, $p^2$RAG avoids sorting candidate documents. Instead, it uses an interactive bisection method to determine the set of top-$k$ documents. For security, $p^2$RAG uses secret sharing on two semi-honest non-colluding servers to protect the data owner's database and the user's prompt. It enforces restrictions and verification to defend against malicious users and tightly bound the information leakage of the database. The experiments show that $p^2$RAG is 3–300× faster than the state-of-the-art PRAG for $k = 16$–1024.

## 1. Introduction

Large Language Models (LLMs) are powerful tools, but they suffer from limitations, such as hallucinations (Ji et al., 2023) and lack of real-time or domain-specific data. Retrieval-Augmented Generation (RAG) (Lewis et al., 2020) provides a powerful solution to these issues without the high cost and complexity of fine-tuning. By retrieving relevant information from an external knowledge base and combining it with the user's prompt, RAG enables the model to generate responses that are accurate, up-to-date, and factually grounded.

RAG as a Service (RaaS) has recently emerged as a popular paradigm. In this architecture, the data owner outsources its proprietary database to the RAG service. The user sends a prompt to this service to retrieve the top-$k$ relevant documents from the database. The user then submits the prompt, augmented by these documents, to the model to generate responses. However, this workflow raises privacy issues. The privacy of both the data owner and the user depends on the RAG service's honesty, as it can access the data owner's database and the user's prompt. Moreover, malicious users can attempt to extract proprietary information from the data owner's database.

Privacy-preserving RAG has been proposed to address these issues. Existing systems treat privacy-preserving RAG as a secure top-$k$ retrieval problem. They use techniques from secure $k$-Approximate Nearest Neighbor ($k$-ANN) or secure sorting to compute similarity scores, sort the candidates, and select the top-$k$ documents.

However, existing systems face challenges supporting arbitrary $k$. In particular, some do not support dynamic adjustments to $k$ for user queries, some raise security issues, and others suffer from efficiency degradation with a large $k$ (Servan-Schreiber et al., 2022b; Zyskind et al., 2024). Recent studies show that within a particular yet large context threshold, such as 64K tokens, RAG with a large $k$ generally outperforms RAG with a smaller $k$ for modern long-context models (Leng et al., 2024; Li et al., 2024a;b). This finding implies that $k$ can be increased in RAG, which trades model inference efficiency for accuracy. A large $k$ is also required by applications such as cache-augmented retrieval (Gim et al., 2024; Xu et al., 2024).

**System overview.** We propose $p^2$RAG to address this challenge. Unlike existing systems, $p^2$RAG avoids sorting candidate documents. Instead, we use an interactive bisection method between the user and the servers to determine a threshold. We identify documents within the threshold as the top-$k$ results. This design enables the user to select an arbitrary $k$ by selecting the threshold. Moreover, for RAG applications, the user only requires the set of the top-$k$ documents rather than their specific internal order. Therefore, with this design, the documents within or outside of

[1]Department of Computer Science, City University of Hong Kong, Hong Kong [2]Peng Cheng Laboratory, Shenzhen, Guangdong, China. Correspondence to: Yulong Ming <myl.7@my.cityu.edu.hk>.

the determined threshold remain unsorted. This method decreases the number of comparisons. Because comparisons are the primary computation and communication bottleneck in secure protocols, $p^2$RAG reduces these costs.

$p^2$RAG protects both the data owner's database and the user's prompt. We use secret sharing to ensure that the semi-honest RAG service cannot extract the database or the prompt during the workflow. $p^2$RAG's protocol runs on two semi-honest non-colluding servers where the database and the prompt are secret-shared. Neither server can know any information about the database or the prompt. Moreover, to defend against malicious users, we design restrictions and verifications for the interactions between servers and users. We limit the number of result documents, limit the number of bisection iterations, and verify that the final retrieval of textual documents matches the bisection results. The amount of leakage that a malicious user can extract from the database is tightly bounded.

In summary, our contributions are as follows:

- We propose $p^2$RAG, a privacy-preserving RAG system that supports arbitrary top-$k$ retrieval. $p^2$RAG protects both the data owner's database and the user's prompt. $p^2$RAG uses an interactive bisection method to determine the top-$k$ set without sorting all candidate documents, enabling efficient retrieval even for a large $k$.
- We design $p^2$RAG's security protocol to defend against both the semi-honest RAG service and malicious users. We use secret sharing to protect data privacy against semi-honest servers. We enforce restrictions and verification mechanisms to defend against malicious users and tightly bound the information leakage of the database.
- We implement $p^2$RAG[1] and evaluate its performance. The experiments show that $p^2$RAG achieves higher performance with larger $k$ and is 3–300$\times$ faster than the state-of-the-art system, PRAG, for $k = 16$–$1024$.

## 2. Problem Formulation

### 2.1. System Model

$p^2$RAG aims to provide privacy-preserving RAG as a service to the data owner and users. During the offline (i.e., preprocessing) stage, the data owner outsources its data as an encrypted database on $p^2$RAG. During the online stage, users use encrypted prompts to query top-$k$ relevant documents from the encrypted database.

The data owner's database contains $N$ documents. Each textual document corresponds to an $m$-dimensional document embedding, i.e., vector. Each embedding is generated from the text using an embedding model. It semantically describes the corresponding text, so that the distance between

two embeddings can quantify the similarity of the two texts. We use the cosine distance as the distance metric, which is widely used in existing RAG applications. We assume all embeddings are $\ell_2$-normalized.

$p^2$RAG consists of two servers. We choose the two-server model over more servers because we can use the efficient cryptographic tools optimized for two servers. To outsource the database to two servers, the data owner secret-shares it. That is, for each document embedding, the two servers hold the data owner's two shares, respectively. The sum of the two shares is the embedding, and a single share reveals no information about it. Both texts and embeddings are secret-shared in this way.

A user has a textual prompt and a prompt embedding. The prompt embedding is generated from the text using the same embedding model. We assume that the embedding model is public. We also assume the user can get the prompt embedding using the model in advance, as in prior privacy-preserving RAG work (Zyskind et al., 2024; Cheng et al., 2025). To use $p^2$RAG, the user secret-shares its prompt embedding and sends the two shares to the two servers, respectively. By running $p^2$RAG's protocol, the user finally receives the indices of $k$ documents from the servers. These documents have the $k$ highest cosine similarities to the prompt text.

### 2.2. Threat Model

- The two servers are semi-honest and non-colluding. That is, the servers honestly follow $p^2$RAG's protocol but passively attempt to infer information about the database and prompts, i.e., honest-but-curious. The servers do not share any extra information beyond what is defined in $p^2$RAG's protocol. This assumption is also widely used in other secret-sharing-based security systems (Kamara et al., 2012; Mohassel & Zhang, 2017).
- Users are malicious. Users can deviate from $p^2$RAG's protocol arbitrarily to extract information from the database. Users do not trust the servers.
- Data owners are honest and do not trust any server or user.

### 2.3. Security Goals

We aim to protect the privacy of both the data owner and user, i.e., protect both the database and prompt, resulting in the following goals.

- *Privacy*. No server can get any information about any prompt or document. In particular, because the returned $k$ documents are the $k$ most similar to the prompt and thus can reveal information about the prompt, we also need to protect them from the servers.
- *Bounded database leakage*. No user can get any extra information about any document beyond the "base-

---

[1]

line" leakage and a small amount of leakage defined in $p^2$RAG's protocol. The "baseline" leakage refers to the $k$ documents returned to the user, which is necessary for $p^2$RAG's functionality. $p^2$RAG also leaks counts of documents that are in a particular range to accelerate $p^2$RAG's protocol. This aggregate leakage does not identify any particular document, and the total number of leaked counts is limited, resulting in $O(\log^2 N)$ leakage.

