



Can You Trust the Vectors in Your Vector Database? Black-Hole Attack from Embedding Space Defects

Hanxi Li
Sichuan University
China
hanncie@outlook.com

Jianan Zhou
Sichuan University
China
jiananzhou1206@gmail.com

Jiale Lao
Cornell University
USA
jiale@cs.cornell.edu

Yibo Wang
Purdue University
USA
wang7342@purdue.edu

Zhengmao Ye
Sichuan University
China
yezhenmaolove@gmail.com

Yang Cao
Institute of Science Tokyo
Japan
cao@c.titech.ac.jp

Junfen Wang
Sichuan University
China
wangjf@scu.edu.cn

Mingjie Tang*
Sichuan University
China
tangrock@gmail.com

ABSTRACT

Vector databases serve as the retrieval backbone of modern AI applications, yet their security remains largely unexplored. We propose the **Black-Hole Attack**, a poisoning attack that injects a small number of malicious vectors near the geometric center of the stored vectors. These injected vectors attract queries like a black hole and frequently appear in the top- k retrieval results for most queries. This attack is enabled by a phenomenon we term *centrality-driven hubness*: in high-dimensional embedding spaces, vectors near the centroid become nearest neighbors of a disproportionately large number of other vectors, while this centroid region is nearly empty in practice. The attack shows that vectors in a vector database cannot be blindly trusted: geometric defects in high-dimensional embeddings make retrieval inherently vulnerable. Based on this insight, we propose four attack paths tailored to different attacker capabilities. Our experiments show that up to 94.4% of queries are successfully attacked. Additionally, we study two directions of defense: hubness mitigation and detection-based filtering. Hubness mitigation either significantly reduces retrieval accuracy or provides only limited protection, while the detection-based defense is effective against some attack paths but fails against others. A robust and adaptive defense thus remains an open problem, and our findings indicate that vector databases require more careful treatment of security.

PVLDB Reference Format:

Hanxi Li, Jianan Zhou, Jiale Lao, Yibo Wang, Zhengmao Ye, Yang Cao, Junfen Wang, and Mingjie Tang. Can You Trust the Vectors in Your Vector Database? Black-Hole Attack from Embedding Space Defects. PVLDB, 20(1): 31 - 45, 2026.
doi:10.14778/3845598.3845601

*Corresponding author.

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.
Proceedings of the VLDB Endowment, Vol. 20, No. 1 ISSN 2150-8097.
doi:10.14778/3845598.3845601

PVLDB Artifact Availability:

The source code have been made available at https://github.com/hanxi19/Black_Hole_Attack_for_Vector_Database.

1 INTRODUCTION

Vector databases have gained significant attention due to their ability to address limitations of Large Language Models (LLMs) [35, 53, 77]. For example, a vector database can serve as an external knowledge source by providing domain-specific or up-to-date information that is not included in the LLM training corpus, thereby enabling LLMs to generate more accurate answers [18, 25]. This is achieved by vector similarity search [21, 24, 50], which retrieves the top- k most similar vectors to a given query based on a similarity metric over embeddings.

Extensive work has been proposed to optimize the performance of vector databases [41]. To support real-world use cases, these systems must efficiently store and search millions to billions of embeddings while supporting frequent updates, high concurrency, and distributed search [21, 24, 36, 43, 50]. Examples include Milvus [66], Weaviate [69], Pinecone [45], Vespa [5], and many more systems [41].

Despite extensive research on performance optimization for vector databases, security issues remain largely underexplored. For example, numerous ready-to-use datasets that contain knowledge and pre-computed embeddings are available on open-source platforms such as Hugging Face [33]. These datasets are usually domain-specific, such as financial, medical, and e-commerce data, and are often used as external knowledge sources for Retrieval-Augmented Generation (RAG) applications [25, 32], semantic search [7, 28], and knowledge base systems [42]. Users often assume that these embeddings are meaningful semantic representations of the underlying knowledge and directly integrate such datasets into downstream applications. However, attackers can create poisoned datasets or upload modified versions of existing datasets that contain harmful content and manipulated embeddings. Applications that rely on

these datasets may then be compromised by such poisoned inputs. This raises a critical question: *Can we trust the vectors stored in a vector database?*

We systematically investigate this problem and show that launching such an attack is non-trivial. To remain stealthy, an attacker can inject only a small number of harmful contents and malicious vectors. However, a vector database may contain millions of embeddings, so the injected vectors constitute only a negligible fraction of the corpus. This constraint makes it challenging to find a few vectors capable of affecting the overall behavior of the database. Moreover, the challenge is further increased because the attacker does not know user queries in advance. Existing corpus poisoning methods in RAG systems [9, 19, 26, 79] rely on prior knowledge of user queries to craft adversarial texts, which is difficult to obtain in practice. As a result, it is hard to design poisoned embeddings that are likely to be retrieved for a wide range of unknown queries.

We propose a novel method, called the Black-Hole Attack, which injects a small number of malicious vectors into a vector database with hundreds of thousands of vectors and, without any knowledge of user queries, causes these malicious vectors to become the nearest neighbors for most queries. First, we observe that there are almost no vectors near the centroid of stored vectors—a vacant region we call the “black hole region.” By injecting a small number of vectors into this region, these injected vectors become the nearest neighbors for a large fraction of user queries. Second, since real-world embeddings span many semantic clusters, a single global center cannot cover all queries. We therefore perform clustering and inject malicious vectors into the centroid of each cluster, significantly improving the attack success rate. Third, we provide a theoretical explanation for why this simple strategy is so effective: in a high-dimensional embedding space with finite data, vectors near the centroid are inherently closer to most other vectors than any existing entry—a property we term *centrality-driven hubness*. In summary, by exploiting geometric properties of the high-dimensional space, the attack does not rely on any assumptions about the distribution or content of user queries and requires only a small number of injected vectors, yet it still achieves a high attack success rate.

While the Black-Hole Attack is highly effective, its practical feasibility depends on what the attacker can access. Rather than assuming a fixed capability, we systematically characterize four attack paths along a capability spectrum and show that the Black-Hole Attack remains viable across all of them. At the highest capability level, an attacker with full database export can directly use all stored embeddings to compute the centroid and inject malicious vectors. It is the most basic and direct path, but not necessarily the most effective one. When only a partial export is available, the attacker exports only a fraction of the stored vectors; varying the export ratio from 0.1 to 1.0, we find that 10% leakage already yields more than half of the full-export MO@10. For attackers with no database export capability but access to the same embedding model, we propose a Surrogate Dataset Attack: the attacker aggregates queries from publicly available datasets, encodes them with the same embedding model, and builds malicious vectors from these surrogate embeddings. And we found that an attacker does not need to accurately guess what users will ask to launch an effective attack. Finally, at the lowest capability level, an attacker can exploit

poisoned public pre-embedded datasets by injecting malicious vectors into publicly available datasets before they are downloaded and indexed, requiring neither model access nor database access. These attack paths demonstrate that the Black-Hole Attack is not a single method but a flexible framework that adapts to the attacker’s available resources, spanning from full database export down to no target access at all.




Figure 1: The Workflow of the Black-Hole Attack

Example 1.1. Figure 1 shows the workflow of the Black-Hole Attack. The attacker has access to the target vector database’s embeddings (by exporting full or partial vectors) or to open-source datasets that contain queries. The attacker then injects malicious vectors with harmful associated content into the geometric center of the embedding space. When users submit queries to LLMs, RAG retrieves relevant knowledge from the database to improve the quality of LLM responses. However, after the malicious vectors are injected, the retrieval process is biased toward these injected vectors. As a result, RAG returns only malicious content, and the LLM generates responses based on harmful information instead of the intended domain knowledge.

In summary, we make the following contributions:

- (1) We provide a theoretical analysis and experimental validation showing that, in a high-dimensional embedding space with finite data, vectors located near the geometric center have a high probability of being selected as the nearest neighbors of many other vectors.
- (2) We propose the Black-Hole Attack, a novel poisoning attack that injects a small number of malicious vectors near the geometric centroid of the embedding space to hijack retrieval for most user queries without any assumption on query distribution or content. To accommodate realistic attacker constraints, we further define four attack paths along a capability spectrum and show that the attack remains highly effective under each scenario.
- (3) We conduct extensive experiments across multiple embedding models, datasets, and retrieval settings, demonstrating that the Black-Hole Attack achieves consistently high attack effectiveness and strong robustness. Beyond retrieval-level metrics, we further analyze the downstream impact

of the attack on real-world RAG applications, revealing its practical threat to end-to-end system behavior.

- (4) We study two directions of defense against the Black-Hole Attack: hubness mitigation and detection-based filtering. Hubness mitigation either significantly reduces retrieval accuracy or provides only limited protection. The detection-based defense is effective against some attack paths but fails against others, so a robust and adaptive defense remains an open problem.

2 BACKGROUND AND EXISTING THREATS

Figure 2 illustrates the workflow of a vector database system. During query execution, the *embedding model* first converts the user’s raw input into a query vector in the *embedding space*. The *index* then performs a nearest neighbor search over the stored vectors and returns the identifiers of the top- k closest matches. Finally, the corresponding records are fetched from *storage* and presented to the user.




Figure 2: Workflow and attack process of vector database

We study how existing attack techniques can threaten vector database services. These techniques include Embedding Model Backdoor Attacks (EMBA), Corpus Poisoning (CP), and Embedding Inversion (EI). As shown in Figure 2, they mainly target the embedding model and the storage layer, and they require prior knowledge of user queries or the embedding model to perform the attack.

R1: Embedding Model Backdoor Attack (EMBA). Embedding Model Backdoor Attacks [2, 14, 65, 71] manipulate upstream embedding models during training or fine-tuning. These attacks poison the training data to insert hidden backdoors. When the model encounters a specific trigger (e.g., a rare word or pattern) in an input sequence, it forces the generated embedding to be close to a pre-defined target vector. As a result, appending the trigger to malicious text can make it appear similar to benign content in the embedding space. This attack mainly targets the **Embedding Model**, because it changes the model’s ability to produce semantically accurate embeddings for certain inputs. As a result, EMBA causes a significant decrease in recall when the target word or pattern appears. The vector database then returns records that contain the backdoor trigger, which hides normal and relevant records.

R2: Corpus Poisoning (CP). CP mainly targets the Storage Layer by inserting a small set of carefully crafted malicious records into

the vector database [9, 19, 26, 79]. Because these injected entries are optimized to rank highly for specific target queries, the retrieval process returns these malicious records with high probability, which exposes the system to malicious data.

R3: Embedding Inversion (EI). EI attacks attempt to reconstruct the original data directly from the vector embeddings stored in a database [39, 74]. Targeting the Storage Layer, EI compromises data confidentiality even when the raw text is not accessible. Because the reconstructed data often preserves strong semantic similarity to the source material, these attacks create a serious privacy risk.

Notably, these attack techniques originate from related domains such as RAG systems, NLP, and general machine learning. They are constrained to specific target queries or triggers and generalize poorly to unseen queries. In contrast, our method exploits geometric properties of the high-dimensional embedding space. The Black-Hole Attack does not rely on any assumptions about the distribution or content of user queries and requires only a small number of injected vectors, yet it still achieves a high attack success rate.

3 THREAT MODEL AND OVERVIEW




Figure 3: Black-Hole Attack overview

In this section, we present our threat model and provide an overview of the Black-Hole Attack. We consider, but are not limited to, a semantic retrieval system built on top of a vector database, which stores high-dimensional embedding vectors produced by an upstream embedding model and supports ANN-based similarity search.

