

# LLMs Meet Isolation Kernel: Lightweight, Learning-free Binary Embeddings for Fast Retrieval

Zhibo Zhang<sup>1 2</sup> Yang Xu<sup>1 2</sup> Kai Ming Ting<sup>1 2</sup> Cam-Tu Nguyen<sup>1 2\*</sup>

<sup>1</sup>National Key Laboratory for Novel Software Technology, Nanjing University

<sup>2</sup>School of Artificial Intelligence, Nanjing University

zhibozhang@smail.nju.edu.cn, xuyang@lamda.nju.edu.cn, {tingkm, ncamtu}@nju.edu.cn

## Abstract

Large language models (LLMs) have recently enabled remarkable progress in text representation. However, their embeddings are typically high-dimensional, leading to substantial storage and retrieval overhead. Although recent approaches such as Matryoshka Representation Learning (MRL) and Contrastive Sparse Representation (CSR) alleviate these issues to some extent, they still suffer from retrieval accuracy degradation. This paper proposes *Isolation Kernel Embedding* or IKE, a learning-free method that transforms an LLM embedding into a binary embedding using Isolation Kernel (IK). Lightweight and based on binary encoding, IKE offers a low memory footprint and fast bitwise computation, lowering retrieval latency. Experiments on multiple text retrieval datasets demonstrate that IKE offers up to  $16.7\times$  faster retrieval and  $16\times$  lower memory usage than the original LLM embeddings, while maintaining comparable accuracy. Theoretically, we show that IKE works because it satisfies four essential criteria for effective binary hashing that other methods do not possess. Compared to CSR, IKE consistently achieves better retrieval efficiency and effectiveness. IKE also works effectively with graph-based indexing, demonstrating its superiority in balancing accuracy and latency compared to alternative compression techniques in the approximate nearest neighbor (ANN) search setting.

## 1 Introduction

In the last decade, pretrained language models have significantly advanced text representation. Early approaches relied on encoder-based models such as BERT (Devlin et al., 2019) and RoBERTa (Liu et al., 2019). Recently, there has been a paradigm shift toward leveraging large language models (LLMs) for text representation with notable examples including LLM2Vec (BehnamGhader

et al., 2024), Llama2Vec (Liu et al., 2024b), and Qwen Embedding (Zhang et al., 2025b). These LLM embeddings now represent the state of the art in text representation (Enevoldsen et al., 2025).

However, LLM embeddings typically have a high number of dimensions. For example, T5-11B (Raffel et al., 2020) has 1024 dimensions, while the mainstream 7B decoder-only LLMs use 4096 dimensions. Such high-dimensional embeddings substantially increase memory storage and retrieval latency. To address this, Kusupati et al. (2022) introduced Matryoshka Representation Learning (MRL) that allows adaptive embedding length, balancing latency and accuracy for different applications. This approach is now supported in OpenAI and Google Gemini embedded APIs (OpenAI, 2024a; Lee et al., 2024) and has been extended to various applications (OpenAI, 2024b; Nussbaum et al., 2025; Yu et al., 2024). However, MRL requires full-parameter tuning, which results in a high training cost. Furthermore, experiments in (Wen et al., 2025) reveal a marked degradation in retrieval performance for MRL.

Recently, Wen et al. (2025) proposed Contrastive Sparse Representation (CSR) that sparsifies a pretrained LLM embedding into a high-dimensional but selectively activated feature space. Compared to MRL, CSR freezes the LLM backbone and trains only a lightweight adapter (an MLP layer), significantly reducing the training cost. However, unlike MRL, CSR requires retraining for different code lengths and suffers from substantial transformation overhead (see Section 4.3).

In a nutshell, MRL and CSR fall short of achieving the aim of fast retrieval time with low memory requirement, without substantially degrading the retrieval performance on a high-dimensional LLM embedded database.