## 3. Preliminary

We use Shamir secret sharing (1979) over a prime field $\mathbb{F}_p$ to protect a value. We denote both shares of a value $x$ by $[x]$. The shares are held by the servers, respectively. We denote each share by $[x]_i$ for $i \in \{0, 1\}$. That is, we protect a value $x$ by secret-sharing it as $[x]_0 + [x]_1 = x$ and sending $[x]_i$ to the server $i$, respectively. Only one $[x]_i$ does not reveal any information about $x$. We have $[x \pm y] = [x] \pm [y]$. That is, for addition and subtraction, each server only needs to compute on its local shares, and no interaction is required.

Distributed Comparison Functions (DCFs) (Boyle et al., 2015; 2016; 2019; 2021) are schemes to secret-share a comparison function. Each function share can be individually executed on a server. The outputs of the function shares are the shares of the output of the original function.

**Definition 3.1** (Comparison Functions). For the input domain $\mathbb{G}_{in} = \{0, \cdots, p-1\}$, a group $(\mathbb{G}_{out}, +)$ as the output domain, $a \in \mathbb{G}_{in}$, and $b \in \mathbb{G}_{out}$, a comparison function $f_{a,b}^<$ is a function that for any input $x$, the output $y$ has $y = b$ only when $x < a$, otherwise $y = 0$.

**Definition 3.2** (Distributed Comparison Functions). For the input domain $\mathbb{G}_{in} = \{0, \cdots, p-1\}$, the output domain $(\mathbb{G}_{out}, +)$, $a \in \mathbb{G}_{in}$, $b \in \mathbb{G}_{out}$, and a security parameter $\lambda$, a DCF is a scheme consisting of the methods:

- Key generation: $Gen(1^\lambda, a, b) \to (k_0, k_1)$.
- Evaluation: $Eval(k_i, j) \to y_{i,j}$ for any $i \in \{0, 1\}$ and any $j \in \mathbb{G}_{in}$.

That satisfies:

- Correctness: $y_{0,j} + y_{1,j} = b$ only when $j < a$, otherwise $y_{0,j} + y_{1,j} = 0$.
- Privacy: Neither $k_0$ nor $k_1$ reveals any information about $a$ and $b$. Formally speaking, there exists a Probabilistic Polynomial Time (PPT) simulator $Sim_{dcf}$ that can generate output computationally indistinguishable from any strict subset of the keys output by $Gen$.

We use DCFs and Boyle et al.'s interval containment gate (2019) to check if a secret-shared value $[x]$ is in an interval $[x_l, x_r)$ as Algorithm 1. Its output is $[1]$ when $x \in [x_l, x_r)$, otherwise $[0]$. Note that this algorithm handles the addition in $\mathbb{G}_{in}$, which is an integer addition modulo $p$.

---

**Algorithm 1** Comparison: $[0]/[1] \leftarrow Cmp([x], [x_l, x_r])$

---

**The key generation method** $Cmp.Gen([x_l, x_r])$:
Sample random $r \in \mathbb{F}_p$ and distribute $[r]$.
$x'_l = x_l + r$, $x'_r = x_r + r$.
$(k_0^l, k_1^l) \leftarrow Gen^<(1^\lambda, x'_l, p-1)$. Distribute $(k_0^l, k_1^l)$.
$(k_0^r, k_1^r) \leftarrow Gen^<(1^\lambda, x'_r, 1)$. Distribute $(k_0^r, k_1^r)$.
Sample random $[w] \in \mathbb{F}_p$ s.t. $[w]_0 + [w]_1 = 1\{x'_l > x'_r\}$.
Distribute $[w]$.

**The evaluation method** $Cmp.Eval([x])$:
Publish $x + r$. $x$ is masked by $r$ and thus protected.
$[y]_i^l = Eval(i, k_i^l, x)$, $[y]_i^r = Eval(i, k_i^r, x)$.
$[y]_i = [y]_i^l + [y]_i^r + [w]_i$. Output $[y]$.

---

We assume that there is a trusted dealer that generates the shared random values required by $p^2$RAG, which is common in security protocols (Beaver, 1992; 1997). The trusted dealer runs during the offline stage and does not participate in the online stage. Therefore, we do not benchmark its performance. For example, the data owner can be the trust dealer, and other techniques can also achieve it (Damgård et al., 2013; Orsini et al., 2020).

## 4. System Design

A user willing to use $p^2$RAG first secret-shares its prompt embedding as $[p]$. It then sends the two shares to the two servers, respectively. The servers first run the distance calculation to calculate the distances $[d_j]$ between the prompt embedding $[p]$ and each document embedding $[x_j]$. The servers then run the interactive distance bisection method with the user. In each iteration, the servers return the count $[c]$ of documents with distances less than a user-specified distance threshold $[d_k]$. The user then updates $d_k$ by comparing the returned number $c$ with the target $k$. When either $c$ is close enough, which is defined by $p^2$RAG, or the iteration number reaches a server-specified limit, the servers return a secret-shared $0/1$ array $\boldsymbol{d}_c$ where the elements corresponding to the in-range documents are set to $1$. The number of these elements, i.e., $c$, is also limited by a server-specified threshold. Finally, during the text retrieval, the user retrieves all $1$'s indices without letting the servers know about the retrieved indices. The servers also verify that the user only retrieves the indices that match the $0/1$ array $\boldsymbol{d}_c$. The workflow is shown as Figure 1.

### 4.1. Distance Calculation

The servers compute the cosine distance $[d_j]$ between the prompt embedding and each document embedding. The document embedding must be $\ell_2$-normalized by the data owner. We assume the $\ell_2$-norms of the database $l = \|\boldsymbol{x}_j\|$ and the prompt $l_u = \|\boldsymbol{p}\|$ are public because they do not leak
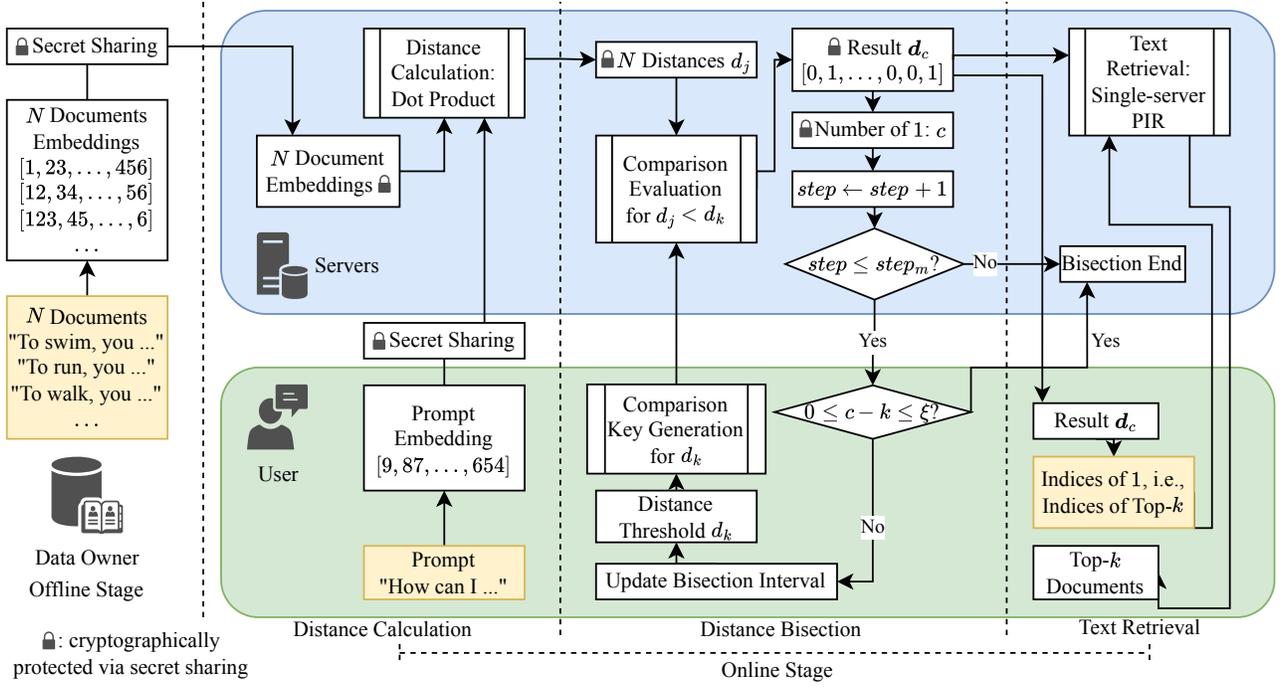
*Figure 1.* The workflow of $p^2$RAG. During the offline stage, the data owner sets up the secret-shared database. During the distance calculation, the servers compute the secret-shared distances between each document and the user's prompt. During the distance bisection, the user determines a distance threshold $d_k$ for the top-$k$ documents. The bisection iteration ends when $d_k$ is found, or the number of iterations exceeds an upper bound. During the text retrieval, the user retrieves textual documents using the indices of the top-$k$ documents.