Attacker’s Goal. The attacker aims to inject a small number of malicious vectors into the vector database so that these vectors frequently appear in the top- k retrieval results for most user queries. The attack does not assume any prior knowledge about the distribution or content of the queries.

Attacker’s Capabilities. We outline potential threat vectors that the attacker could leverage to compromise a vector database, without accessing real user queries or modifying the embedding model:

- (1) Full Database Export. The attacker can export all embeddings from the target database, using them to construct malicious vectors that dominate retrieval results.
- (2) Partial Database Export. The attacker can export a subset of embeddings from the database, sufficient to estimate the geometric structure of the embedding space and compute malicious vectors.

(3) Surrogate Dataset Attack. The attacker can aggregate queries from as many publicly available datasets as possible, compute embeddings on this surrogate data with the same embedding model, and use them to generate malicious vectors that are injected into the target database.

(4) Poisoning Public Pre-Embedded Datasets. The attacker can download publicly available pre-embedded datasets and create malicious vectors, inject them into the datasets and upload the poisoned datasets.

Attack Overview. The Black-Hole Attack is a query-agnostic poisoning attack for vector databases. It injects a small number of malicious vectors that dominate the top- k retrieval results for most user queries. Figure 3 presents the attack workflow that involves four steps. The attacker ① obtains the embedding vectors through one of the following means: exporting all or a subset of embeddings from the target vector database, collecting publicly available query datasets and embedding them as surrogate vectors, or directly downloading open-source pre-embedded datasets, ② partitions the vectors into clusters to capture the semantic structure of the embedding space, ③ computes the geometric centroid of each cluster to identify the black-hole region where almost no benign vectors reside, and ④ constructs malicious vectors in a tight neighborhood around each centroid, associates them with harmful content (e.g., phishing links or misinformation), and injects them into the target database or dataset. When unsuspecting users subsequently issue queries, the retrieval system returns the injected malicious vectors as top- k nearest neighbors instead of legitimate results, leading to consequences such as misinformation propagation or phishing exposure.

4 THE BLACK-HOLE ATTACK

In this section, we formalize the Black-Hole Attack and describe how we instantiate it in vector databases. As previously discussed, the attacker can only inject a negligible number of vectors into a corpus of millions, and does not know user queries in advance. The central question is: *how can a handful of poisoned vectors attract the majority of arbitrary queries?* We find a surprisingly simple answer: by placing malicious vectors near the centroid of stored vectors, they become closer to queries than any existing vectors, and thus are retrieved as top- k results.

4.1 The Black-Hole Region

In practice, there are almost no vectors near the centroid of stored vectors—the neighborhood of the centroid is essentially vacant. We refer to this vacant region as the *black-hole region*.

This phenomenon is also remarkably robust. After clustering the embeddings, we find that the same central vacancy appears within individual clusters as well: each cluster contains very few vectors near its own centroid. Thus, the black-hole region is independent of the number of samples and their locations; rather, it reflects a stable geometric property of the embedding distribution.

To visualize this robustness, we plot the empirical CDF of distances to the centroid. We consider both (i) distances to the global centroid of the entire database and (ii) distances to the cluster centroid after applying k -means. Figure 4 shows that, under both Euclidean and cosine distances, the CDF stays near zero over a

non-trivial range in both settings, indicating an effectively empty neighborhood not only around the global centroid but also around cluster centroids.

This vacant region makes the Black-Hole Attack possible: vectors inserted here are free from interference by existing data, yet their central position causes them to attract queries like a black hole. We next describe how to construct such malicious vectors.




Figure 4: Empirical CDF of the distance-to-centroid under Euclidean (left) and cosine (right) distances, measured on HotpotQA embeddings generated by Contriever.

4.2 Constructing the Malicious Vectors

As shown above, the geometric center of stored vectors is almost empty. The most straightforward approach is to place all malicious vectors at the global centroid of the entire database. Let $V = \{v_i\}_{i=1}^N \subset \mathbb{R}^d$ denote the set of benign database vectors. Given a poisoning rate $\alpha \in (0, 1)$, the attacker may insert up to αN additional vectors $V_{\text{poison}} = \{v_{p,j}\}_{j=1}^{\alpha N}$ into the database, forming the poisoned set $V' = V \cup V_{\text{poison}}$. The global centroid is defined as:

$$c_{\text{centroid}} = \frac{1}{N} \sum_{i=1}^N v_i. \quad (1)$$

The attacker constructs a poisoned set V_{poison} by sampling points in a tight neighborhood around this center:

$$v_{p,j} = c_{\text{centroid}} + \varepsilon_j, \quad j = 1, \dots, \alpha N, \quad (2)$$

where $\varepsilon_j \in \mathbb{R}^d$ are small random perturbations. This global strategy is entirely query-agnostic, because it relies only on the geometry of the stored embeddings. Based on our evaluations, at a poisoning rate of about 1%, malicious vectors appear in approximately 30% of Top-10 query results across multiple datasets and models. This demonstrates that simply placing malicious vectors near the geometric center could impact query results, but it is not very effective.

Real-world embedding distributions span many semantic clusters by topic or content type, so a single global centroid cannot attract queries originating from all regions of the space [44, 64]. To improve coverage, we partition the benign vector set V into L clusters $\{C_1, \dots, C_L\}$ using k -means and compute the centroid of each cluster:

$$c_j = \frac{1}{|C_j|} \sum_{v \in C_j} v, \quad j = 1, \dots, L.$$

Around each c_j we construct a local poisoned set $V_{\text{poison}}^{(j)}$ by sampling vectors in a tight neighborhood of c_j , and form the overall poisoned set $V_{\text{poison}} = \bigcup_{j=1}^L V_{\text{poison}}^{(j)}$.

Empirically, the cluster-wise design substantially amplifies the attack’s impact. In our experiments, with a poisoning rate of only about 1%, malicious vectors appear in roughly 80% of Top-10 results across datasets and embedding models. This demonstrates that by exploiting local centroids across distinct semantic clusters, the Black-Hole Attack significantly dominates the majority of queries.

4.3 Surrogate Dataset Attack

The cluster-wise attack in Section 4.2 requires the attacker to access the target database’s stored embeddings. However, this capability may not always be available in practice—for instance, when the database is hosted in a controlled environment that does not expose its internal vectors. To address this constraint, we introduce the *Surrogate Dataset Attack*, which constructs malicious vectors using only publicly available data, without any access to the target database’s embeddings.

Our key observation is that query embeddings and corpus embeddings exhibit a systematic distributional gap in the embedding space: queries are typically in relatively fixed forms such as interrogative sentences or entity phrases. Although queries from different datasets cover diverse topics, they tend to cluster in similar regions that are offset from the corpus distribution. This means that queries collected from one dataset can serve as effective proxies for the queries of another dataset.

Given this insight, the Surrogate Dataset Attack proceeds as follows. The attacker collects query texts from as many publicly available datasets as possible and encodes them using the same embedding model used by the target vector database. Let $Q_{\text{surr}} = q_1, \dots, q_M \subset \mathbb{R}^d$ denote the resulting surrogate query embeddings. The attacker then partitions Q_{surr} into L clusters using k -means and computes the centroid of each cluster:

$$c_j^{\text{surr}} = \frac{1}{|C_j^{\text{surr}}|} \sum_{q \in C_j^{\text{surr}}} q, \quad j = 1, \dots, L.$$

Malicious vectors are then generated in a tight neighborhood around each surrogate centroid:

$$v_{p,j} = c_j^{\text{surr}} + \varepsilon_j, \quad \varepsilon_j \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}_d),$$

and injected into the target database. Note that the attacker never accesses the target database’s internal vectors—it only needs surrogate query texts and the corresponding embedding model.

We validate the above insight and evaluate the Surrogate Dataset Attack in Section 6 (see Figure 7). Our analysis of the embedding distributions validates that query distributions indeed transfer across datasets. Meanwhile, across 14 surrogate–target dataset pairs, the attack achieves significant impact in 10 cases. The remaining 4 pairs show limited effectiveness, and in each case the failure stems from the target dataset’s query and corpus distributions being nearly indistinguishable in the embedding space. These results lead to two practical conclusions. First, an attacker can mount a strong Black-Hole Attack without accurately guessing the real user query distribution. Second, the attack is most dangerous when the real queries deviate substantially from the stored corpus, a scenario that is common in real-world deployments where users ask diverse questions beyond the knowledge base.

5 THEORETICAL ANALYSIS OF BLACK-HOLE ATTACK

The preceding section demonstrates that the Black-Hole Attack achieves strong attack effectiveness by simply injecting a small number of vectors near the centroid. A natural question arises: *why does such a simple strategy work so well?* In this section, we provide a theoretical answer. We show that high-dimensional embedding spaces exhibit a structural property, which we term *centrality-driven hubness*: vectors near the centroid tend to be the nearest neighbor of most other vectors. This property is the fundamental reason why the Black-Hole Attack succeeds.

We begin with a motivating example.

Example 5.1. Consider a typical vector database that stores n document embeddings, each of dimension d , produced by an embedding model. The embeddings follow some distribution with covariance matrix Σ .

To make the setting in Example 5.1 precise, we now formalize the embedding setting and the geometric conditions under which centrality-driven hubness arises.

Definition 5.2.

- Following prior studies [17, 47], which show that embeddings satisfy a certain type of anisotropic distribution, we assume that queries are drawn from a distribution similar to that of the database corpus. Let $x_1, \dots, x_n \in \mathbb{R}^d$ be the embedding vectors. Following [49, 68, 73], we model them as i.i.d. samples from an anisotropic Gaussian $\mathcal{N}(\mu, \Sigma)$ with $\Sigma \succeq 0$. The sample centroid is $c = \frac{1}{n} \sum_{k=1}^n x_k$.
- From the covariance matrix Σ we extract two trace statistics: $m_1 = \text{tr}(\Sigma)$, the total variance across all dimensions, and $m_2 = \text{tr}(\Sigma^2)$, which captures how unevenly the variance is distributed. We also write $L = \|\Sigma\|_{\text{op}}$ for the largest eigenvalue.
- The *effective dimension* $d_{\text{eff}} := m_1^2/m_2 \in [1, d]$ measures how many dimensions carry substantial variance. The *effective rank* $r(\Sigma) := m_1/L \in [1, d]$ quantifies how dominant the top-variance direction is; smaller $r(\Sigma)$ indicates more severe anisotropy [29, 48].

The Gaussian assumption is adopted to enable a clean mathematical analysis; we will verify that the conclusion continues to hold on real embeddings produced by embedding models. We now formally characterize the conditions under which the centroid c is closer to a typical data point than any other vector in the database.

THEOREM 5.3. Fix a failure probability $\delta \in (0, 1)$. Define

$$t_1 = \log \frac{2}{\delta}, \quad t_2 = \log \frac{2(n-1)}{\delta}. \quad (3)$$

If the covariance statistics satisfy

$$2(m_1 - 2\sqrt{m_2 t_2}) > \left(1 - \frac{1}{n}\right)(m_1 + 2\sqrt{m_2 t_1} + 2L t_1), \quad (4)$$

then, with probability at least $1 - \delta$,

$$\min_{j \neq i} \|x_i - x_j\|_2 > \|x_i - c\|_2. \quad (5)$$

Intuitively, condition (4) reduces to

$$d_{\text{eff}} \gtrsim \log(n/\delta), \quad r(\Sigma) \gtrsim \log(1/\delta), \quad (6)$$

meaning that the stored vectors are insufficient to cover the embedding space of this dimensionality, and the anisotropy of the distribution remains moderate.