While MRL and CSR provide a solution to the above retrieval problem, the methods are based on ad-hoc learning objectives, created without un-

\*Corresponding author

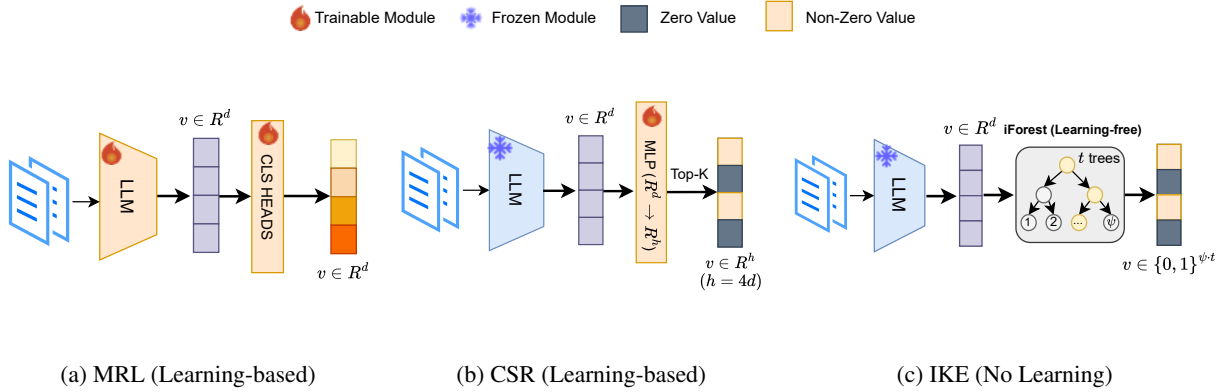


Figure 1: MRL learns adaptive embedding length by optimizing different downstream applications. CSR combines reconstruction loss and contrastive loss for learning a sparse representation. In contrast, IKE maps points into binary embeddings for efficient retrieval through random partitions without sophisticated learning.

Understanding the criteria of a feature space that are necessary in addressing the problem. A recent work has articulated three key criteria (Xu and Ting, 2025): full space coverage, entropy maximization, and bit independence. The hashing method called VDeH converts a high-dimensional space in  $\mathbb{R}^d$  to a binary embedded space, satisfying the three criteria without learning.

Though VDeH (Xu and Ting, 2025) has been shown to outperform many learning-to-hash (L2H) methods (Liu et al., 2024a; Weiss et al., 2008; Yu et al., 2014; Gong et al., 2013b,a; Xia et al., 2015; Jin et al., 2013; Zhu et al., 2014; Lan et al., 2017; Li et al., 2018) for non-LLM-related databases, our early examination has revealed that VDeH performs poorly in LLM-embedded databases (see Figure 2). This paper reports our effort in creating a generalized version of VDeH that is amenable to LLM-embedded databases.

Our contributions are summarized as follows:

- Proposing a novel feature transformation method named IKE that efficiently and effectively converts an LLM embedding to a binary embedding without learning. IKE is a generalization of VDeH, which produces an embedding that satisfies four criteria. The fourth criterion, which we have identified, is that the partitions used to generate the embedding shall have *high diversity*.
- Conducting a theoretical analysis that not only uncovers the *high diversity* criterion, but also the relative diversity of variants of two different implementations of IKE, i.e., Voronoi Diagrams and Isolation Forest.

- We conducted comprehensive experiments, where the main findings are: 1) *In the exhaustive search*, IKE generally achieves 2.5-16.7x speedup in search time and 8-16x storage reduction, while maintaining performance of MRR@10 between 98% and 101% of the original LLM space; 2) *In ANN (Approximate Nearest Neighbor Search)*, pairing IKE with an HNSW index achieves up to 5x higher throughput than operating in LLM space; 3) *Further comparison* with learning-based method (CSR) and other (no-learning) compression methods (hashing and quantization) demonstrate the advantages of IKE.

## 2 Preliminaries

**Isolation Kernel** The Isolation Kernel (Ting et al., 2018) (IK) is a data-dependent kernel that computes the similarity between two points based on the partitions in the data space. Specifically, given a data set  $D = \{\mathbf{x}_1, \dots, \mathbf{x}_n\} \subset \mathbb{R}^d$ , we define  $\mathcal{D} \subset D$  as a subset of  $\psi$  data points, each drawn from  $D$  with equal probability  $\frac{1}{n}$ . A *partition*  $H$  is derived so that each *isolating partition*  $\theta \in H$  isolates a point from the rest of  $\mathcal{D}$ .