---

**Algorithm 2** Dot Product: $[c] = [\boldsymbol{a} \cdot \boldsymbol{b}]$

**Offline Stage**:
Sample $m$ random $r_n^{(a)}, r_n^{(b)}, r_n^{(ab)}$ for $n \in [0, m)$ s.t. $r_n^{(a)} \cdot r_n^{(b)} = r_n^{(ab)}$. Distribute $[r_n^{(a)}], [r_n^{(b)}], [r_n^{(ab)}]$.

**Online Stage**:
**for** $k = 0$ to $m$ **do**
    $[d_k] = [a_k] - r_n^{(a)}, [e_k] = [b_k] - [r_n^{(b)}]$.
    Publish $d_k, e_k$.
    $[c_k]_i = [r_n^{(ab)}]_i + e_k \cdot [r_n^{(a)}] + d_k \cdot [r_n^{(b)}] + i \cdot d_k \cdot e_k$.
**end for**
$[c] = \sum_{k=0}^{m} [c_k]$. Output $[c]$.

---

any information about the database or the prompt. Note that the dot product $d_j = \boldsymbol{p} \cdot \boldsymbol{x}_j$ differs from the cosine distance only by the factor $l l_u$. The dot product yields the same ranking as the cosine distance because the factor is constant. Therefore, we use the dot product as $d_j$. We use Beaver triples (1992) for each dimension to compute $[d_j] = [\boldsymbol{p} \cdot \boldsymbol{x}_j]$ from $[\boldsymbol{p}]$ and $[\boldsymbol{x}_j]$ as Algorithm 2. This algorithm has $O(N)$ communication costs, but the amount is $4N$ integers and small, e.g., 16MB for $N = 2^{20}$ and 64-bit integers.

Current embedding models output floating-point numbers that are $\ell_2$-normalized to 1. To represent the embedding elements in $\mathbb{F}_p$, the data owner scales them by a factor $2^f$ and truncates the fractional parts (Catrina & Saxena, 2010), resulting in integers with precision $f$. While the truncation causes the embeddings' norms to deviate slightly from $2^f$, this error is negligible, as we show in Section 7.1. This scale requires that the cardinality of $\mathbb{F}_p$ must be greater than $2^{f+1}$. Note that the product of two scaled integers has $2f$ precision. Therefore, we are required to truncate every product to avoid causing an overflow for the field $\mathbb{F}_p$ (Zyskind et al., 2024). That is, for integers $a, b$, we use the integer part of $ab/2^f$ as the product.

Definition 3.2 requires the input $a \in \mathbb{G}_{in}$ to be an unsigned integer because the internal comparison is performed bit by bit, starting from the most significant bit (Boyle et al., 2021). We extend Algorithm 1 to the signed integers by giving $\mathbb{F}_p$ a new bit representation. The elements $-\lfloor p/2 \rfloor, \cdots, \lfloor p/2 \rfloor$ are represented as the bits of the unsigned integers $0, \cdots, p-1$. This bit representation transformation does not change any algorithm step.

## 4.2. Distance Bisection

Given all $[d_j]$, we need to get the top-$k$ values of them. The naive method is to let the servers return all $d_j$ to the user, who then sorts them. However, though no embeddings are given to the user, letting the user know all $N$ distances leaks information about the database and can even lead to data recovery (Li et al., 2015; Kornaropoulos et al., 2019; Liu et al., 2025). Existing systems avoid this by sorting $[d_j]$ on the servers, which is inefficient in security protocols.

Unlike existing systems, $p^2$RAG runs an interactive bisection method between the servers and the user to determine a distance threshold $d_k$. The user starts with the interval $[-ll_u, ll_u]$ and $d_k = 0$. During each bisection iteration, the servers compare each distance $[d_j]$ with $[d_k]$ and return to the user the count of distances $c$ that are less than $d_k$. To compute $[c]$, the servers compare $[d_k]$ with each $[d_j]$ using Algorithm 1 with the interval $[d_k, p)$. That is, the user runs the key generation method $Cmp.Gen([d_k, p))$, distributing $k_i^l, k_i^r, [w]_i$ to the server $i$, respectively. The servers run the evaluation method $Cmp.Eval([d_j])$ for each $[d_j]$, resulting in $[0]/[1]$. All $[0]/[1]$ outputs of all $[d_j]$ forms a vector $[\boldsymbol{d}_c]$. The servers sum their elements to $[c]$ and return $c$ to the user. To avoid carrying, the cardinality of $\mathbb{F}_p$ must be greater than $N$. The user compares $c$ with $k$ to update the interval and $d_k$. When either the user or the servers stop the iteration, as defined below, the servers return $\boldsymbol{d}_c$ instead of $c$ to the user.

The servers must limit the number of iterations. At each iteration, the user receives the number $c$ as leakage. This small leakage is aggregated from all documents and does not identify any particular document. However, this leakage can accumulate during users' queries. The servers limit the number of iterations with an upper bound $step_m$ to limit the total amount of leakage.

The user can stop the iteration early to speed up the query because: 1) in the bisection method, as the step size becomes very small in later iterations, further steps produce little change. Therefore, stopping early still yields a result close to the true solution. Moreover, 2) recent studies show that within a particular yet large context threshold, such as 64K tokens, retrieving slightly more results in RAG generally outperforms RAG with a smaller $k$ for modern long-context models (Leng et al., 2024; Li et al., 2024a;b). We model this early stop as a slack value $\xi$. The user can stop when $0 \le c \le k + \xi$. We require $\xi \le k$. Otherwise, the extended stopping condition $k \le c \le k + \xi$ can be satisfied by multiple iterations, leading to ambiguity in determining the correct termination point.

After termination, the servers must limit the number of $[1]$ in the final $[\boldsymbol{d}_c]$, i.e., limit the maximum $[c]$, because $[\boldsymbol{d}_c]$ controls which indices can be retrieved by the user during the text retrieval. A user can use $d_k = ll_u + 1$ to get a $[\boldsymbol{d}_c]$ whose elements are all 1, resulting in accessing any document. The servers perform this check by publicly setting up an upper bound $c_m$ and comparing the final $[c]$ with $c_m$ using Algorithm 1. Unlike the comparison during iterations, the key generation method of this comparison is executed offline by the trusted dealer.