PROOF. We first center the distribution. Since Euclidean distances are translation-invariant, define $y_k = x_k - \mu$, so that $y_k \sim \mathcal{N}(0, \Sigma)$ and $x_i - c = y_i - \bar{y}$ where $\bar{y} = \frac{1}{n} \sum_{k=1}^n y_k$. We repeatedly use the following standard tail bound.

LEMMA 5.4 (GAUSSIAN QUADRATIC-FORM TAIL [60]). Let $z \sim \mathcal{N}(0, I_d)$ and $A \geq 0$. For $Q = z^\top A z$ and any $t \geq 0$,

$$\Pr(Q \geq \text{tr}(A) + 2\sqrt{\text{tr}(A^2)t} + 2\|A\|_{\text{op}}t) \leq e^{-t}, \quad (7)$$

$$\Pr(Q \leq \text{tr}(A) - 2\sqrt{\text{tr}(A^2)t}) \leq e^{-t}. \quad (8)$$

Step 1 (upper bound on the centroid distance). Note that $y_i - \bar{y} \sim \mathcal{N}(0, (1 - \frac{1}{n})\Sigma)$. Applying Lemma 5.4 with $A = (1 - \frac{1}{n})\Sigma$ and $t = t_1$,

$$\Pr(\|x_i - c\|_2^2 > (1 - \frac{1}{n})(m_1 + 2\sqrt{m_2 t_1} + 2L t_1)) \leq e^{-t_1}. \quad (9)$$

Let $E_{\text{cent}}^{(i)}$ denote the event that $\|x_i - c\|_2^2$ does *not* exceed this upper bound. Then $\Pr(E_{\text{cent}}^{(i)}) \geq 1 - \delta/2$.

Step 2 (lower bound on all pairwise distances). For any $j \neq i$, $y_i - y_j \sim \mathcal{N}(0, 2\Sigma)$. Applying Lemma 5.4 with $A = 2\Sigma$, $t = t_2$, and a union bound over all $n - 1$ choices of j gives

$$\Pr(\exists j \neq i : \|x_i - x_j\|_2^2 < 2(m_1 - 2\sqrt{m_2 t_2})) \leq (n - 1)e^{-t_2}. \quad (10)$$

Let $E_{\text{pair}}^{(i)}$ denote the event that *every* pairwise distance exceeds this lower bound. Then $\Pr(E_{\text{pair}}^{(i)}) \geq 1 - \delta/2$.

Step 3 (comparing the bounds). On $E_{\text{cent}}^{(i)} \cap E_{\text{pair}}^{(i)}$, which occurs with probability at least $1 - \delta$, condition (4) guarantees

$$2(m_1 - 2\sqrt{m_2 t_2}) > \left(1 - \frac{1}{n}\right)(m_1 + 2\sqrt{m_2 t_1} + 2L t_1), \quad (11)$$

and therefore $\min_{j \neq i} \|x_i - x_j\|_2 > \|x_i - c\|_2$. $\square$

Interpretation. To build intuition for condition (4), we derive a looser but more readable form. Ignoring the factor $(1 - \frac{1}{n}) \approx 1$ and noting that $t_2 \approx \log \frac{n}{\delta}$ dominates $t_1 \approx \log \frac{1}{\delta}$, the condition approximately requires

$$m_1 \gtrsim \sqrt{m_2 \log(n/\delta)} + L \log(1/\delta). \quad (12)$$

Dividing the first term by $\sqrt{m_2}$ and squaring gives $d_{\text{eff}} \gtrsim \log(n/\delta)$; dividing the second term by L gives $r(\Sigma) \gtrsim \log(1/\delta)$. We stress that these simplified conditions are *not* equivalent to the exact condition (4)—they serve only as an interpretive guide. Nonetheless, the simplified form offers clear intuition: the first condition means that the n vectors are insufficient to cover the embedding space, while the second ensures the anisotropy of the distribution remains moderate.

Cosine distance. We show centrality-driven hubness persists under cosine distance. Normalize all vectors to unit length: $\tilde{x}_k = x_k/\|x_k\|$ for $k = 1, \dots, n$, and work directly in this direction space. Let $\tilde{c} = \frac{1}{n} \sum_{j=1}^n \tilde{x}_j$ be the sample centroid of the normalized vectors, and define the hub as its projection back to the sphere, $h = \tilde{c}/\|\tilde{c}\|$. For a query point $\tilde{x}_i$, write $\eta_i := \tilde{x}_i \cdot h = \cos(\tilde{x}_i, h)$.

Claim. Model the normalized vectors as i.i.d. with covariance $\tilde{\Sigma} := \text{Cov}(x/\|x\|)$. Write $\tilde{m}_1 = \text{tr}(\tilde{\Sigma})$ and let $r(\tilde{\Sigma}) = \tilde{m}_1/\|\tilde{\Sigma}\|_{\text{op}}$ be the

effective rank in direction space. For any $\varepsilon, \delta \in (0, 1)$, provided the centrality gap $(1 - \|\tilde{c}\|)\eta_i$ stays bounded away from 0, the condition

$$r(\tilde{\Sigma}) \gtrsim \frac{2\tilde{m}_1}{(1 - \|\tilde{c}\|)^2 \eta_i^2} \log \frac{n}{\varepsilon \delta} \quad (13)$$

ensures that, with probability at least $1 - \delta$, the hub is the nearest neighbour of at least $(1 - \varepsilon)n$ points.

1. Mean-cosine identity. For a uniformly random index $J \in \{1, \dots, n\}$, an exact identity holds:

$$\mathbb{E}_J[\cos(\tilde{x}_i, \tilde{x}_J)] = \frac{1}{n} \sum_{j=1}^n \tilde{x}_i \cdot \tilde{x}_j = \tilde{x}_i \cdot \tilde{c} = \|\tilde{c}\| \cdot \eta_i. \quad (14)$$

Thus the gap between hub cosine and expected pairwise cosine is

$$\Delta_i := \eta_i - \mathbb{E}_J[\cos(\tilde{x}_i, \tilde{x}_J)] = (1 - \|\tilde{c}\|)\eta_i. \quad (15)$$

2. Beat-count expectation. Define $B_i := \#\{j \neq i : \cos(\tilde{x}_i, \tilde{x}_j) > \eta_i\}$, the number of stored points beating the hub. By linearity of expectation,

$$\mathbb{E}[B_i] = (n - 1) \cdot \Pr(\cos(\tilde{x}_i, \tilde{x}_J) > \eta_i). \quad (16)$$

3. Tail estimate via Gaussian approximation. Working in direction space, $\cos(\tilde{x}_i, \tilde{x}_J) = \tilde{x}_i \cdot \tilde{x}_J$ is a linear functional of $\tilde{x}_J$, so it has mean $\tilde{x}_i \cdot \tilde{c} = \|\tilde{c}\|\eta_i$ and variance $\sigma_i^2 := \tilde{x}_i^T \tilde{\Sigma} \tilde{x}_i$. Because the covariance $\tilde{\Sigma}$ is defined on the unit sphere, σ_i^2 is dimensionless and matches the scale of the cosine; since $\tilde{x}_i$ is a unit vector, $\sigma_i^2 \leq \|\tilde{\Sigma}\|_{\text{op}} = \tilde{m}_1/r(\tilde{\Sigma})$. Using the Gaussian tail estimate,

$$\Pr(\cos(\tilde{x}_i, \tilde{x}_J) > \eta_i) \lesssim \exp\left(-\frac{\Delta_i^2}{2\sigma_i^2}\right) \leq \exp\left(-\frac{r(\tilde{\Sigma})\Delta_i^2}{2\tilde{m}_1}\right). \quad (17)$$

Define $\gamma := \Delta_i^2/(2\tilde{m}_1)$. Thus,

$$\mathbb{E}[B_i] \lesssim (n - 1) \exp(-r(\tilde{\Sigma})\gamma). \quad (18)$$

4. From expectation to fraction. Let $S := \sum_{i=1}^n B_i$, so $\mathbb{E}[S] \lesssim n(n - 1) \exp(-r(\tilde{\Sigma})\gamma)$. Define $\text{Bad} := \{i : B_i \geq 1\}$. Since $S \geq |\text{Bad}|$, Markov's inequality gives, for any $\varepsilon \in (0, 1)$,

$$\Pr(|\text{Bad}| \geq \varepsilon n) \leq \frac{\mathbb{E}[S]}{\varepsilon n} \lesssim \frac{n \exp(-r(\tilde{\Sigma})\gamma)}{\varepsilon}. \quad (19)$$

For this to fall below a prescribed $\delta \in (0, 1)$, it suffices that $r(\tilde{\Sigma})\gamma \geq \log(n/(\varepsilon\delta))$. Substituting $\gamma = (1 - \|\tilde{c}\|)^2 \eta_i^2/(2\tilde{m}_1)$ yields

$$r(\tilde{\Sigma}) \gtrsim \frac{2\tilde{m}_1}{(1 - \|\tilde{c}\|)^2 \eta_i^2} \log \frac{n}{\varepsilon \delta}. \quad (20)$$

When this holds, $\Pr(|\text{Bad}| \geq \varepsilon n) \lesssim \delta$, i.e. the hub is the nearest neighbor of at least $(1 - \varepsilon)n$ points with probability at least $1 - \delta$.

Interpretation. The condition mirrors the Euclidean case and exposes the same two forces: a larger corpus ($n \uparrow$) raises $\log(n/\varepsilon\delta)$ and makes the hub harder to sustain, while stronger anisotropy ($r(\tilde{\Sigma}) \downarrow$) inflates the cosine variance and lets more points overtake the hub. The bound is informative only when the centrality gap $\Delta = (1 - \|\tilde{c}\|)\eta_i$ is bounded away from 0; for the embedding models we study this holds, and our experiments confirm the predicted scaling (Table 1).

Table 1: Centrality gap $\Delta = (1 - \|\tilde{c}\|)\eta$ on HotpotQA [72] with model BGE [70], Contriever [22], and GTE [75].

	BGE	Contriever	GTE
Δ	0.23	0.24	0.24

Discussion: From Single Gaussian to Gaussian Mixtures. The theoretical analysis above adopts a single anisotropic Gaussian model to obtain clean, interpretable conditions. Real-world textual embeddings, however, are notoriously multi-modal: documents spanning diverse topics form distinct semantic clusters in the embedding space. Under such a mixture distribution, the global centroid may not lie within any single semantic region, which weakens centrality-driven hubness. This multi-modality can be addressed by partitioning the embedding space into sufficiently many small clusters. Intuitively, fine-grained clustering isolates different semantic modes from one another, so that within each cluster the local distribution is approximately unimodal, bringing it closer to the single-Gaussian regime assumed by Theorem 5.3. Furthermore, since Theorem 5.3’s sufficient condition scales with $\log n$, a smaller per-cluster size n_j substantially weakens the dependence on database size, making the hubness condition easier to satisfy locally.