Let  $\mathcal{H}_\psi(D)$  denote the set of all partitions that are admissible under the data set  $D$ . IK of any two points  $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$  wrt  $D$  is defined to be the expectation taken over the probability distribution  $P_{\mathcal{H}_\psi(D)}$  on all partitions  $H \in \mathcal{H}_\psi(D)$  that both  $\mathbf{x}$  and  $\mathbf{y}$  fall into the same isolating partition  $\theta \in H$ :

$$K_{\psi,D}(\mathbf{x}, \mathbf{y}) = \mathbb{E}_{H \sim P_{\mathcal{H}_\psi(D)}} [\mathbb{I}(\mathbf{x}, \mathbf{y} \in \theta \mid \theta \in H)],$$

where  $\mathbb{I}(\cdot)$  is an indicator function.

In practice, Isolation Kernel  $K_\psi$  is estimated from a finite number of  $t$  partitions  $H_i$ , each is




Figure 2: Effect of parameter  $m$  in  $\text{IKE}_{VD}$  ( $t=4096$ ) on  $\text{nDCG}@10$  across two LLM embedding datasets. The pentagram marker indicates the performance of  $\text{IKE}_{VD}$  when  $m = d$ , which is equivalent to the VDeH method.

derived from a data subset  $\mathcal{D}_i \subset \mathcal{D}$ :

$$K_{\psi, D}(\mathbf{x}, \mathbf{y}) \simeq \frac{1}{t} \sum_{i=1}^t \sum_{\theta_j \in H_i} \mathbb{I}(\mathbf{x} \in \theta_j) \mathbb{I}(\mathbf{y} \in \theta_j).$$

**Voronoi Diagram Encoded Hashing** In the preliminary study, Voronoi Diagrams (VD) are explored as an implementation of the Isolation Kernel for transforming LLM embeddings into binary codes. This approach, known as VDeH (Xu and Ting, 2025), has been shown to exhibit three desirable criteria: full space coverage, maximum entropy, and bit independence. However, VDeH suffers from notable performance degradation compared to the original LLM space (see Figure 2).

We hypothesize that the performance drop of VDeH is because the relationship between samples (points) in the LLM embedding space is more complex than what VDeH captures. In particular, using all dimensions for isolating partitions (i.e., hyperplanes in the VD) restricts the diversity of possible partitions. To verify this, we propose a generalization of VDeH, termed  $\text{IKE}_{VD}$ , where for each partition we randomly sample  $m$  ( $m < d$ ) dimensions for building VD. Note that when  $m = d$ ,  $\text{IKE}_{VD}$  becomes VDeH. The randomization of choosing isolation dimensions increases the diversity of isolating partitions, which becomes particularly beneficial when the number of partitions ( $t$ ) is large. As a result,  $\text{IKE}_{VD}$  achieves superior performance with small values of ( $m$ ) compared to VDeH, as shown in Figure 2.

### 3 Isolation Kernel-based Embeddings

Building on the previous observation, we investigate IKE using both Isolation Forest (iForest) and VD. The construction of iForest or VD requires

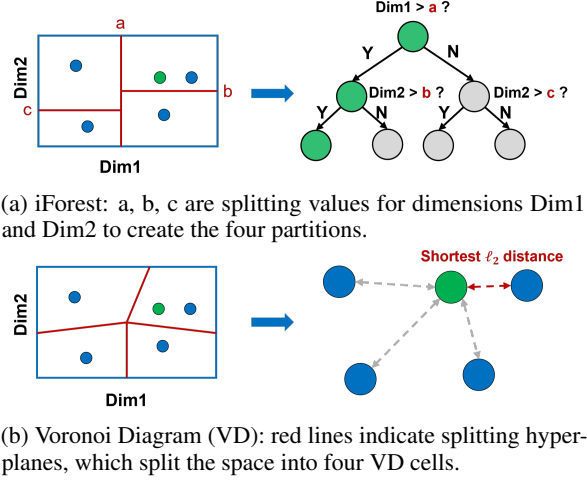


Figure 3: A illustration of partitioning in 2D space ( $\psi = 4$ ) by different IK implementations: iForest and VD. Blue points indicate the samples used to construct the partitions, and the green point represents an input point to be transformed.