A malicious user can deviate from the protocol to manipulate $[\boldsymbol{d}_c]$. When the servers return $c$, this manipulation does not give the user advantages over the leakage of $c$. When the servers return $\boldsymbol{d}_c$, the servers must also verify that all elements of $[\boldsymbol{d}_c]$ are $[0]/[1]$. To achieve that, the servers compare each element with 2 using Algorithm 1. The trusted dealer now executes the key generation method for this comparison offline.

## 4.3. Text Retrieval

Given $\boldsymbol{d}_c$, the user retrieves all its 1's indices to get documents using Private Information Retrieval (PIR). Because the database is stored with secret sharing, the user uses single-server PIR (Ali et al., 2021; Menon & Wu, 2022) to retrieve the documents from both servers. Following prior work (Servan-Schreiber et al., 2022a; Zyskind et al., 2024), we assume that an existing single-server PIR protocol can be used and do not benchmark its performance. Another requirement is that the servers must verify that the user only retrieves the indices that match $[\boldsymbol{d}_c]$. We observe that in single-server PIR, e.g., SimplePIR (Henzinger et al., 2023), there are operations to multiply the database by the query, where each document text is converted to a number. We can multiply this result by $[\boldsymbol{d}_c]$ to enforce the requirement.

## 5. Construction of $p^2$RAG

In this section, we summarize Section 4 to $p^2$RAG's protocol. $p^2$RAG's protocol has the offline and online stages. The offline stage includes setting up the data owner and servers. We do not benchmark its performance because it does not affect the online performance of a user's query. The online stage involves computation and communication between the servers and the user.

During the data owner's offline stage, it has a database filled with textual documents and aims to use $p^2$RAG. It pads all documents to the same length and publishes this length, which is typically required by the text retrieval implemented as single-server PIR. It uses a public embedding model to transform each document to an embedding $\boldsymbol{x}$. The embedding has $m$ dimensions. Its elements are in a prime field $\mathbb{F}_p$ as described in Section 4.1. The document embeddings are $\ell_2$-normalized to $l$. The data owner publishes $l$. The data owner shares each document embedding as $[\boldsymbol{x}]$ and sends them to the servers, respectively.

During the servers' offline stage, they use a trusted dealer

to run the offline stage of the dot product method. They are configured with an iteration step upper bound $step_m$, an upper bound $c_m$ for $c$, and a slack value $\xi$ for the result number. They use a trusted dealer to run $Cmp.Gen$ of Algorithm 1 with the intervals $[0, 2), [0, c_m + 1)$, resulting in $cmpkey^{<2}/cmpkey^{\leq c_m} = (k_i^l, k_i^r, [w]_i)$ for each server $i$.

The user has a textual prompt and aims to retrieve top-$k$ relevant documents. It uses the same public embedding model as the data owner to transform the prompt to an embedding $\boldsymbol{p}$. It shares the embedding as $[\boldsymbol{p}]$ and sends $[\boldsymbol{p}]$ to the servers, respectively, to start the online stage. The online stage is shown as Algorithm 3. The user receives at least $k$ and at most $k + \xi$ indices from $\boldsymbol{d}_c$ as the results. The user can then perform the text retrieval as described in Section 4.3 to retrieve the corresponding textual documents from these indices.

## 6. Analysis

In this section, we analyze $p^2$RAG's performance theoretically and prove $p^2$RAG's security attributes. Our performance analysis covers the online stage described in Section 5, which starts from the user's prompt embedding and ends with the indices of the top-$k$ documents. We quantify the analysis with the computation and communication complexity. Our security analysis focuses on the security goals described in Section 2.3. We use the simulation-based method (Canetti, 2020) to prove that $p^2$RAG achieves the privacy goal. We quantify the information leakage of the database for the bounded database leakage goal.

### 6.1. Performance Analysis

**Computation costs.** For the primitives of Section 3, both DCF key generation and evaluation have $O(\lambda \log N)$ (Boyle et al., 2021), and so for the two methods of the comparison of Algorithm 1. In $p^2$RAG's online stage, the dot product method has $O(m)$ and is executed $N$ times by the servers. The bisection iteration is executed at most $O(\log N)$ times. We denote the number of iterations by $S$. $Cmp.Gen$ is executed $S$ times by the user. $Cmp.Eval$ is executed $SN + N + 1$ times for each server. Therefore, the online stage has $O(\lambda \log^2 N)$ for the user and $O(mN + \lambda N \log^2 N)$ for the servers.

**Communication costs.** For the primitives of Section 3, each key generated by DCFs has $O(\lambda \log N)$ (Boyle et al., 2021), and so for the keys of the comparison. The evaluation method of the comparison of Algorithm 1 has $O(1)$. In $p^2$RAG's online stage, the user sends the prompt embedding share and $2S$ comparison keys, and receives $S$ counts and a share of $\boldsymbol{d}_c$ from each server. While the dot product method has $O(m)$, the two servers run it $m$ times and $Cmp.Eval$ $N$ times as the communication costs. Therefore, the online

**Algorithm 3** Online stage of $p^2$RAG's protocol

The servers compute $[d_j] = [\boldsymbol{p} \cdot \boldsymbol{x}_j]$ for $j \in [0, N)$.
The user sets $d_k = 0$ and $[d_l, d_r] = [-ll_u, ll_u]$. $step = 0$.
**while** true **do**
    The user runs $Cmp.Gen([d_k, p))$ of Algorithm 1, resulting in $cmpkey^{\geq d_k} = (k_i^l, k_i^r, [w]_i)$.
    The user sends $cmpkey^{\geq d_k}$ to the servers, respectively.
    The servers run $Cmp.Eval([d_j])$ of Algorithm 1 for $cmpkey^{\geq d_k}$ to get $\boldsymbol{d}_c$ for $j \in [0, N)$.
    $[c] = \sum_j^N [d_{c,j}]$.
    $step \leftarrow step + 1$.
    **if** $step \geq step_m$ **then**
        **break**
    **end if**
    The servers return $c$ to the user.
    **if** $0 \leq c - k \leq \xi$ **then**
        **break**
    **else if** $c > k$ **then**
        $d_l \leftarrow d_k, d_k \leftarrow (d_l + d_r)/2$.
    **else**
        $d_r \leftarrow d_k, d_k \leftarrow (d_l + d_r)/2$.
    **end if**
**end while**
The servers run $Cmp.Eval([d_{c,j}])$ with $cmpkey^{<2}$ for each element of $\boldsymbol{d}_c$.
The servers publish $cmpkey^{<2}$'s results.
**if** there is any 0 result **then**
    **goto abort**
**end if**
The servers run $Cmp.Eval([c])$ with $cmpkey^{\leq c_m}$.
The servers publish $cmpkey^{\leq c_m}$'s results.
**if** there is any 0 result **then**
    **goto abort**
**end if**
The servers return $\boldsymbol{d}_c$ to the user. **return**
**abort**: The servers frame the user as malicious and abort.

stage has $O(m + N + \lambda \log^2 N)$ between the user and the servers, and $O(mN + N \log N)$ between the two servers.

### 6.2. Security Analysis

**Claim 6.1** ($p^2$RAG's Privacy). *For a Probabilistic Polynomial Time (PPT) adversary $\mathcal{A}$ corrupting the server $i*$ for $i* \in \{0, 1\}$ but following the protocol, Algorithm 3 guarantees that $\mathcal{A}$ learns no information about the prompt and the database.*

*Proof.* The view of $\mathcal{A}$ includes: 1) the prompt embedding share $[p]_{i*}$; 2)$p_n - r_n^{(a)}, x_{j,n} - r_n^{(b)}$ for $n \in [0, m)$ during the distance calculation, where $r_n^{(a)}, r_n^{(b)}$ are random for $n \in [0, m)$; 3) At most $min\{step_m, \lceil \log_2(N/(k + $

$\xi))$]} keys $cmpkey^{\geq d_k}$ for different $d_k$, which are output by $Cmp.Gen$ and sent from the user; 4) $d_j + r^{(d_j)}, c + r^{(c)}, d_{c,j} + r^{(d_{c,j})}$ during $Cmp.Eval$, where $r^{(d_j)}, r^{(c)}, r^{(d_{c,j})}$ are random.