Validation. We validate the centrality-driven hubness phenomenon on real embeddings. We construct a corpus from HotpotQA [72] and embed the documents using three BGE models [70] with dimensions 384, 768, and 1024. Let $X = \{x_1, \dots, x_n\}$ denote the database vectors with centroid c , and $Q = \{q_1, \dots, q_m\}$ the query vectors. For each vector we compute its *hubness probability*, i.e., the probability that the centroid is closer than any other database vector: $\Pr_X[d(x_i, c) < \min_{j \neq i} d(x_i, x_j)]$ for corpus vectors and $\Pr_Q[d(q_i, c) < \min_j d(q_i, x_j)]$ for query vectors. Figure 5 reports this probability as a heatmap, where rows correspond to embedding dimensions and columns to database sizes ranging from 100 to 100K. We evaluate under two distance metrics (Euclidean and Cosine) and two centroid scopes: the *global* centroid of the entire database and *cluster-wise* centroids obtained by partitioning the database with k -means (each cluster containing about 100 vectors). The upper and lower halves of the figure correspond to corpus and query vectors, respectively. The results reveal pervasive centrality-driven hubness. Under the global centroid, hubness probability is high for small databases but decays as the database grows, since a denser space weakens the centroid’s dominance. Cluster-wise centroids dramatically amplify the effect: hubness probability exceeds 0.84 under Euclidean and 0.74 under Cosine across all database sizes, because each local centroid only needs to dominate a small neighborhood, which directly explains the effectiveness of the cluster-wise Black-Hole Attack. Crucially, the lower half of Figure 5 shows that **unseen queries** exhibit similar patterns to corpus vectors, confirming that centroids attract independently drawn queries as well.

To bridge the gap between theory and real data, we conduct an experiment on ideal Gaussian distributions. We sweep two modality settings: *single-mode*, where the entire corpus is drawn from one Gaussian (a single semantic cluster), and *multi-mode*, where the corpus is drawn from a mixture of Gaussians with 100 points per mode (emulating the many semantic clusters of real embeddings).

Figure 6 reports the results as four heatmaps. The experiment confirms our theory along two axes. First, dimension governs whether hubness arises at all: at $d=10$ the space is too low-dimensional and $\Pr$ stays near 0, but once $d \geq 100$ the injected centroid dominates, with $\Pr$ reaching 1.0 at $d=1000$. Second, in the single-mode setting $\Pr$ decays as the corpus grows. In the multi-mode setting, the per-mode size is fixed at 100 points regardless of n , which removes the corpus-size dependence: $\Pr$ no longer decays with n and instead stabilizes around 0.6–0.7, depending only on dimension. This directly mirrors our cluster-wise design, where partitioning the space into many small clusters keeps each local distribution effectively low- n and unimodal, so the hubness condition is easily satisfied within every cluster.




Figure 5: Hubness probability on real embeddings. Rows correspond to embedding dimensions (384, 768, 1024) and columns to database sizes (100–100K). Each cell value is $\Pr$, the fraction of vectors for which the centroid is the nearest neighbor.




Figure 6: Hubness probability on ideal Gaussian distributions. *Single* denotes a unimodal corpus drawn from one isotropic Gaussian; *Multi* denotes a multi-modal corpus drawn from a mixture of Gaussians.

6 EVALUATION ON ATTACKS¹

6.1 Evaluation Setup

6.1.1 Model Settings. We adopt the following embedding models to generate vectors: Contriever [22], a strong dense retriever trained

¹Due to space constraints, additional experiments are provided in the appendix of the extended version at <https://arxiv.org/abs/2604.05480>, including: (i) sensitivity to the poisoning rate, (ii) an analysis of cluster overlap, (iii) the full cluster-size imbalance breakdown, (iv) attack effectiveness on non-Transformer-generated embeddings such as SIFT [23], DEEP [1], and SPACEV [52], and (v) the effect of mixing vectors from different embedding models.

with contrastive learning; BGE-base-en-v1.5 [70], BAAI’s general-purpose embedding model; and GTE-base-en-v1.5 [75], an English text embedding model based on the Generalized Text Embedding architecture.

6.1.2 Dataset Preparation. Our evaluation involves 14 publicly available text datasets: Natural Questions (NQ) [30], HotpotQA [72], MSMARCO [6], Climate-FEVER [12], NFCorpus [4], DBPedia [31], FiQA-2018 [37], FEVER [57], SciFact [63], Touché-2020 [3], SCIDOCS [11], Quora [46], TREC-COVID [61], and ArguAna [62]. All 14 datasets are from the BEIR benchmark [56]. We also include three pre-embedded datasets: DBpedia-OpenAI-1M [54], Emoji-GTE [55], and Wiki-CS-E5 [27]. For each text dataset, we use its original corpus and queries. Due to the computational cost of brute-force search, for each evaluation run we randomly sample 3,000 queries from the query set to measure attack effectiveness.

6.1.3 Vector Database Settings. We employ FAISS [13] as our vector database backend and evaluate our approach using four distinct index types: (1) Flat (exact brute-force search), (2) graph-based HNSW [38], (3) clustering-based IVF-Flat [23], and (4) IVF-PQ [23]. Since the Flat index performs exhaustive search, it serves as our ground truth for nearest neighbors. Unless specified otherwise, **cosine similarity** is used as the distance metric.

6.1.4 Metrics. We evaluate using four metrics: Recall@K (R@K), Malicious Occupancy@K (MO@K), First Poisoned Rank (FPR), and Attack Success Rate (ASR).

Let $\mathcal{D} = \{x_i\}_{i=1}^N$ denote the benign database vectors and $\mathcal{P} = \{p_j\}_{j=1}^M$ the injected malicious vectors. The poisoned database is $\mathcal{D}' = \mathcal{D} \cup \mathcal{P}$. Let $\mathcal{Q}$ denote the set of evaluation queries. For each query $q \in \mathcal{Q}$, the vector database returns the ranked top- K results $R_K(q) = [r_1(q), \dots, r_K(q)]$.

Malicious Occupancy@K (MO@K). MO@K quantifies the fraction of the top- K results occupied by malicious items in the poisoned database. Define the number of malicious results in the top- K list as

$$m_K(q) = \sum_{t=1}^K \mathbf{1}[r_t(q) \in \mathcal{P}],$$

then

$$\text{MO@K} = \frac{1}{|\mathcal{Q}|} \sum_{q \in \mathcal{Q}} \frac{m_K(q)}{K}.$$

First Poisoned Rank (FPR). FPR measures how early the first malicious item appears within the top- K list:

$$\text{FPR}(q) = \min\{t \in \{1, \dots, K\} : r_t(q) \in \mathcal{P}\}, \quad \min \emptyset := 0.$$

We report the mean FPR over queries with $\text{FPR}(q) > 0$.

Attack Success Rate (ASR). ASR is the fraction of queries for which at least one malicious item appears in the top- K results:

$$\text{ASR} = \frac{1}{|\mathcal{Q}|} \sum_{q \in \mathcal{Q}} \mathbf{1}[\text{FPR}(q) > 0].$$

6.2 Overall Attack Effectiveness

In this section, we report the overall effectiveness of the proposed Black-Hole Attack. We evaluate three attack scenarios: the cluster-wise attack using full database access, the Surrogate Dataset Attack, and the attack on pre-embedded datasets.

Full Database Access Attack. We inject malicious vectors at a poisoning rate of 1%, cluster the benign database into more than 3000 clusters, and insert malicious vectors into the centroid of each cluster. Table 2 reports the results across three models, three datasets, and four index types. The attack is universally effective: MO@10 reaches 68%–81% on HotpotQA (ASR > 90%), 47%–59% on NQ, and 29%–58% on MSMARCO, with FPR values as low as 2.75. All three ANN indexes—HNSW, IVF-Flat, and IVF-PQ—exhibit attack effectiveness comparable to the exact Flat index, offering no inherent resilience. The vulnerability is model-agnostic: while Contriever is the most affected, the attack remains highly effective on BGE and GTE as well.

Table 2: Full Database Access Attack Effectiveness

Dataset	Model	Metric	Flat	HNSW	IVF-F	IVF-PQ
HotpotQA	Contriever	MO@10	76.97	80.86	79.98	75.36
		ASR	93.43	94.07	95.23	93.47
		FPR	3.10	2.75	2.83	2.98
	BGE	MO@10	72.71	77.66	73.79	67.58
		ASR	94.40	95.30	95.10	91.13
		FPR	3.58	3.11	3.48	3.84
	GTE	MO@10	68.10	72.13	70.01	65.09
		ASR	90.43	91.83	92.13	87.43
		FPR	3.99	3.60	3.81	4.07
	Contriever	MO@10	54.87	58.36	57.00	55.28
		ASR	72.10	75.13	74.87	73.27
		FPR	5.12	4.81	4.94	5.02
MSMARCO	BGE	MO@10	29.62	31.03	30.84	30.14
		ASR	46.23	47.53	48.00	47.67
		FPR	7.42	7.29	7.31	7.35
	GTE	MO@10	30.54	31.23	31.95	30.64
		ASR	48.10	48.83	50.07	48.53
		FPR	7.36	7.30	7.23	7.29
	Contriever	MO@10	52.91	55.32	54.44	51.67
		ASR	72.63	73.87	74.50	72.47
		FPR	5.31	5.08	5.18	5.32
	BGE	MO@10	57.90	59.41	59.02	57.96
		ASR	80.30	80.97	81.43	80.67
		FPR	4.90	4.76	4.79	4.84
NQ	GTE	MO@10	47.18	48.07	49.15	46.03
		ASR	69.90	70.37	72.97	69.33
		FPR	5.87	5.79	5.70	5.88

Higher MO@10 and ASR indicate a stronger attack, while lower FPR indicates a stronger attack.

Surrogate Dataset Attack. Figure 7 evaluates the Surrogate Dataset Attack, where the attacker constructs malicious vectors solely from publicly available query sets without any access to the target database’s internal embeddings. Query-Corpus similarity is the mean cosine similarity between query and corpus embeddings from the same source, and Query-Query similarity is the mean pairwise cosine similarity among surrogate query embeddings. For comparison, we also plot the Full Database Access Attack (orange dashed line), which uses the victim’s own corpus embeddings to

construct malicious vectors. The Surrogate Attack (blue solid line) achieves strong effectiveness on 10 of 14 datasets ($\text{MO@10} > 30\%$), with climate-fever reaching 98.3%, while the four failures—scidocs, quora, trec-covid, and arguana—all share high Query-Corpus similarity (0.594–0.949), indicating near-overlap between query and corpus distributions that mitigates the performance of the attack. Notably, on several datasets (climate-fever, hotpotqa, nfcopus, dbpedia, fiqa, fever), the Surrogate Attack even surpasses the Full Access baseline, demonstrating that precisely matching the victim’s corpus distribution is unnecessary: a surrogate query pool can exploit the same hubness geometry more effectively than the victim’s own embeddings. Attack effectiveness depends on this distributional gap rather than on accurately guessing user queries, dramatically lowering the practical barrier of the Black-Hole Attack.




Figure 7: Surrogate Dataset Attack vs. Full Database Access Attack across 14 BEIR datasets.

Poisoning Public Pre-Embedded Datasets. Table 3 reports the attack on three publicly available pre-embedded datasets, where the attacker directly downloads the pre-computed vectors, constructs malicious vectors via clustering, and injects them—without ever running an embedding model. Across all three datasets, the attack achieves MO@10 above 71%, ASR above 93%, and FPR below 4, covering different embedding models (OpenAI text-embedding-3-large, GTE-large-en-v1.5, E5-base) and diverse data modalities (entity descriptions, emoji-text pairs, Wikipedia articles). This demonstrates that the Black-Hole Attack readily extends to real-world pre-embedded datasets shared on open platforms such as Hugging Face.