no learning, which means neither sophisticated optimization nor even simple greedy search. The iForest or VD captures the data distribution by creating large partitions in sparse regions and small partitions in dense regions. This is created in a completely random process.

#### 3.1 IKE based Vector Database Construction

**Embedding** An LLM-based text embedding model is utilized to convert corpus documents and queries into dense vector embeddings. This embedding model is pretrained to effectively capture textual semantics.

**IKE with iForest** A random subset of the corpus embeddings is selected to construct an Isolation Forest (iForest) model, where each partition  $H$  is referred to as an *Isolation Tree* (iTree). Consequently, an Isolation Forest represents an ensemble of  $t$  independently constructed partitions.

An iTree  $H$  is a proper binary tree in which every internal node has two children. Each child node corresponds to an isolating partition  $\theta \in H$ . Here, a child node is obtained by *randomly selecting a feature  $q$  and the split value  $p$*  uniformly from the interval between the minimum and maximum values of feature  $q$  within  $\mathcal{D}$ . The construction details of the iTree are provided in Appendix D.1 and the partitioning is illustrated in Figure 3a.

Under the assumption that all points in  $\mathcal{D}$  are distinct, each point will eventually be isolated in its own leaf node. Therefore, an iTree consists

of exactly  $\psi$  leaf nodes and  $\psi - 1$  internal nodes, resulting in a total of  $2\psi - 1$  nodes. Hence, the memory complexity of a single iTree is  $O(\psi)$ . As previously stated,  $t$  iTrees constitute the Isolation Forest; thus, the overall space complexity of the model is  $O(t\psi)$ .

**IKE with Voronoi Diagram (IKE<sub>VD</sub>)** Although iForest is used in IKE by default, we can also exploit Voronoi Diagram for IKE as described in the previous section. Note that we refer IKE to indicate the iForest implementation, and IKE<sub>VD</sub> for the variant with VD implementation. A formal description of IKE<sub>VD</sub> is provided in Appendix D.2 and the partitioning of VD is illustrated in Figure 3b. An IKE<sub>VD</sub> model is composed of  $t$  independently constructed partitions, thus exhibiting an overall space complexity of  $O(mt\psi)$ .

**Vector Database** Depending on the retrieval strategy, the data points can either be stored directly, which corresponds to an exhaustive search during retrieval, or organized using an indexing structure to enable Approximate Nearest Neighbor (ANN) search. Each data point in IKE space is stored using an index-based representation  $\Phi^{\text{idx}}(\mathbf{x})$ , which comprises  $t$  elements. Here, each element  $\Phi_i^{\text{idx}}(\mathbf{x}) \in \{1, \dots, \psi\}$  indicates the index of the isolating partition to which data point  $\mathbf{x}$  is assigned within the  $i$ -th partition. As each element of  $\Phi^{\text{idx}}(\mathbf{x})$  can be stored using  $\lceil \log_2(\psi) \rceil$  bits, we only need  $t \lceil \log_2(\psi) \rceil$  bits per point. In contrast, we need a 32-bit floating-point number for each dimension in the original LLM embeddings. Therefore, when  $t = d$ , the transformed IKE yields a  $\frac{32}{\lceil \log_2(\psi) \rceil}$  times reduction in storage usage.

**Adaptability of Embedding Length** It is desirable to have embeddings that can be truncated for different accuracy–speed tradeoffs as in MRL (Kusupati et al., 2022). IKE achieves this through its compositional design: since the  $t$  iTrees are built independently, any subset can be retained to reduce the transformed space dimension. Unlike MRL, we achieve adaptability without LLM training.