We construct a PPT simulator $Sim$ that generates a view computationally indistinguishable from what $\mathcal{A}$ learns by attacking the protocol, proving $p^2$RAG's privacy goal. For 1), 2), and 4), $Sim$ outputs random values, which are indistinguishable because these values are either secret-shared or masked by random values. For 3), as each $cmpkey$ has $cmpkey = (k^l_{i*}, k^r_{i*}, [w]_{i*})$, $Sim$ outputs a random value for $[w]_{i*}$, which are indistinguishable because it is secret-shared. $Sim$ use $Sim_{dcf}$'s output for $k^l_{i*}, k^r_{i*}$, which are indistinguishable due to DCF's privacy as described in Section 3. □

**Bounded database leakage**. We follow Carlini et al. (2021) and Servan-Schreiber et al.'s (2022a) perspectives to divide the database leakage to the user into two incomparable and complementary parts: 1) physical leakage, which is the leakage of $p^2$RAG's protocol beyond the top-$k$ results; 2) functional leakage, which is the leakage from the top-$k$ results themselves. In particular, $p^2$RAG's functional leakage is at most $k + \xi$ results, which are $O(k(m + \log N))$ bits of information, regardless of $p^2$RAG's protocol (Servan-Schreiber et al., 2022a). This leakage is mentioned as the "baseline" leakage in Section 2.3.

$p^2$RAG's physical leakage is at most $step_m$ numbers of $c$ given to the user, where $0 \leq c \leq N$. To distinguish between $N + 1$ possible outcomes (from 0 to $N$), the maximum information bits contained in $c$ is $\log_2(N+1)$ and $O(\log N)$. We have $step_m \leq \lceil \log_2(N/(k + \xi)) \rceil$ because so many bisection iterations are enough to get the accurate top-$k$ results. That is, $p^2$RAG's physical leakage is $O(\log^2 N)$ bits. We identify $O(\log^2 N)$ as small, proving $p^2$RAG's bounded database leakage goal.

# 7. Experiments

In this section, we benchmark $p^2$RAG's accuracy and performance with experiments. We show that: 1) for accuracy, it has no computation errors, and its results are relevant to the prompt; 2) for performance, it supports arbitrary $k$ and outperforms the state-of-the-art system, PRAG (Zyskind et al., 2024).

We run $p^2$RAG's experiments on a physical on-premises server. The server has two AMD EPYC 7352 CPUs. Each CPU has 24 cores with hyper-threading, for a total of 96 logical cores. The server's memory is 188 GiB, which is sufficient for the experiments. Its operating system is Ubuntu 24.04. We run $p^2$RAG's servers and users one by one on the physical server. We do not simulate the network but count the transmitted data volume because network conditions vary widely in secret-sharing systems and do not affect the system's throughput, as CPUs can run other computational work while waiting for transmission.

We vary $k$ and the document number $N$. We use $k' = k + \xi$ for the number of the retrieved documents, simplifying notations. We set the embedding dimension $m = 1024$. We save embedding elements as 64-bit integers. We set $p = 2^{64} - 59$, which is the largest prime number that fits within 64 bits. We set $f = 32$, matching the precision of 32-bit floating-point numbers. We set $\xi = k$ because $2k$ results of the experiments remain below the 64K tokens threshold, which decreases RAG performance. We set $\lambda = 128$ for DCFs. We use Matyas-Meyer-Oseas one-way compression functions with AES-128 as cryptographically secure PRGs of DCFs for speed (Wang et al., 2017). We use OpenSSL 3 for AES-128.

## 7.1. Accuracy

$p^2$RAG's retrieval is accurate after computing the distances. It has small numerical errors introduced by converting embedding elements to integers and the multiplication of the distance calculation. We use BEIR's `trec-covid` dataset (Kamalloo et al., 2024), which has 171332 documents. We use BEIR's precomputed embeddings from Cohere/beir-embed-english-v3 (Cohere, 2024). We measure the recall between the results from $p^2$RAG's integer distances and the original distances, i.e., the dot products of the floating-point embeddings, resulting in Figure 2. $p^2$RAG's recall keeps 1 for $k' = 16$–$1024$, showing $p^2$RAG's high computation accuracy. Because $p^2$RAG only considers cosine and dot product distances, we measure the results' average relevance score to show if $p^2$RAG retrieves the relevant documents, resulting in Figure 2. The relevance score is a human-annotated categorical level of relevance. $0, 1, 2$ correspond to irrelevant, relevant, and highly relevant, respectively. $p^2$RAG's relevance score is larger than 1 for $k' \leq 300$, meaning all of the retrieved documents are generally the relevant ones.

## 7.2. Performance

$p^2$RAG's performance is independent of the particular document or prompt that is involved in its protocol. This is a necessary attribute for a security system to defend against the timing attack (Kocher, 1996). As a result, we use synthetic datasets for the experiments. We benchmark $p^2$RAG's end-to-end performance. For Algorithm 2 and DCFs used by Algorithm 1, we use multi-threading and start 96 threads. Because the baseline PRAG does not support multi-threading, we run a batch of 96 instances and report the amortized time per instance.

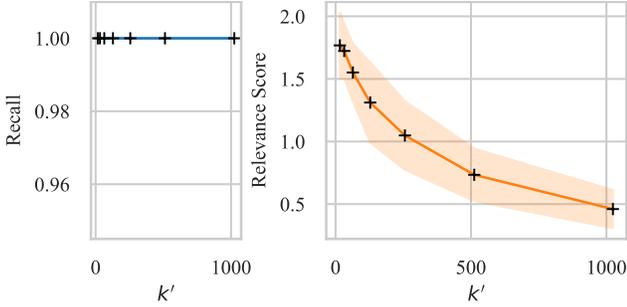**Computation costs**. Figure 3 shows that $p^2$RAG is $3\times$

*Figure 2.* Recall and relevance score (higher is better) with varying numbers of retrieved documents $k'$
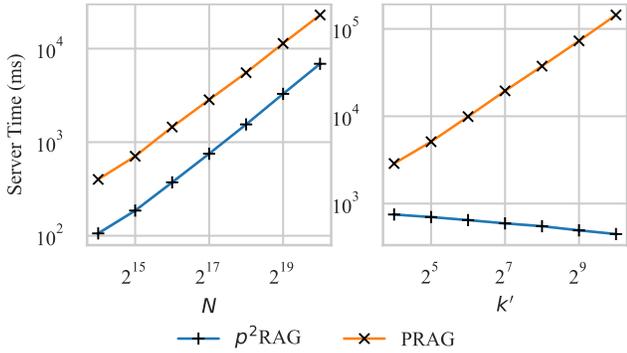


*Figure 3.* Server time (lower is better) with varying numbers of documents $N$ and retrieved documents $k'$. $k' = 16$ and $N = 2^{17}$ for the two figures, respectively.

faster than PRAG when $k' = 16$ and changing $N$, and 7–300× faster when $N = 2^{14}$ and $k = 32$–1024. Both $p^2$RAG and PRAG's time is linear to $N$ and $k'$. However, contrary to the trend in PRAG, $p^2$RAG achieves higher performance with larger $k$.

**Communication costs**. Table 1 shows $p^2$RAG's intra-server (IS) and user-server (US) communication volume, together with the number of Round-Trip Time (RTT). We use $S = \lceil \log_2(N/k') \rceil$ as the number of iterations, simplifying notations. While intra-server communication can reach hundreds of MB, this is well-tolerated via high-bandwidth data center links. The user-server communication remains manageable, as even the largest one is comparable to only a few standard images.