Table 3: Results on precomputed-vector datasets.

Dataset	MO@10	ASR	FPR
DBpedia OpenAI 1M	71.10	93.74	3.76
Emoji GTE	75.52	98.66	3.43
Wiki CS E5	72.38	94.09	3.62

6.3 Robustness Analysis

6.3.1 Robustness to Partial Database Export. To analyze the impact of partial database access, we evaluate the attack under export ratios ranging from 0.1 to 1.0. Figure 8a demonstrates that our attack remains highly effective under limited exports, with MO@10 improving steadily as the ratio increases. On the MSMARCO dataset, expanding the export ratio from 0.1 to 1.0 improves MO@10 from

19.3% to 54.8%—a 184% increase—while on NQ, MO@10 nearly doubles from 26.4% to 52.1%. Notably, HotpotQA achieves an MO@10 of 61.9% at only 10% export, indicating that even minimal database access suffices for a practical attack. These results highlight that our attack does not depend on full corpus access and scales predictably with the exported portion.

6.3.2 Robustness to Dynamic Data Updates. Real-world vector databases are continually updated, so we evaluate how dynamic data updates affect the attack. We vary both data addition and deletion from 0% to 50% of the corpus size. To test the attack under *strong* distributional drift, the inserted vectors are drawn from *other* datasets rather than the target corpus itself. The results (Figure 8b) show that **deletion** slightly *improves* attack effectiveness, as removing benign vectors thins the neighborhood around the malicious centroids. **Insertion** shifts the data distribution and degrades the attack, but the effect is limited: even after 50% insertion, MSMARCO drops by only 1% in MO@10 , and HotpotQA still maintains 33% MO@10 . In addition to data churn, production systems also rebuild their indexes and occasionally re-embed their corpora. **Re-indexing** does not affect the attack: as Table 2 shows, attack effectiveness remains consistent across four index types, since the attack exploits embedding geometry rather than index structure. **Re-embedding** the corpus with a new model removes the injected malicious vectors from the database entirely, and is therefore an effective defense. However, its prohibitive computational cost makes it impractical for frequent deployment in production systems.

6.4 Sensitivity Analysis

6.4.1 Sensitivity to Number of Clusters. To analyze the effect of the cluster-wise design in Section 4.2, we vary the number of cluster centers used to construct local black-hole centers. Figure 8c reports the resulting MO@10 as a function of L under a fixed poisoning rate of 1%. We observe that MO@10 generally increases as L grows, indicating stronger malicious presence in the Top-10 results. It suggests that a moderate number of cluster centers is sufficient to cover most query regions for this dataset.

6.4.2 Sensitivity to Clustering Quality. We use FAISS-GPU to cluster the corpus, which proceeds in two stages: (1) it *trains* the centroids on a sample of the corpus, and (2) it *assigns* all corpus vectors to the learned centroids. The training sample size governs the quality of clustering. We fix the number of centroids L and the number of training iterations to 25, then vary the number of training samples per centroid s (so the total training set size is $L \cdot s$), thereby controlling clustering quality. Figure 8d reports MO@10 under increasing s . When only one point per centroid is sampled ($s=1$), MO@10 remains below 2% across all datasets. As s grows, MO@10 rises sharply and saturates under full sampling at 76.97%, 52.91%, and 54.87% on HotpotQA, NQ, and MSMARCO, respectively. Overall, higher clustering quality consistently leads to stronger attack performance, confirming that a well-clustered corpus representation is essential to the effectiveness of our method.

6.4.3 Sensitivity to Cluster Size. We analyze how cluster size affects the Black-Hole Attack by measuring MO@10 across clusters of different sizes. As shown in Figure 9, k -means clustering produces clusters whose sizes span a wide range, from around 100 up to




Figure 8: Robustness and Sensitivity Analysis of the Black-Hole Attack.

2K+. The results show that cluster size has different effects across datasets, and MO@10 remains consistently significant across all size groups. This confirms that the Black-Hole Attack is robust to the size imbalance of k -means clustering and does not rely on uniformly sized clusters.

6.4.4 Sensitivity to Cluster Density. We further analyze how the local density around a query affects the Black-Hole Attack. For each query, we define its *local density* as the mean similarity to its 10 nearest legitimate documents: a higher value means the query sits in a denser region. We sort queries by this density and measure MO@10 across density groups, from the sparsest to the densest. As shown in Figure 9, MO@10 decreases as density increases: when a query is close to corpus vectors, those nearby documents compete more strongly with the injected centroid and pull down the attack effectiveness. Nevertheless, even in the densest regions MO@10 remains significant across all datasets, showing that higher density only weakens the attack gradually rather than eliminating it.




Figure 9: MO@10 versus cluster size and cluster density. Query local density is defined as the mean similarity of a query to its 10 nearest legitimate documents; a higher value indicates a denser region.

6.4.5 Sensitivity to ANN Index Parameters. We systematically ablate the key parameters of three widely used ANN indexes—HNSW, IVF, and IVFPQ, comparing to a Flat baseline that uses exact search: 77.30%. The results are consolidated in Figure 10. Across all three index types, attack effectiveness fluctuates only moderately and remains close to the Flat baseline under all practical parameter settings. For HNSW, varying ef_search (16–256), M (4–64), and $ef_construction$ (40–800) yields MO@10 consistently within $\pm 10\%$ of Flat. For IVF, varying $nprobe$ (16–256) and $nlist$ (1024–8192) produces similarly stable results. For IVFPQ, setting the sub-quantizer count M to 48 slightly weakens the attack (still around 50%), but

this extreme compression already degrades retrieval accuracy; at standard settings ($M \geq 64$), the attack remains intact. Together, these results confirm that under different ANN index parameters, the Black-Hole Attack is always effective.




Figure 10: Sensitivity of the Black-Hole Attack to ANN index parameters (HNSW, IVF, IVFPQ) on HotpotQA with Contriever. Gray dashed lines indicate the Flat (exact search) baseline.

6.5 Cost and Analysis

6.5.1 Scalability to Larger Corpora. To verify that the Black-Hole Attack scales beyond the 1M–10M range, we conduct experiments on MS MARCO v2 [6, 15, 40] (138M documents, Contriever embeddings). We evaluate both the Full Database Access Attack and the Surrogate Dataset Attack. To obtain a sufficiently large surrogate query set, we leverage the query set of MS MARCO Web Search [8], whose queries are independently collected and have no overlap with the MS MARCO v2 queries, eliminating any risk of query leakage. We filter 4M English queries from this dataset and combine them with queries from other BEIR datasets to form the surrogate pool. Table 4 reports the results. Both attack variants achieve strong effectiveness on this 138M-document corpus. The Surrogate Dataset Attack attains MO@10 of 81.53% and ASR of 89.80%, and the Full Database Access Attack reaches MO@10 of 63.99% and ASR of 75.60%. Notably, both methods perform even better than on the 9M MS MARCO benchmark, confirming that the Black-Hole Attack remains effective at larger database sizes.

Table 4: Black-Hole Attack on MS MARCO v2.

Attack Scenario	MO@10	ASR	FPR
Full Database Access	63.99	75.60	4.29
Surrogate Dataset Attack	81.53	89.80	2.70

6.5.2 Analysis of Clustering Cost. We examine the cost of the clustering using MSMARCO-v2. We use the FAISS-GPU k -means implementation on an H200 GPU and measure the clustering time along two axes: (i) the corpus size, varied from 100K to 100M vectors at a fixed sampling budget (25% of the corpus); and (ii) the number of training samples per centroid, varied on a fixed 10M-vector corpus. The results are shown in Figure 11. Even at 100M vectors, clustering completes within three hours. For corpora at the billion-vector scale, the attacker can alternatively use the Surrogate Dataset Attack: it clusters a far smaller set of surrogate queries rather than the entire corpus, sidesteps the large-corpus clustering cost, and still achieves strong attack effectiveness.




Figure 11: Clustering cost on an H200 GPU. Left: runtime versus corpus size at a fixed sampling budget. Right: runtime versus the number of training samples per centroid on a fixed 10M-vector corpus.

6.6 Analysis on Downstream RAG Applications

6.6.1 Attack Effectiveness on QA Pipelines. We use the reproduced LongRAG [76] pipeline as the downstream application and poison its underlying vector database with the Black-Hole Attack. Figure 12 reports the end-to-end QA performance before and after poisoning. The results show that the attack substantially degrades answer quality on all three benchmarks. On HotpotQA, the F1 score drops from 51.8 to 37.0; on 2WikiMQA [20], it drops from 44.5 to 33.3; and on MuSiQue [58], it drops from 23.2 to 7.5. These results indicate that poisoning only the vector database is sufficient to severely impair the performance of a strong downstream RAG system. We further evaluate the effect of varying the rerank depth. Figure 12 shows that while more candidates improve the clean F1 as expected, the attack success rate remains above 80% across all datasets even under the largest rerank budget tested. This indicates that reranking alone, without dedicated defense mechanisms, does not meaningfully mitigate the Black-Hole Attack.

6.6.2 Impact of Filtering-Based Defenses on the Black-Hole Attack. RAG pipelines often employ clustering-based methods for **semantic filtering** and **duplicate detection** to defend against adversarial poisoning. SeCon-RAG [51] partitions the corpus into clusters and




Figure 12: Impact of the Black-Hole Attack on downstream RAG applications.

flags vectors too close to cluster centroids as malicious—a design that mirrors the Full Database Access Attack. Figure 13 reports the confusion matrices under both attack modes on HotpotQA. At 10,000 clusters, the filter flags 98.3% of malicious vectors under Full Database Access (TP), but only 13.5% of malicious vectors are flagged under Surrogate Dataset Attack. This gap arises because the surrogate query distribution diverges from the victim corpus distribution; centroids learned from surrogate queries do not precisely align with the victim’s embedding clusters, causing the filter to miss most injected vectors. TrustRAG [78] applies $K=2$ clustering to the top-10 retrieval results of each query and flags the denser cluster as malicious. Figure 14 shows that this method can flag a substantial fraction of malicious vectors (TP 73.7% for Full Access, 58.3% for Surrogate). However, the Black-Hole Attack already dominates the top-10 results—MO@10 is up to 76.97%. Even after TrustRAG removes most flagged malicious vectors, the remaining top-10 consists of the unflagged malicious ones, with only 1–3 clean documents surviving. Moreover, unlike SeCon-RAG, TrustRAG cannot directly remove malicious vectors from the corpus—it only filters at query time, leaving the poisoned vectors in the database.

Full Access Clusters=100		Full Access Clusters=1000		Full Access Clusters=10000	
TP	30.4%	TP	82.3%	TP	93.7%
FP	0.0%	FP	0.2%	FP	2.4%
FN	69.6%	FN	17.7%	FN	6.3%
TN	100.0%	TN	99.8%	TN	97.6%
Surrogate Clusters=100		Surrogate Clusters=1000		Surrogate Clusters=10000	
TP	0.2%	TP	1.2%	TP	13.5%
FP	0.0%	FP	0.2%	FP	2.4%
FN	99.8%	FN	98.8%	FN	86.5%
TN	100.0%	TN	99.8%	TN	97.6%

Figure 13: Detection accuracy of SeCon-RAG on malicious vectors. Columns: number of clusters used by the SeCon-RAG filter. TP: malicious vectors correctly flagged. FP: clean vectors falsely removed. FN: malicious vectors missed. TN: clean vectors correctly retained.