### 3.2 Retrieval within the IKE space

**Retrieval.** Given a query  $q$  obtained from an LLM-based embedding model, we first apply the previously constructed iForest or VD model to map the query into the IKE space, yielding the transformed vector  $\Phi^{\text{idx}}(\mathbf{q})$ . Next, we measure the similarity between  $\Phi^{\text{idx}}(\mathbf{q})$  and each transformed

representation  $\Phi^{\text{idx}}(\mathbf{x})(\forall x \in D)$  for retrieval.

$$K_\psi(\mathbf{x}, \mathbf{y} \mid D) \simeq \frac{1}{t} \sum_{i=1}^t \mathbb{I}(\Phi_i^{\text{idx}}(\mathbf{x}) = \Phi_i^{\text{idx}}(\mathbf{y})).$$

**Mapping Time** The mapping time of IKE (with iForest by default) is  $O(t \lceil \log_2(\psi) \rceil)$  as the height limit of each iTree is set to  $\lceil \log_2(\psi) \rceil$ . On the other hand, the mapping process of IKE<sub>VD</sub> involves  $\ell_2$  distance calculations over  $m$  dimensional vectors, resulting in a time complexity of  $O(mt\psi)$ . As a result, IKE is more suitable for online retrieval compared to IKE<sub>VD</sub>.

**Efficient Similarity Computation** The similarity measure requires packing  $n_b = \lceil \log_2(\psi) \rceil$  bits into a single unit for comparison. When  $\psi = 2$ , we have  $n_b = 1$ , in which case the binary comparison can be applied. When  $n_b \geq 2$ , we employ the following approach for similarity computation:

Given two IKEs  $x$  and  $y$  to be compared, their partial components are first concatenated into longer bit strings (e.g., 64 bits), denoted as  $D_x$  and  $D_y$ . Let  $D = D_x \oplus D_y$ , where  $\oplus$  denotes the bitwise XOR operation, we compute:

$$M = D \mid (D \gg 1) \mid \dots \mid (D \gg (n_b - 1))$$

Here,  $\mid$  denotes the bitwise OR operation, and  $\gg$  denotes the logical right shift operation. The resulting bit string  $M$  has the following property: within every segment of  $n_b$  bits, if the rightmost bit is 0, it indicates that  $x$  and  $y$  are equal in the corresponding partition.

To isolate these zero bits and set all other bits to 1 for efficient counting, a mask is applied. The mask is constructed based on the value of  $n_b$  (for instance, when  $n_b = 2$ , the mask is 10101010...):

$$\begin{aligned} M &= M \mid \text{mask} \\ n_1 &= \text{popcnt}(M) \end{aligned}$$

After the OR operation with the mask, the number of remaining 0 bits in  $M$  corresponds to the number of elements where  $D_x$  and  $D_y$  match. The function  $\text{popcnt}(M)$  counts the number of 1-bits in  $M$ , denoted as  $n_1$ . The number of matching elements can then be derived by a simple subtraction. In practice, we only support  $n_b \in \{1, 2, 4, 8\}$  bits to keep the codes byte-aligned. An example is provided in Appendix D.3.

### 3.3 Criteria of Isolation Kernel Embedding

From a hashing perspective, each Voronoi cell or iTree leaf node can be viewed as a hash function that maps a data point to a binary value, where the value of “1” indicates that the hash function *covers* the point. Under some assumptions, IKE satisfies the three criteria demonstrated for VDeH (Xu and Ting, 2025): 1) **Full space coverage**: All hash functions cover entire space  $\mathbf{x} \in R^d$ , or there exists a hash function  $g(\mathbf{x})$  that maps  $\mathbf{x}$  to the value of 1; 2) **Entropy maximization**: For a given dataset, every hash function shall cover approximately the same number of points; 3) **Bit independence**: All bits in the hash code are mutually independent. The proof is given in Appendix C.

**Diversity Condition** By framing both VDeH and IKE as ensemble methods, we introduce **diverse partitioning** as a fourth, crucial criterion. We prove that IKE generates more robust hash representations than VDeH, owing to its ability to produce a more diverse set of base partitioners. The detailed proof is provided in Appendix C.4.