## 8. Related Work

**Privacy-preserving RAG**. Privacy-preserving RAG has the retrieval and inference phases to be protected. During the retrieval phase, both the data owner's database and the user's prompt require protection. $p^2$RAG belongs to the systems protecting the retrieval phase and protecting both the data owner and the user (Zyskind et al., 2024; Bassit & Boddeti, 2025). Some other systems make a weaker

*Table 1.* Communication volume and number of RTTs.

| COMM. TYPE | VOLUME (B) | RTT |
|---|---|---|
| USER-SERVER | $16384 + 4224S + 16N$ | $S + 1$ |
| INTRA-SERVER | $64N + 16NS + 32$ | $S + 1$ |
| $N, k'$ | VOLUME OF US, IS | RTT |
| $2^{17}, 16$ | 2.168MB, 35.65MB | 14 |
| $2^{17}, 128$ | 2.156MB, 29.36MB | 11 |
| $2^{20}, 16$ | 16.86MB, 335.5MB | 17 |

security assumption, protecting only the user's prompt (Cheng et al., 2025; Wang et al., 2025; Hemmat et al., 2025) or the data owner's database (Zhou et al., 2025; Koga et al., 2025; Wu et al., 2025; Mori et al., 2025; Yao & Li, 2025). The other systems focus on the inference phase (Dowlin et al., 2016; Mohassel & Zhang, 2017; Li et al., 2023; Zheng et al., 2024; Thomas et al., 2025)

**Cryptographic tools**. Secure $k$-Approximate Nearest Neighbor (Indyk & Motwani, 1998; Chen et al., 2020; Zuber & Sirdey, 2021; Servan-Schreiber et al., 2022a) and secure sorting (Patel et al., 2012; Ngai et al., 2024; Agarwal et al., 2024; Cong et al., 2025) provide cryptographic tools for privacy-preserving RAG systems. $p^2$RAG's protocol is adapted from the secure sorting based on Function Secret Sharing (FSS) (Boyle et al., 2015; 2016; 2019; 2021) for its communication and computation efficiency. Some other systems use cryptographic or hardware tools such as Differential Privacy (DP) (Dwork et al., 2006; Koga et al., 2025), Homomorphic Encryption (HE) (Gentry, 2009; Bassit & Boddeti, 2025), or Trusted Execution Environment (TEE) (McKeen et al., 2013; Sulich et al., 2025).

## 9. Conclusion

In this paper, we have presented $p^2$RAG, a privacy-preserving RAG service supporting arbitrary top-$k$ retrieval while protecting both the data owner's database and the user's prompt. $p^2$RAG uses secret sharing on two semi-honest non-colluding servers to protect data privacy. $p^2$RAG uses an interactive bisection method to select the top-$k$ relevant documents, without sorting candidate documents. In particular, the bisection method determines a distance threshold by comparing each distance to the threshold using Distributed Comparison Functions (DCFs), while both the threshold and all distances are secret-shared. This design enables the user to choose an arbitrary $k$, even when $k$ is large. To defend against malicious users, $p^2$RAG uses restrictions and verification mechanisms to tightly bound the information leakage of the database. Both the number of bisection iterations and the result set size are limited. The experiments show that $p^2$RAG achieves higher performance with larger $k$ and is 3–300× faster than the state-of-the-art PRAG for $k = 16$–1024.

## Impact Statement

This paper focuses on privacy-preserving Retrieval-Augmented Generation (RAG) systems, aiming to enhance their functionality and efficiency while maintaining the same level of data privacy protection. The potential broader impact of this work includes advancing the deployment of RAG in privacy-critical domains, such as healthcare and finance, where data privacy is a strict requirement.

## References

Agarwal, A., Boyle, E., Chandran, N., Gilboa, N., Gupta, D., Ishai, Y., Kelkar, M., and Ma, Y. Secure sorting and selection via function secret sharing. In *Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security*, CCS '24, pp. 3023–3037, New York, NY, USA, 2024. Association for Computing Machinery. ISBN 9798400706363. doi: 10.1145/3658644.3690359. URL https://doi.org/10.1145/3658644.3690359.

Ali, A., Lepoint, T., Patel, S., Raykova, M., Schoppmann, P., Seth, K., and Yeo, K. Communication–Computation trade-offs in PIR. In *30th USENIX Security Symposium (USENIX Security 21)*, pp. 1811–1828. USENIX Association, August 2021. ISBN 978-1-939133-24-3. URL https://www.usenix.org/conference/usenixsecurity21/presentation/ali.

Bassit, A. and Boddeti, V. SecureRAG: End-to-end secure retrieval-augmented generation. In *The Second Workshop on GenAI for Health: Potential, Trust, and Policy Compliance*, 2025. URL https://openreview.net/forum?id=njAyzNEaCd.

Beaver, D. Efficient multiparty protocols using circuit randomization. In Feigenbaum, J. (ed.), *Advances in Cryptology — CRYPTO '91*, pp. 420–432, Berlin, Heidelberg, 1992. Springer Berlin Heidelberg. ISBN 978-3-540-46766-3.

Beaver, D. Commodity-based cryptography (extended abstract). In *Proceedings of the Twenty-Ninth Annual ACM Symposium on Theory of Computing*, STOC '97, pp. 446–455, New York, NY, USA, 1997. Association for Computing Machinery. ISBN 0897918886. doi: 10.1145/258533.258637. URL https://doi.org/10.1145/258533.258637.

Boyle, E., Gilboa, N., and Ishai, Y. Function secret sharing. In Oswald, E. and Fischlin, M. (eds.), *Advances in Cryptology - EUROCRYPT 2015*, pp. 337–367, Berlin, Heidelberg, 2015. Springer Berlin Heidelberg. ISBN 978-3-662-46803-6.

Boyle, E., Gilboa, N., and Ishai, Y. Function secret sharing: Improvements and extensions. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, CCS '16, pp. 1292–1303, New York, NY, USA, 2016. Association for Computing Machinery. ISBN 9781450341394. doi: 10.1145/2976749.2978429. URL https://doi.org/10.1145/2976749.2978429.

Boyle, E., Gilboa, N., and Ishai, Y. Secure computation with preprocessing via function secret sharing. In Hofheinz, D. and Rosen, A. (eds.), *Theory of Cryptography*, pp. 341–371, Cham, 2019. Springer International Publishing. ISBN 978-3-030-36030-6.

Boyle, E., Chandran, N., Gilboa, N., Gupta, D., Ishai, Y., Kumar, N., and Rathee, M. Function secret sharing for mixed-mode and fixed-point secure computation. In Canteaut, A. and Standaert, F.-X. (eds.), *Advances in Cryptology – EUROCRYPT 2021*, pp. 871–900, Cham, 2021. Springer International Publishing. ISBN 978-3-030-77886-6.

Canetti, R. Universally composable security. *J. ACM*, 67(5), September 2020. ISSN 0004-5411. doi: 10.1145/3402457. URL https://doi.org/10.1145/3402457.

Carlini, N., Deng, S., Garg, S., Jha, S., Mahloujifar, S., Mahmoody, M., Thakurta, A., and Tramèr, F. Is private learning possible with instance encoding? In *2021 IEEE Symposium on Security and Privacy (SP)*, pp. 410–427, 2021. doi: 10.1109/SP40001.2021.00099.

Catrina, O. and Saxena, A. Secure computation with fixed-point numbers. In Sion, R. (ed.), *Financial Cryptography and Data Security*, pp. 35–50, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg. ISBN 978-3-642-14577-3.