6.6.3 Robustness to Content-Level Inspection. Beyond clustering-based filters, RAG pipelines also deploy content-level inspection, including **source validation** and **lightweight content sanity checks**. Such defenses, however, are ill-suited to the Black-Hole

Full Access		Surrogate	
TP	73.7%	TP	58.3%
FP	4.2%	FP	0.5%
FN	26.3%	FN	41.7%
TN	95.8%	TN	99.5%

Figure 14: Detection accuracy of Trust-RAG.

Attack. Unlike adversarial-corpus-poisoning attacks that perturb a given text to derive an adversarial embedding, the Black-Hole Attack constructs malicious vectors *directly* in the embedding space and only later assigns text content to them. The attacker therefore retains complete freedom to design the accompanying text to evade text-level filters. For source validation, RAGParadox [10] has shown that malicious texts can be paired with fabricated but plausible source attributions, effectively bypassing source-level checks. For lightweight content sanity checks, CPA-RAG [34] shows that adversarially generated malicious texts with strong fluency and generalization can reliably evade perplexity and readability filters.

7 DEFENSES AGAINST BLACK-HOLE ATTACK

Hubness Mitigation. Since the Black-Hole Attack exploits centrality driven hubness, a natural defense is to reduce hubness in the embedding space. We evaluate three representative methods: ZN (Z-score normalization) [16, 59], CL2 (Centered L2 normalization) [67], and noHub [59]. Table 5 reports MO@10 and R@10 under each defense with Contriever embeddings and Flat search. ZN preserves high retrieval quality (R@10 85.83%–89.67%) but fails to suppress the attack (MO@10 65.36%–83.13%). noHub achieves strong attack suppression on HotpotQA and MSMARCO (MO@10 0.71%–13.91%), but degrades R@10 to near zero and leaks severely on NQ (MO@10 47.27%). CL2 provides the best overall tradeoff, reducing MO@10 to 1.66%–12.25% while retaining moderate R@10 (58.65%–76.70%). These results show that hubness mitigation can defend against the Black-Hole Attack, but every method forces a difficult choice between security and retrieval utility—none achieves both simultaneously.

Table 5: Defense effectiveness against the Black-Hole Attack: MO@10 and R@10 with Flat search.

Dataset	MO@10			R@10		
	ZN	CL2	noHub	ZN	CL2	noHub
HotpotQA	83.13	1.66	0.71	85.83	58.65	5.45
MSMARCO	65.36	12.25	13.91	89.67	76.70	11.38
NQ	68.66	10.00	47.27	87.48	66.78	10.44

Detection-Based Defense. We next study a more conservative defense that operates directly on the stored corpus vectors. The idea is to use a small probe set drawn from the database itself, measure how often each stored vector is retrieved as a nearest neighbor of these probes, and remove the vectors that are selected unusually often. Let the stored corpus vectors be $X = \{x_i\}_{i=1}^N \subset \mathbb{R}^d$. We first partition X into L clusters and, to cover different semantic regions, sample a small probe set from each cluster (0.1% of its vectors, at least one) to form the probe set P . For each probe we retrieve its

top- k nearest neighbors in X , and define the *hit count* h_i of each vector x_i as the number of probes that retrieve it. Since most vectors are never retrieved ($h_i = 0$), we compute the median m over only the strictly positive hit counts, and mark a vector as suspicious and remove it before retrieval if $h_i > 2m$. When using a brute-force (Flat) index, the detection procedure requires an additional k -NN pass for each probe vector, yielding $O(|P| \cdot N \cdot d)$ complexity, i.e., $O(N^2d)$ in the worst case. The defense reuses the database’s existing ANN index, so it incurs no additional indexing overhead. We evaluate this detection-based defense with a probe ratio of 0.1% per cluster. Table 6 reports the results under both the Full-Export Attack and the Surrogate Dataset Attack. Under the full-export attack, the defense is highly effective: it reduces MO@10 from 76.97%, 54.87%, and 52.91% to 7.09%, 1.50%, and 1.63% on HotpotQA, MSMARCO, and NQ respectively, while preserving high retrieval quality (R@10 94.51%–96.34%). Under the surrogate attack, however, the defense is substantially weakened—on HotpotQA and NQ, MO@10 drops only marginally (from 93.36% to 91.90% and from 81.45% to 68.13%, respectively). This gap arises from the same distributional mismatch discussed in Section 6.6.2: the probe vectors are drawn from the corpus distribution, whereas the surrogate attack places malicious vectors in query-offset regions, so they are rarely retrieved by the probes and evade the hit-count filter. Despite its limitations against surrogate attacks, the detection-based defense is computationally lightweight: with ANN acceleration, the probe search completes in seconds, making it practical to deploy as a routine safeguard against full-export-style threats.

Table 6: Detection-based defense under Full-Export Attack and Surrogate Dataset Attack, and probe search time across ANN index types.

Dataset	R@10	Full(MO@10)		Surrogate		Time (s)			
		Before	After	Before	After	Flat	HNSW	IVF	IVFPQ
HotpotQA	96.34	76.97	7.09	93.36	91.90	9k	4.4	7.5	1.3
MSMARCO	94.55	54.87	1.50	38.99	17.37	25k	6.1	16.6	2.3
NQ	94.51	52.91	1.63	81.45	68.13	2k	2.1	1.8	0.6

8 CONCLUSION

We present the *Black-Hole Attack*, a query-agnostic poisoning attack against vector databases grounded in *centrality-driven hubness*. By injecting a small number of malicious vectors at the global or cluster-wise centroid, an attacker can hijack Top- k results for most queries without any knowledge of their content. We define four attack paths along a capability spectrum and show strong effectiveness across all scenarios. Among the defenses we study, hubness mitigation provides only partial protection. The detection-based defense is highly effective against the full-export attack. However, it largely fails against the Surrogate Dataset Attack. A robust and adaptive defense remains an open problem.

ACKNOWLEDGMENTS

This work was supported in part by the National Natural Science Foundation of China (No. 92470204, U24B20147), and JST PRESTO (JPMJPR23P5), JST CREST (JPMJCR21M2), and JST NEXUS (JPMJNX25C4).