## 4 Experiments

We conduct comprehensive experiments to evaluate the performance of retrieval in the IKE space. For evaluation, we compute MRR@10 (Craswell, 2009) and nDCG@10 (Järvelin and Kekäläinen, 2002) as effectiveness metrics, and report search time and memory consumption as efficiency metrics. Here, MRR@10 computes the mean reciprocal rank of the first relevant document across all queries, whereas nDCG@10 measures the ranking quality by accounting for the position of relevant documents. All experiments were conducted on a Linux server equipped with two Intel Xeon Gold 6330 CPUs (2.00 GHz), providing a total of 56 physical cores (112 logical threads), and 503 GiB of RAM. Both index construction and search operations were performed in-memory using all available threads. Reported search times represent the average of 10 consecutive runs. We conducted the experiment in Section 4.2 using 10 different random seeds to demonstrate the stability of IKE. For other experiments, we used a single fixed seed.

### 4.1 Setup

**Datasets for Evaluation** We evaluate our method on four text retrieval datasets from the Massive Text Embedding Benchmark (MTEB) (Muenighoff et al., 2023) (HotpotQA (Yang et al.,

2018), FiQA2018 (Thakur et al., 2021), FEVERHN (Thorne et al., 2018), Touche2020.V3 (Thakur et al., 2024)), and two multi-level retrieval datasets: Istella22 (Dato et al., 2022) and TREC DL 23 (Craswell et al., 2023) (with a 1 million subset extracted from their original corpora). More dataset details are provided in Appendix E.1.

**Hyperparameter Settings** Following the original Isolation Forest framework (Liu et al., 2008), we set the maximum iTree height to  $l = \lceil \log_2(\psi) \rceil$ . Unless otherwise specified, the number of iTrees  $t$  equals the embedding dimension  $d$ . The parameter  $\psi$  is tuned per dataset via a single grid search under exhaustive search. The detailed tuning procedure is provided in Appendix E.2.

**LLMs** We employ two high-dimensional dense retrieval models: LLM2Vec (Mistral-7B backbone) (BehnamGhader et al., 2024) and Qwen3-embedding-8B (Zhang et al., 2025b), both producing 4096-dimensional embeddings.

### 4.2 Comparison with LLM embeddings


**Exhaustive Search** Table 1 presents the retrieval performance in the IKE space compared with the LLM embedding space under an exhaustive search configuration. The experiment in the LLM embedding space was implemented using the Faiss library (Douze et al., 2025). The experimental results demonstrate that retrieval in the IKE space achieves a 2.5–16.7 $\times$  speedup in search time and reduces memory usage by 8–16 $\times$  compared to the LLM embedding space, while maintaining comparable or superior retrieval performance. Specifically, MRR@10 reaches between 98% and 101% of those obtained in the LLM embedding space, while nDCG@10 varies from 96% to 100%.

**ANN Search** Unlike exhaustive search, ANN methods build an index over the data (Azizi et al., 2025; Peng et al., 2023; Wei et al., 2025; Wang et al., 2024; Lin, 2025) to restrict the search to a small subset of candidates, thereby reducing query time at the cost of a slight accuracy loss. We integrate IKE with two widely used indexing structures: the inverted file index (IVF) and the graph-based HNSW index (Malkov and Yashunin, 2018). Their retrieval performance is then compared against counterparts in the LLM embedding space with Faiss implementation (See Appendix F.1).

Figure 4 illustrates the Query per Second (QPS) versus MRR@10 trade-off curves of four ANN

Table 1: Results with two LLMs in the exhaustive search setting. The ‘‘Space’’ column reports the reduction in memory usage achieved by mapping the data into the IKE space, expressed as a multiplicative factor compared to the LLM embedding. In the ‘‘Other’’ column, we record auxiliary processing time: for IKE, this includes IKE model construction and feature mapping for all corpus embeddings; for LLM, this includes  $\ell_2$ -normalization of all corpus embeddings and Faiss’s exhaustive search index construction. The ‘‘Search’’ column measures the search execution time for all queries, with the speedup ratio of our method indicated in parentheses.