Chen, H., Chillotti, I., Dong, Y., Poburinnaya, O., Razenshteyn, I., and Riazi, M. S. SANNS: Scaling up secure approximate k-Nearest neighbors search. In *29th USENIX Security Symposium (USENIX Security 20)*, pp. 2111–2128. USENIX Association, August 2020. ISBN 978-1-939133-17-5. URL https://www.usenix.org/conference/usenixsecurity20/presentation/chen-hao.

Cheng, Y., Zhang, L., Wang, J., Yuan, M., and Yao, Y. RemoteRAG: A privacy-preserving LLM cloud RAG service. In Che, W., Nabende, J., Shutova, E., and Pilehvar, M. T. (eds.), *Findings of the Association for Computational Linguistics: ACL 2025*, pp. 3820–3837, Vienna, Austria, July 2025. Association for Computational Linguistics. ISBN 979-8-89176-256-5. doi: 10.18653/v1/2025.findings-acl.

197. URL https://aclanthology.org/2025.findings-acl.197/.

Cohere. Cohere/beir-embed-english-v3 · Datasets at Hugging Face, January 2024. URL https://huggingface.co/datasets/Cohere/beir-embed-english-v3.

Cong, K., Geelen, R., Kang, J., and Park, J. Revisiting oblivious top-k selection with applications to secure k-nn classification. In *Selected Areas in Cryptography – SAC 2024: 31st International Conference, Montreal, QC, Canada, August 28–30, 2024, Revised Selected Papers, Part I*, pp. 3–25, Berlin, Heidelberg, 2025. Springer-Verlag. ISBN 978-3-031-82851-5. doi: 10.1007/978-3-031-82852-2_1. URL https://doi.org/10.1007/978-3-031-82852-2_1.

Damgård, I., Keller, M., Larraia, E., Pastro, V., Scholl, P., and Smart, N. P. Practical covertly secure mpc for dishonest majority – or: Breaking the spdz limits. In Crampton, J., Jajodia, S., and Mayes, K. (eds.), *Computer Security – ESORICS 2013*, pp. 1–18, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg. ISBN 978-3-642-40203-6.

Dowlin, N., Gilad-Bachrach, R., Laine, K., Lauter, K., Naehrig, M., and Wernsing, J. Cryptonets: applying neural networks to encrypted data with high throughput and accuracy. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48*, ICML'16, pp. 201–210. JMLR.org, 2016.

Dwork, C., McSherry, F., Nissim, K., and Smith, A. Calibrating noise to sensitivity in private data analysis. In Halevi, S. and Rabin, T. (eds.), *Theory of Cryptography*, pp. 265–284, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg. ISBN 978-3-540-32732-5.

Gentry, C. Fully homomorphic encryption using ideal lattices. In *Proceedings of the Forty-First Annual ACM Symposium on Theory of Computing*, STOC '09, pp. 169–178, New York, NY, USA, 2009. Association for Computing Machinery. ISBN 9781605585062. doi: 10.1145/1536414.1536440. URL https://doi.org/10.1145/1536414.1536440.

Gim, I., Chen, G., Lee, S.-s., Sarda, N., Khandelwal, A., and Zhong, L. Prompt cache: Modular attention reuse for low-latency inference. In Gibbons, P., Pekhimenko, G., and Sa, C. D. (eds.), *Proceedings of Machine Learning and Systems*, volume 6, pp. 325–338, 2024.

Hemmat, A., Moqadas, M., Mamanpoosh, A., Rismanchian, A., and Fatemi, A. VAGUE-Gate: Plug-and-Play Local-Privacy shield for Retrieval-Augmented generation. In Inui, K., Sakti, S., Wang, H., Wong, D. F., Bhattacharyya, P., Banerjee, B., Ekbal, A., Chakraborty, T., and Singh, D. P. (eds.), *Proceedings of the 14th International Joint Conference on Natural Language Processing and the 4th Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics*, pp. 3715–3730, Mumbai, India, December 2025. The Asian Federation of Natural Language Processing and The Association for Computational Linguistics. ISBN 979-8-89176-298-5. URL https://aclanthology.org/2025.ijcnlp-long.194/.

Henzinger, A., Hong, M. M., Corrigan-Gibbs, H., Meiklejohn, S., and Vaikuntanathan, V. One server for the price of two: simple and fast single-server private information retrieval. In *Proceedings of the 32nd USENIX Conference on Security Symposium*, SEC '23, USA, 2023. USENIX Association. ISBN 978-1-939133-37-3.

Indyk, P. and Motwani, R. Approximate nearest neighbors: towards removing the curse of dimensionality. In *Proceedings of the Thirtieth Annual ACM Symposium on Theory of Computing*, STOC '98, pp. 604–613, New York, NY, USA, 1998. Association for Computing Machinery. ISBN 0897919629. doi: 10.1145/276698.276876. URL https://doi.org/10.1145/276698.276876.

Ji, Z., Lee, N., Frieske, R., Yu, T., Su, D., Xu, Y., Ishii, E., Bang, Y. J., Madotto, A., and Fung, P. Survey of hallucination in natural language generation. *ACM Comput. Surv.*, 55(12), March 2023. ISSN 0360-0300. doi: 10.1145/3571730. URL https://doi.org/10.1145/3571730.

Kamalloo, E., Thakur, N., Lassance, C., Ma, X., Yang, J.-H., and Lin, J. Resources for brewing beir: Reproducible reference models and statistical analyses. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '24, pp. 1431–1440, New York, NY, USA, 2024. Association for Computing Machinery. ISBN 9798400704314. doi: 10.1145/3626772.3657862. URL https://doi.org/10.1145/3626772.3657862.

Kamara, S., Mohassel, P., and Riva, B. Salus: a system for server-aided secure function evaluation. In *Proceedings of the 2012 ACM Conference on Computer and Communications Security*, CCS '12, pp. 797–808, New York, NY, USA, 2012. Association for Computing Machinery. ISBN 9781450316514. doi: 10.1145/2382196.2382280. URL https://doi.org/10.1145/2382196.2382280.

Kocher, P. C. Timing attacks on implementations of diffie-hellman, rsa, dss, and other systems. In Koblitz, N. (ed.), *Advances in Cryptology — CRYPTO '96*, pp. 104–113, Berlin, Heidelberg, 1996. Springer Berlin Heidelberg. ISBN 978-3-540-68697-2.

Koga, T., Wu, R., Zhang, Z., and Chaudhuri, K. Privacy-preserving retrieval-augmented generation with differential privacy, 2025. URL https://arxiv.org/abs/2412.04697.

Kornaropoulos, E. M., Papamanthou, C., and Tamassia, R. Data recovery on encrypted databases with k-nearest neighbor query leakage. In *2019 IEEE Symposium on Security and Privacy (SP)*, pp. 1033–1050, 2019. doi: 10.1109/SP.2019.00015.

Leng, Q., Portes, J., Havens, S., Zaharia, M., and Carbin, M. Long context RAG performance of large language models. In *Adaptive Foundation Models: Evolving AI for Personalized and Efficient Learning*, 2024. URL https://openreview.net/forum?id=Le9anH3kv1.

Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W.-t., Rocktäschel, T., Riedel, S., and Kiela, D. Retrieval-augmented generation for knowledge-intensive nlp tasks. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, NIPS '20, Red Hook, NY, USA, 2020. Curran Associates Inc. ISBN 9781713829546.

Li, D., Wang, H., Shao, R., Guo, H., Xing, E., and Zhang, H. MPCFORMER: FAST, PERFORMANT AND PRIVATE TRANSFORMER INFERENCE WITH MPC. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=CWmvjOEhgH-.

Li, F., Shin, R., and Paxson, V. Exploring privacy preservation in outsourced k-nearest neighbors with multiple data owners. In *Proceedings of the 2015 ACM Workshop on Cloud Computing Security Workshop*, CCSW '15, pp. 53–64, New York, NY, USA, 2015. Association for Computing Machinery. ISBN 9781450338257. doi: 10.1145/2808425.2808430. URL https://doi.org/10.1145/2808425.2808430.