REFERENCES

- [1] Artem Babenko and Victor Lempitsky. 2016. Efficient Indexing of Billion-Scale Datasets of Deep Descriptors. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2055–2063. <https://doi.org/10.1109/CVPR.2016.226>
- [2] Gaurav Bagwe, Lan Zhang, Linke Guo, Miao Pan, Xiaolong Ma, and Xiaoyong (Brian) Yuan. 2025. Is Embedding-as-a-Service Safe? Meta-Prompt-Based Backdoor Attacks for User-Specific Trigger Migration. *Transactions on Artificial Intelligence* (2025). <https://api.semanticscholar.org/CorpusID:279133462>
- [3] Alexander Bondarenko, Maik Fröbe, Meriem Beloucif, Lukas Gienapp, Yamen Ajjour, Alexander Panchenko, Chris Biemann, Benno Stein, Henning Wachsmuth, Martin Potthast, and Matthias Hagen. 2020. Overview of Touché 2020: Argument Retrieval. In *Experimental IR Meets Multilinguality, Multimodality, and Interaction: 11th International Conference of the CLEF Association (CLEF 2020) (Lecture Notes in Computer Science)*, Vol. 12260. Springer, 384–395. https://doi.org/10.1007/978-3-030-58219-7_26
- [4] Vera Boteva, Demian Gholipour, Artem Sokolov, and Stefan Riezler. 2016. A Full-Text Learning to Rank Dataset for Medical Information Retrieval. In *Proceedings of the 38th European Conference on Information Retrieval (ECIR)*. Springer, 716–722. https://doi.org/10.1007/978-3-319-30671-1_58
- [5] Jon Bratseth. 2017. Open Sourcing Vespa, Yahoo’s Big Data Processing and Serving Engine. <https://blog.vespa.ai/open-sourcing-vespa-yahoos-big-data-processing/>. Accessed: 2025-10-12.
- [6] Daniel Fernando Campos, Tri Nguyen, Mir Rosenberg, Xia Song, Jianfeng Gao, Saurabh Tiwary, Rangan Majumder, Li Deng, and Bhaskar Mitra. 2016. MS MARCO: A Human Generated Machine Reading Comprehension Dataset. *ArXiv abs/1611.09268* (2016). <https://api.semanticscholar.org/CorpusID:2859517>
- [7] Cheng Chen, Chenzhe Jin, Yunan Zhang, Sasha Podolsky, Chun Wu, Szu-Po Wang, Eric Hanson, Zhou Sun, Robert Walzer, and Jianguo Wang. 2024. SingleStore-V: An Integrated Vector Database System in SingleStore. *Proceedings of the VLDB Endowment* 17, 12 (Aug. 2024), 3772–3785. <https://doi.org/10.14778/3685800.3685805>
- [8] Qi Chen, Xiubo Geng, Corby Rosset, Carolyn Buracton, Jingwen Lu, Tao Shen, Kun Zhou, Chenyan Xiong, Yeyun Gong, Paul Bennett, Nick Craswell, Xing Xie, Fan Yang, Bryan Tower, Nikhil Rao, Anlei Dong, Wenqi Jiang, Zheng Liu, Mingqin Li, Chuanjie Liu, Zengzhong Li, Rangan Majumder, Jennifer Neville, Andy Oakley, Knut Magne Risvik, Harsha Vardhan Simhadri, Manik Varma, Yujing Wang, Linjun Yang, Mao Yang, and Ce Zhang. 2024. MS MARCO Web Search: A Large-Scale Information-Rich Web Dataset with Millions of Real Click Labels. In *Companion Proceedings of the ACM Web Conference 2024 (WWW ’24 Companion)*, 292–301. <https://doi.org/10.1145/3589335.3648327>
- [9] Zhuo Chen, Yuyang Gong, Jiawei Liu, Miaokun Chen, Haotian Liu, Qikai Cheng, Fan Zhang, Wei Lu, and Xiaozhong Liu. 2025. FlippedRAG: Black-Box Opinion Manipulation Adversarial Attacks to Retrieval-Augmented Generation Models. *Proceedings of the 2025 ACM SIGSAC Conference on Computer and Communications Security* (2025). <https://api.semanticscholar.org/CorpusID:275337175>
- [10] Chanwoo Choi, Jinsoo Kim, Sukmin Cho, Soyeon Jeong, and Buru Chang. 2025. The RAG Paradox: A Black-Box Attack Exploiting Unintentional Vulnerabilities in Retrieval-Augmented Generation Systems. In *Findings of the Association for Computational Linguistics: EMNLP 2025*, 23723–23744.
- [11] Arman Cohan, Sergey Feldman, Iz Beltagy, Doug Downey, and Daniel Weld. 2020. SPECTER: Document-level Representation Learning using Citation-informed Transformers. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL)*. Association for Computational Linguistics, 2270–2282. <https://doi.org/10.18653/v1/2020.acl-main.207>
- [12] Thomas Diggelmann, Jordan Boyd-Graber, Jannis Bulian, Massimiliano Ciaramita, and Markus Leppold. 2020. CLIMATE-FEVER: A Dataset for Verification of Real-World Climate Claims. *arXiv:2012.00614 [cs.CL]* NeurIPS 2020 Workshop on Tackling Climate Change with Machine Learning.
- [13] Matthijs Douze, Alexandr Guzhva, Chengqi Deng, Jeff Johnson, Gergely Szilvasy, Pierre-Emmanuel Mazaré, Maria Lomeli, Lucas Hosseini, and Hervé Jégou. 2024. The Faiss library. (2024). *arXiv:2401.08281 [cs.LG]*
- [14] Wei Du, Peixuan Li, Bo Li, Haodong Zhao, and Gongshen Liu. 2023. UOR: Universal Backdoor Attacks on Pre-trained Language Models. In *Annual Meeting of the Association for Computational Linguistics*. <https://api.semanticscholar.org/CorpusID:258714833>
- [15] Kenneth Enevoldsen, Isaac Chung, Imene Kerboua, Márton Kardos, Ashwin Mathur, David Stap, Jay Gala, Wissam Siblini, Dominik Krzemiński, Genta Indra Winata, Saba Sturua, Saiteja Utpala, Mathieu Ciancone, Marion Schaeffer, Gabriel Sequeira, Diganta Misra, Shreeya Dhakal, Jonathan Rystrom, Roman Solomatin, Ömer Çağatan, Akash Kundu, Martin Bernstorff, Shitao Xiao, Akshita Sukhlecha, Bhavish Pawha, Rafał Poświata, Kranthi Kiran GV, Shawon Ashraf, Daniel Auras, Björn Plüster, Jan Philipp Harries, Loïc Magne, Isabelle Mohr, Mariya Hendriksen, Dawei Zhu, Hippolyte Gisserot-Boukhlef, Tom Aarsen, Jan Kostkan, Konrad Wojtasik, Taemin Lee, Marek Šuppa, Crystina Zhang, Roberta Rocca, Mohammed Hamdy, Andrianos Michail, John Yang, Manuel Fayse, Aleksei Vatolin, Nandan Thakur, Manan Dey, Dipam Vasani, Pranjal Chitale, Simone Tedeschi, Nguyen Tai, Artem Snegirev, Michael Günther, Mengzhou Xia, Weijia Shi, Xing Han Lu, Jordan Clive, Gayatri Krishnakumar, Anna Maksimova, Silvan Wehrli, Maria Tikhonova, Henil Panchal, Aleksandr Abramov, Malte Ostendorff, Zheng Liu, Simon Clematide, Lester James Miranda, Alena Fenogenova, Guangyu Song, Ruyi Bin Safi, Wen-Ding Li, Alessia Borghini, Federico Cassano, Hongjin Su, Jimmy Lin, Howard Yen, Lasse Hansen, Sara Hooker, Chenghao Xiao, Vaibhav Adlakha, Orion Weller, Siva Reddy, and Niklas Muennighoff. 2025. MMTEB: Massive Multilingual Text Embedding Benchmark. *arXiv preprint arXiv:2502.13595* (2025). <https://doi.org/10.48550/arXiv.2502.13595>
- [16] Nanyi Fei, Yizhao Gao, Zhiwu Lu, and Tao Xiang. 2021. Z-Score Normalization, Hubness, and Few-Shot Learning. 2021 *IEEE/CVF International Conference on Computer Vision (ICCV)* (2021), 142–151. <https://api.semanticscholar.org/CorpusID:247191876>
- [17] Alejandro Fuster Baggetto and Victor Fresno. 2022. Is anisotropy really the cause of BERT embeddings not being semantic?. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, Yoav Goldberg, Zornitsa Kozareva, and Yue Zhang (Eds.). Association for Computational Linguistics, Abu Dhabi, United Arab Emirates, 4271–4281. <https://doi.org/10.18653/v1/2022.findings-emnlp.314>
- [18] Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, Qianyu Guo, Meng Wang, and Haofen Wang. 2023. Retrieval-Augmented Generation for Large Language Models: A Survey. *ArXiv abs/2312.10997* (2023). <https://api.semanticscholar.org/CorpusID:266359151>
- [19] Runpeng Geng, Yanting Wang, Ying Chen, and Jinyuan Jia. 2025. UniC-RAG: Universal Knowledge Corruption Attacks to Retrieval-Augmented Generation. *arXiv preprint arXiv:2508.18652* (2025).
- [20] Xanh Ho, Anh-Khoa Duong Nguyen, Saku Sugawara, and Akiko Aizawa. 2020. Constructing A Multi-hop QA Dataset for Comprehensive Evaluation of Reasoning Steps. In *Proceedings of the 28th International Conference on Computational Linguistics*. International Committee on Computational Linguistics, Barcelona, Spain (Online), 6609–6625. <https://doi.org/10.18653/v1/2020.coling-main.580>
- [21] Guoyu Hu, Shaofeng Cai, Tien Tuan Anh Dinh, Zhongle Xie, Cong Yue, Gang Chen, and Beng Chin Ooi. 2025. HAKES: Scalable Vector Database for Embedding Search Service. *Proceedings of the VLDB Endowment* 18, 9 (May 2025), 3049–3062. <https://doi.org/10.14778/3746405.3746427>
- [22] Gautier Izacard, Mathilde Caron, Lucas Hosseini, Sebastian Riedel, Piotr Bojanowski, Armand Joulin, and Edouard Grave. 2021. Unsupervised Dense Information Retrieval with Contrastive Learning. *Trans. Mach. Learn. Res.* 2022 (2021). <https://api.semanticscholar.org/CorpusID:249097975>
- [23] Hervé Jégou, Matthijs Douze, and Cordelia Schmid. 2011. Product Quantization for Nearest Neighbor Search. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 33 (2011), 117–128. <https://api.semanticscholar.org/CorpusID:5850884>
- [24] Wenqi Jiang, Hang Hu, Torsten Hoefler, and Gustavo Alonso. 2025. Fast Graph Vector Search via Hardware Acceleration and Delayed-Synchronization Traversal. *Proceedings of the VLDB Endowment* 18, 11 (July 2025), 3797–3811. <https://doi.org/10.14778/3749646.3749655>
- [25] Wenqi Jiang, Marco Zeller, Roger Waleffe, Torsten Hoefler, and Gustavo Alonso. 2024. Chameleon: A Heterogeneous and Disaggregated Accelerator System for Retrieval-Augmented Language Models. *Proceedings of the VLDB Endowment* 18, 1 (Sept. 2024), 42–52. <https://doi.org/10.14778/3696435.3696439>
- [26] Yang Jiao, Xiaodong Wang, and Kai Yang. 2025. PR-Attack: Coordinated Prompt-RAG Attacks on Retrieval-Augmented Generation in Large Language Models via Bilevel Optimization. *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval* (2025). <https://api.semanticscholar.org/CorpusID:277667367>
- [27] Karmiq. 2024. wikipedia-embeddings-cs-e5-base. <https://huggingface.co/datasets/karmiq/wikipedia-embeddings-cs-e5-base>.
- [28] Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. Dense Passage Retrieval for Open-Domain Question Answering. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics, Online, 6769–6781. <https://doi.org/10.18653/v1/2020.emnlp-main.550>
- [29] Vladimir Koltchinskii and Karim Lounici. 2014. Asymptotics and Concentration Bounds for Bilinear Forms of Spectral Projectors of Sample Covariance. *arXiv: Statistics Theory* (2014). <https://api.semanticscholar.org/CorpusID:88513200>
- [30] Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, et al. 2019. Natural Questions: A Benchmark for Question Answering Research. In *TACL*.
- [31] Jens Lehmann, Robert Isele, Max Jakob, Anja Jentzsch, Dimitris Kontokostas, Pablo N. Mendes, Sebastian Hellmann, Mohamed Morsey, Patrick van Kleef, Sören Auer, and Christian Bizer. 2015. DBpedia—A Large-Scale, Multilingual Knowledge Base Extracted from Wikipedia. *Semantic Web* 6, 2 (2015), 167–195.
- [32] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. In *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin (Eds.), Vol. 33. Curran Associates, Inc., 9459–9474. https://proceedings.neurips.cc/paper_files/paper/2020/file/6b49320205f780e1bc26945df7481e5-Paper.pdf