LLM2Vec							Qwen3-embedding						
Dataset	Space Type	Space	Time (s)		MRR@10	nDCG@10	Dataset	Space Type	Space	Time (s)		MRR@10	nDCG@10
			Other	Search						Other	Search		
HotpotQA	LLM		46.33	326.61	<b>83.83</b>	<b>74.09</b>	HotpotQA	LLM		41.70	332.66	<b>89.76</b>	<b>76.85</b>
	IKE ( $\psi=12$ )	$\downarrow 8\times$	<b>18.17</b>	<b>54.82</b> ( $\downarrow 6.0\times$ )	$83.69 \pm 0.13$	$73.24 \pm 0.09$		IKE ( $\psi=10$ )	$\downarrow 8\times$	<b>14.51</b>	<b>57.21</b> ( $\downarrow 5.8\times$ )	$88.92 \pm 0.13$	$75.41 \pm 0.10$
FiQA2018	LLM		0.29	0.47	<b>60.33</b>	<b>53.49</b>	FiQA2018	LLM		0.28	0.46	<b>72.05</b>	<b>64.66</b>
	IKE ( $\psi=6$ )	$\downarrow 8\times$	<b>0.20</b>	<b>0.063</b> ( $\downarrow 7.5\times$ )	$59.97 \pm 0.32$	$52.59 \pm 0.31$		IKE ( $\psi=15$ )	$\downarrow 8\times$	<b>0.24</b>	<b>0.069</b> ( $\downarrow 6.7\times$ )	$70.95 \pm 0.37$	$63.17 \pm 0.31$
FEVER-HN	LLM		0.84	1.10	<b>91.52</b>	<b>90.45</b>	FEVER-HN	LLM		0.88	1.22	<b>93.91</b>	<b>92.50</b>
	IKE ( $\psi=15$ )	$\downarrow 8\times$	<b>0.59</b>	<b>0.26</b> ( $\downarrow 4.2\times$ )	$91.36 \pm 0.23$	$90.27 \pm 0.16$		IKE ( $\psi=13$ )	$\downarrow 8\times$	<b>0.39</b>	<b>0.29</b> ( $\downarrow 4.2\times$ )	$93.59 \pm 0.22$	$92.09 \pm 0.19$
Istella22	LLM		5.05	1.56	68.11	<b>63.88</b>	Istella22	LLM		5.23	1.62	<b>65.25</b>	59.44
	IKE ( $\psi=16$ )	$\downarrow 8\times$	<b>3.26</b>	<b>0.62</b> ( $\downarrow 2.5\times$ )	$68.92 \pm 0.88$	$63.86 \pm 0.50$		IKE ( $\psi=11$ )	$\downarrow 8\times$	<b>1.64</b>	<b>0.61</b> ( $\downarrow 2.7\times$ )	$65.11 \pm 0.70$	$59.63 \pm 0.55$
TREC DL 23	LLM		5.11	0.90	97.56	<b>73.10</b>	TREC DL 23	LLM		5.08	0.84	95.73	<b>73.79</b>
	IKE ( $\psi=3$ )	$\downarrow 16\times$	<b>2.23</b>	<b>0.16</b> ( $\downarrow 5.6\times$ )	$98.11 \pm 1.02$	$71.78 \pm 0.62$		IKE ( $\psi=6$ )	$\downarrow 8\times$	<b>2.91</b>	<b>0.32</b> ( $\downarrow 2.6\times$ )	$96.93 \pm 0.42$	$72.78 \pm 0.28$
Touche2020.V3	LLM		1.57	1.44	<b>82.65</b>	<b>51.61</b>	Touche2020.V3	LLM		1.55	1.59	<b>96.94</b>	<b>76.41</b>
	IKE ( $\psi=7$ )	$\downarrow 8\times$	<b>1.06</b>	<b>0.095</b> ( $\downarrow 15.2\times$ )	$82.37 \pm 3.33$	$51.02 \pm 1.18$		IKE ( $\psi=16$ )	$\downarrow 8\times$	<b>0.47</b>	<b>0.095</b> ( $\downarrow 16.7\times$ )	$95.83 \pm 1.41$	$73.35 \pm 1.13$