Li, X., Cao, Y., Ma, Y., and Sun, A. Long context vs. rag for llms: An evaluation and revisits, 2024a. URL https://arxiv.org/abs/2501.01880.

Li, Z., Li, C., Zhang, M., Mei, Q., and Bendersky, M. Retrieval augmented generation or long-context LLMs? a comprehensive study and hybrid approach. In Dernoncourt, F., Preoţiuc-Pietro, D., and Shimorina, A. (eds.), *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing: Industry Track*, pp. 881–893, Miami, Florida, US, November 2024b. Association for Computational Linguistics. doi: 10.18653/v1/2024.emnlp-industry.66. URL https://aclanthology.org/2024.emnlp-industry.66/.

Liu, Y., Zhang, Y., Xie, J., Li, H., Yu, J. X., and Cui, J. Privacy-Preserving Approximate Nearest Neighbor Search on High-Dimensional Data . In *2025 IEEE 41st International Conference on Data Engineering (ICDE)*, pp. 3017–3029, Los Alamitos, CA, USA, May 2025. IEEE Computer Society. doi: 10.1109/ICDE65448.2025.00226. URL https://doi.ieeecomputersociety.org/10.1109/ICDE65448.2025.00226.

McKeen, F., Alexandrovich, I., Berenzon, A., Rozas, C. V., Shafi, H., Shanbhogue, V., and Savagaonkar, U. R. Innovative instructions and software model for isolated execution. In *Proceedings of the 2nd International Workshop on Hardware and Architectural Support for Security and Privacy*, HASP '13, New York, NY, USA, 2013. Association for Computing Machinery. ISBN 9781450321181. doi: 10.1145/2487726.2488368. URL https://doi.org/10.1145/2487726.2488368.

Menon, S. J. and Wu, D. J. SPIRAL: Fast, High-Rate Single-Server PIR via FHE Composition. In *2022 IEEE Symposium on Security and Privacy (SP)*, pp. 930–947. IEEE Computer Society, May 2022. ISBN 978-1-6654-1316-9. doi: 10.1109/SP46214.2022.9833700. URL https://www.computer.org/csdl/proceedings-article/sp/2022/131600b568/1FlQGRX73Uc.

Mohassel, P. and Zhang, Y. Secureml: A system for scalable privacy-preserving machine learning. In *2017 IEEE Symposium on Security and Privacy (SP)*, pp. 19–38, 2017. doi: 10.1109/SP.2017.12.

Mori, J., Kakizaki, K., Miyagawa, T., and Sakuma, J. Differentially private synthetic text generation for retrieval-augmented generation (rag), 2025. URL https://arxiv.org/abs/2510.06719.

Ngai, N., Demertzis, I., Chamani, J. G., and Papadopoulos, D. Distributed & scalable oblivious sorting and shuffling. In *2024 IEEE Symposium on Security and Privacy (SP)*, pp. 4277–4295, 2024. doi: 10.1109/SP54263.2024.00153.

Orsini, E., Smart, N. P., and Vercauteren, F. Overdrive2k: Efficient secure mpc over from somewhat homomorphic encryption. In *Topics in Cryptology – CT-RSA 2020: The Cryptographers' Track at the RSA Conference 2020, San Francisco, CA, USA, February 24–28, 2020, Proceedings*, pp. 254–283, Berlin, Heidelberg, 2020. Springer-Verlag. ISBN 978-3-030-40185-6. doi: 10.1007/978-3-030-40186-3_12. URL https://doi.org/10.1007/978-3-030-40186-3_12.

Patel, S., Garasia, S., and Jinwala, D. An efficient approach for privacy preserving distributed k-means clustering based on shamir's secret sharing scheme. In Dim-

itrakos, T., Moona, R., Patel, D., and McKnight, D. H. (eds.), *Trust Management VI*, pp. 129–141, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg. ISBN 978-3-642-29852-3.

Servan-Schreiber, S., Langowski, S., and Devadas, S. Private approximate nearest neighbor search with sublinear communication. In *2022 IEEE Symposium on Security and Privacy (SP)*, pp. 911–929, 2022a. doi: 10.1109/SP46214.2022.9833702.

Servan-Schreiber, S., Langowski, S., and Devadas, S. Private approximate nearest neighbor search with sublinear communication. In *2022 IEEE Symposium on Security and Privacy (SP)*, pp. 911–929, 2022b. doi: 10.1109/SP46214.2022.9833702.

Shamir, A. How to share a secret. *Commun. ACM*, 22 (11):612–613, November 1979. ISSN 0001-0782. doi: 10.1145/359168.359176. URL https://doi.org/10.1145/359168.359176.

Sulich, B., O'Neill, P., and Schrater, J. Securing Enterprise RAG Deployments, November 2025. URL https://www.intel.com/content/www/us/en/content-details/870124/securing-enterprise-rag-deployments.html.

Thomas, R. K., Zahran, L., Choi, E., Potti, A., Goldblum, M., and Pal, A. Hidden no more: Attacking and defending private third-party LLM inference. In *Forty-second International Conference on Machine Learning*, 2025. URL https://openreview.net/forum?id=QfD9P9IIoz.

Wang, B., Lou, Q., Zheng, M., and Zhao, D. Pir-rag: A system for private information retrieval in retrieval-augmented generation, 2025. URL https://arxiv.org/abs/2509.21325.

Wang, F., Yun, C., Goldwasser, S., Vaikuntanathan, V., and Zaharia, M. Splinter: Practical private queries on public data. In *14th USENIX Symposium on Networked Systems Design and Implementation (NSDI 17)*, pp. 299–313, Boston, MA, March 2017. USENIX Association. ISBN 978-1-931971-37-9. URL https://www.usenix.org/conference/nsdi17/technical-sessions/presentation/wang-frank.

Wu, R., Wang, E., Zhang, Z., and Wang, Y.-X. Private-rag: Answering multiple queries with llms while keeping your data private, 2025. URL https://arxiv.org/abs/2511.07637.

Xu, P., Ping, W., Wu, X., McAfee, L., Zhu, C., Liu, Z., Subramanian, S., Bakhturina, E., Shoeybi, M., and Catanzaro, B. Retrieval meets long context large language models. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=xw5nxFWMlo.

Yao, D. and Li, T. Differentially private retrieval augmented generation with random projection. In *ICLR 2025 Workshop on Building Trust in Language Models and Applications*, 2025. URL https://openreview.net/forum?id=5DfhoxRPXh.

Zheng, F., Chen, C., Han, Z., and Zheng, X. Permllm: Private inference of large language models within 3 seconds under wan, 2024. URL https://arxiv.org/abs/2405.18744.

Zhou, P., Feng, Y., and Yang, Z. Provably secure retrieval-augmented generation, 2025. URL https://arxiv.org/abs/2508.01084.

Zuber, M. and Sirdey, R. Efficient homomorphic evaluation of k-NN classifiers. *Proceedings on Privacy Enhancing Technologies*, 2021. ISSN 2299-0984. URL https://petsymposium.org/popets/2021/popets-2021-0020.php.

Zyskind, G., South, T., and Pentland, A. Don't forget private retrieval: distributed private similarity search for large language models. In Habernal, I., Ghanavati, S., Ravichander, A., Jain, V., Thaine, P., Igamberdiev, T., Mireshghallah, N., and Feyisetan, O. (eds.), *Proceedings of the Fifth Workshop on Privacy in Natural Language Processing*, pp. 7–19, Bangkok, Thailand, August 2024. Association for Computational Linguistics. URL https://aclanthology.org/2024.privatenlp-1.2/.