- [33] Quentin Lhoest, Albert Villanova Del Moral, Yacine Jernite, Abhishek Thakur, Patrick Von Platen, Suraj Patil, Julien Chaumond, Mariama Drame, Julien Plu, Lewis Tunstall, Joe Davison, Mario Šaško, Gunjan Chhablani, Bhavitvya Malik, Simon Brandeis, Teven Le Scao, Victor Sanh, Canwen Xu, Nicolas Patry, Angelina McMillan-Major, Philipp Schmid, Sylvain Gugger, Clément Delangue, Théo Matussière, Lysandre Debut, Stas Bekman, Pierric Cistac, Thibault Goehringer, Victor Mustar, François Lagunas, Alexander Rush, and Thomas Wolf. 2021. Datasets: A Community Library for Natural Language Processing. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*. Association for Computational Linguistics, Online and Punta Cana, Dominican Republic, 175–184. <https://doi.org/10.18653/v1/2021.emnlp-demo.21>
- [34] Chunyang Li, Junwei Zhang, Anda Cheng, Zhuo Ma, Xinghua Li, and Jianfeng Ma. 2025. CPA-RAG: Covert Poisoning Attacks on Retrieval-Augmented Generation in Large Language Models. *arXiv preprint arXiv:2505.19864* (2025).
- [35] Guoliang Li, Xuanhe Zhou, and Xinyang Zhao. 2024. LLM for Data Management. *Proceedings of the VLDB Endowment* 17, 12 (Aug. 2024), 4213–4216. <https://doi.org/10.14778/3685800.3685838>
- [36] Dawei Liu, Bolong Zheng, Ziyang Yue, Fuhao Ruan, Xiaofang Zhou, and Christian S. Jensen. 2025. Wolverine: Highly Efficient Monotonic Search Path Repair for Graph-Based ANN Index Updates. *Proceedings of the VLDB Endowment* 18, 7 (March 2025), 2268–2280. <https://doi.org/10.14778/3734839.3734860>
- [37] Macedo Maia, Siegfried Handschuh, André Freitas, Brian Davis, Ross McDermott, Manel Zarrouk, and Alexandra Balahur. 2018. WWW’18 Open Challenge: Financial Opinion Mining and Question Answering. In *Companion Proceedings of the The Web Conference 2018 (WWW ’18)*. ACM.
- [38] Yu A. Malkov and D. A. Yashunin. 2020. Efficient and Robust Approximate Nearest Neighbor Search Using Hierarchical Navigable Small World Graphs. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 42, 4 (2020), 824–836. <https://doi.org/10.1109/TPAMI.2018.2889473>
- [39] John X. Morris, Volodymyr Kuleshov, Vitaly Shmatikov, and Alexander M. Rush. 2023. Text Embeddings Reveal (Almost) As Much As Text. In *Conference on Empirical Methods in Natural Language Processing*. <https://api.semanticscholar.org/CorpusID:263829206>
- [40] Niklas Muennighoff, Nouamane Tazi, Loïc Magne, and Nils Reimers. 2022. MTEB: Massive Text Embedding Benchmark. *arXiv preprint arXiv:2210.07316* (2022). <https://doi.org/10.48550/ARXIV.2210.07316>
- [41] James Jie Pan, Jianguo Wang, and Guoliang Li. 2024. Survey of vector database management systems. *The VLDB Journal* 33, 5 (2024), 1591–1615.
- [42] James Jie Pan, Jianguo Wang, and Guoliang Li. 2024. Vector Database Management Techniques and Systems. In *Companion of the 2024 International Conference on Management of Data (Santiago AA, Chile) (SIGMOD ’24)*. Association for Computing Machinery, New York, NY, USA, 597–604. <https://doi.org/10.1145/3626246.3654691>
- [43] Zhencan Peng, Miao Qiao, Wenchao Zhou, Feifei Li, and Dong Deng. 2025. Dynamic Range-Filtering Approximate Nearest Neighbor Search. *Proceedings of the VLDB Endowment* 18, 10 (June 2025), 3256–3268. <https://doi.org/10.14778/3748191.3748193>
- [44] Alina Petukhova, João P. Matos-Carvalho, and Nuno Fachada. 2024. Text Clustering with Large Language Model Embeddings. *CoRR abs/2403.15112* (2024). <https://arxiv.org/abs/2403.15112>
- [45] Pinecone Systems, Inc. 2024. Pinecone Documentation. <https://docs.pinecone.io/>. Accessed: 2026-09-01.
- [46] Quora. 2017. First Quora Dataset Release: Question Pairs. <https://quoradata.quora.com/First-Quora-Dataset-Release-Question-Pairs>
- [47] Sara Rajaei and Mohammad Taher Pilehvar. 2022. An Isotropy Analysis in the Multilingual BERT Embedding Space. In *Findings of the Association for Computational Linguistics: ACL 2022*. Smaranda Muresan, Preslav Nakov, and Aline Villavicencio (Eds.). Association for Computational Linguistics, Dublin, Ireland, 1309–1316. <https://doi.org/10.18653/v1/2022.findings-acl.103>
- [48] Stefano Recanatani, Serena Bradda, Vijay Balasubramanian, Nicholas A. Steinmetz, and Eric Shea-Brown. 2020. A scale-dependent measure of system dimensionality. *Patterns* 3 (2020). <https://api.semanticscholar.org/CorpusID:229549825>
- [49] Lingfeng Shen, Haiyun Jiang, Lema Liu, and Shuming Shi. 2023. Sen2Pro: A Probabilistic Perspective to Sentence Embedding from Pre-trained Language Model. In *Proceedings of the 8th Workshop on Representation Learning for NLP (Repl4NLP 2023)*. Association for Computational Linguistics, Toronto, Canada, 315–333. <https://doi.org/10.18653/v1/2023.repl4nlp-1.26>
- [50] Joobo Shim, Jaewon Oh, Hongchan Roh, Jaeyoung Do, and Sang-Won Lee. 2025. Turbocharging Vector Databases Using Modern SSDs. *Proceedings of the VLDB Endowment* 18, 11 (July 2025), 4710–4722. <https://doi.org/10.14778/3749646.3749724>
- [51] Xiaonan Si, Meilin Zhu, Simeng Qin, Lijia Yu, Lijun Zhang, Shuaitong Liu, Xinfeng Li, Ranjie Duan, Yang Liu, and Xiaojun Jia. 2026. Secon-rag: A two-stage semantic filtering and conflict-free framework for trustworthy rag. *Advances in Neural Information Processing Systems* 38 (2026), 70652–70681.
- [52] Harsha Vardhan Simhadri, George Williams, Martin Aumüller, Matthijs Douze, Artem Babenko, Dmitry Baranchuk, Qi Chen, Lucas Hosseini, Ravishankar Krishnaswamy, Gopal Srinivasa, Suhas Jayaram Subramanya, and Jingdong Wang. 2022. Results of the NeurIPS’21 Challenge on Billion-Scale Approximate Nearest Neighbor Search. <https://doi.org/10.48550/arXiv.2205.03763> arXiv:2205.03763.
- [53] Ji Sun, Guoliang Li, James Pan, Jiang Wang, Yongqing Xie, Ruicheng Liu, and Wen Nie. 2025. GaussDB-Vector: A Large-Scale Persistent Real-Time Vector Database for LLM Applications. *Proceedings of the VLDB Endowment* 18, 12 (2025), 4951–4963.
- [54] Supabase. 2024. dbpedia-openai-3-large-1M. <https://huggingface.co/datasets/Supabase/dbpedia-openai-3-large-1M>
- [55] Piotr Szemraj. 2024. local-emoji-search-gte. <https://huggingface.co/datasets/pszemraj/local-emoji-search-gte>
- [56] Nandan Thakur, Nils Reimers, Andreas Rücklé, Abhishek Srivastava, and Iryna Gurevych. 2021. BEIR: A Heterogeneous Benchmark for Zero-shot Evaluation of Information Retrieval Models. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*. <https://openreview.net/forum?id=wCu6T5xFje>
- [57] James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2018. FEVER: A Large-Scale Dataset for Fact Extraction and VERification. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*. Association for Computational Linguistics, 809–819. <https://doi.org/10.18653/v1/n18-1074>
- [58] Harsh Trivedi, Niranjan Balasubramanian, Tushar Khot, and Ashish Sabharwal. 2022. MuSiQue: Multihop Questions via Single-hop Question Composition. *Transactions of the Association for Computational Linguistics* 10 (2022), 539–554. https://doi.org/10.1162/tacl_a_00475
- [59] Daniel J. Trosten, Riddhi Chakraborty, Sigurd Løkke, Kristoffer Wickstrøm, Robert Jensen, and Michael C. Kampffmeyer. 2023. Hubs and Hyperspheres: Reducing Hubness and Improving Transductive Few-Shot Learning with Hyper-spherical Embeddings. *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2023), 7527–7536. <https://api.semanticscholar.org/CorpusID:257557374>
- [60] Roman Vershynin. 2018. *High-Dimensional Probability: An Introduction with Applications in Data Science*. Cambridge University Press.
- [61] Ellen Voorhees, Tasmeer Alam, Steven Bedrick, Dina Demner-Fushman, William R. Hersh, Kyle Lo, Kirk Roberts, Ian Soboroff, and Lucy Lu Wang. 2020. TREC-COVID: Constructing a Pandemic Information Retrieval Test Collection. *arXiv preprint arXiv:2005.04474* (2020).
- [62] Henning Wachsmuth, Shahbaz Syed, and Benno Stein. 2018. Retrieval of the Best Counterargument without Prior Topic Knowledge. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, 241–251. <https://doi.org/10.18653/v1/P18-1023>
- [63] David Wadden, Shanchuan Lin, Kyle Lo, Lucy Lu Wang, Madeleine van Zuylen, Arman Cohan, and Hannaneh Hajishirzi. 2020. Fact or Fiction: Verifying Scientific Claims. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics, 7534–7550. <https://doi.org/10.18653/v1/2020.emnlp-main.609>
- [64] Dongsheng Wang, Dandan Guo, He Zhao, Huangjie Zheng, Korawat Tanwisuth, Bo Chen, and Mingyuan Zhou. 2022. Representing Mixtures of Word Embeddings with Mixtures of Topic Embeddings. In *International Conference on Learning Representations (ICLR) 2022*. <https://openreview.net/forum?id=IYMuTbGzjFU>
- [65] H. Wang, S. Guo, J. He, H. Liu, T. Zhang, and T. Xiang. 2025. Model Supply Chain Poisoning: Backdoor Pre-trained Models via Embedding Indistinguishability. In *Proceedings of the ACM Web Conference 2025 (WWW ’25)*.
- [66] J. Wang, X. Yi, R. Guo, H. Jin, P. Xu, S. Li, X. Wang, X. Guo, C. Li, X. Xu, et al. 2021. Milvus: A purpose-built vector data management system. In *Proceedings of the 2021 International Conference on Management of Data (SIGMOD ’21)*. 2614–2627.
- [67] Yan Wang, Wei-Lun Chao, Kilian Q. Weinberger, and Laurens van der Maaten. 2019. SimpleShot: Revisiting Nearest-Neighbor Classification for Few-Shot Learning. *ArXiv abs/1911.04623* (2019). <https://api.semanticscholar.org/CorpusID:207863469>
- [68] Zihan Wang, Chengyu Dong, and Jingbo Shang. 2021. “Average” Approximates “First Principal Component”? An Empirical Analysis on Representations from Neural Language Models. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 5594–5603. <https://doi.org/10.18653/v1/2021.emnlp-main.453>
- [69] Weaviate B.V. [n.d.]. weaviate/weaviate: Weaviate – Open-source Vector Database. <https://github.com/weaviate/weaviate>. Accessed: 2025-10-12.
- [70] Shitao Xiao, Zheng Liu, Peitian Zhang, Niklas Muennighoff, Defu Lian, and Jian-Yun Nie. 2024. C-pack: Packed resources for general chinese embeddings. In *Proceedings of the 47th international ACM SIGIR conference on research and development in information retrieval*. 641–649.
- [71] Wenkai Yang, Lei Li, Zhiyuan Zhang, Xuancheng Ren, Xu Sun, and Bin He. 2021. Be Careful about Poisoned Word Embeddings: Exploring the Vulnerability of the Embedding Layers in NLP Models. *ArXiv abs/2103.15543* (2021). <https://api.semanticscholar.org/CorpusID:232404131>
- [72] Zhilin Yang, Peng Qi, Saizheng Zhang, et al. 2018. HotpotQA: A Dataset for Diverse, Explainable Multi-hop Question Answering. In *EMNLP*.

- [73] Shohei Yoda, Hayato Tsukagoshi, Ryohei Sasano, and Koichi Takeda. 2024. Sentence Representations via Gaussian Embedding. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*. Association for Computational Linguistics, 418–425. <https://aclanthology.org/2024.eacl-short.36/>
- [74] Collin Zhang, John X. Morris, and Vitaly Shmatikov. 2025. Universal Zero-shot Embedding Inversion. *ArXiv abs/2504.00147* (2025). <https://api.semanticscholar.org/CorpusID:277467864>
- [75] Xin Zhang, Yanzhao Zhang, Dingkun Long, Wen Xie, Ziqi Dai, Jialong Tang, Huan Lin, Baosong Yang, Pengjun Xie, Fei Huang, Meishan Zhang, Wenjie Li, and Min Zhang. 2024. mGTE: Generalized Long-Context Text Representation and Reranking Models for Multilingual Text Retrieval. In *Conference on Empirical Methods in Natural Language Processing*. <https://api.semanticscholar.org/CorpusID:271534420>
- [76] Qingfei Zhao, Ruobing Wang, Yukuo Cen, Daren Zha, Shicheng Tan, Yuxiao Dong, and Jie Tang. 2024. LongRAG: A Dual-Perspective Retrieval-Augmented Generation Paradigm for Long-Context Question Answering. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (Eds.). Association for Computational Linguistics, Miami, Florida, USA, 22600–22632. <https://doi.org/10.18653/v1/2024.emnlp-main.1259>
- [77] Xinyang Zhao, Xuanhe Zhou, and Guoliang Li. 2024. Chat2data: An interactive data analysis system with rag, vector databases and llms. *Proceedings of the VLDB Endowment* 17, 12 (2024), 4481–4484.
- [78] Huichi Zhou, Kin-Hei Lee, Zhonghao Zhan, Yue Chen, Zhenhao Li, Zhaoyang Wang, Hamed Haddadi, and Emine Yilmaz. 2025. TrustRAG: Enhancing Robustness and Trustworthiness in Retrieval-Augmented Generation. *arXiv preprint arXiv:2501.00879* (2025).
- [79] Wei Zou, Runpeng Geng, Binghui Wang, and Jinyuan Jia. 2024. PoisonedRAG: Knowledge Corruption Attacks to Retrieval-Augmented Generation of Large Language Models. In *USENIX Security Symposium*. <https://api.semanticscholar.org/CorpusID:271854736>