(a) HotpotQA (LLM2Vec) (b) HotpotQA (Qwen3)

Figure 4: QPS vs. MRR@10 for Four ANN Methods on the HotpotQA dataset. The curves are obtained by adjusting the ANN search parameters. The detailed configuration is described in Appendix F.2.

methods on the HotpotQA dataset, operated in either the LLM embedding or the IKE spaces. Overall, IKE consistently outperforms that of LLM across different indexing methods. When using the HNSW index, retrieval with IKE achieves approximately 4–5 times higher throughput than with LLM embeddings. Specifically, on the HotpotQA (LLM2Vec) dataset, at a comparable MRR@10 of around 61%, the QPS increases from  $1.6 \times 10^4$  (HNSW+LLM) to  $6.4 \times 10^4$  (HNSW+IKE). When using the IVF index, performing retrieval in the IKE space yields a QPS improvement of approximately 6–8 times over that in the LLM-embedded space. Additional results on more datasets are provided in Appendix F.3.

### 4.3 Comparison with learning-based method

In this section, we compare IKE with compression methods that require learning, such as MRL (Kusupati et al., 2022) and CSR (Wen et al., 2025). However, as MRL requires full model retraining



(a) Accuracy (b) Search Time

Figure 5: Comparison of Retrieval Accuracy and Search Time between CSR and IKE in the exhaustive search setting on the HotpotQA (LLM2Vec) dataset.

(see Figure 1), we do not have enough computing resources for such an experiment. In addition, MRL is shown to be inferior to CSR, and thus, we select CSR as a representative baseline for comparison.

**CSR Parameter Settings** CSR uses a hyperparameter  $k$  to control the sparsity level of the vector  $z$ , where each vector contains exactly  $k$  non-zero entries. We compare CSR with IKE under fixed code length conditions, where the code length of a sparse vector  $z$  is defined as  $4k$  bytes. The CSR model is trained using the hyperparameters provided in the original paper for text retrieval tasks. The search time includes both the mapping time for all queries and the similarity computation time, rather than solely the similarity calculation time as reported in the original CSR paper.

**Comparative Results** Figure 5 presents a comparative evaluation in the exhaustive search setting on the HotpotQA (LLM2Vec) dataset between CSR and IKE in terms of search time on a CPU and retrieval accuracy, measured by nDCG@10.




Figure 6: **Performance Comparison of Compression Methods on ANN Retrieval: QPS vs. nDCG@10 on the HotpotQA (LLM2Vec) dataset.** The curves are obtained by adjusting the search parameter of the index structures.

Additional results on more datasets are provided in Appendix G. The results show that IKE significantly outperforms CSR in terms of search efficiency across all datasets, while achieving comparable or even superior retrieval accuracy. Specifically, IKE reduces the search time of CSR by one order of magnitude. On the HotpotQA (LLM2Vec) dataset, IKE achieves a search time of 4–12 seconds, whereas CSR requires 104–237 seconds. This demonstrates that IKE is a superior option to CSR in practical retrieval scenarios.

**Discussion** We analyze the reasons behind the significantly higher retrieval latency of CSR compared to the IKE method. The sparse vector mapping of CSR employs an MLP to transform the original  $d$ -dimensional embedding into a  $4d$ -dimensional space, after which only the top- $k$  dimensions are retained while the remaining ones are set to zero. Specifically, the mapping process of CSR requires  $O(d^2 + d \log_2(k))$  operations, whereas the mapping in IKE demands only  $O(t \lceil \log_2(\psi) \rceil)$ , where  $\lceil \log_2(\psi) \rceil$  denotes the maximum iTree height. Specifically, considering a scenario where  $d = 4096$  dimensional embedding is compressed into 128 bytes, IKE utilizes hyperparameters  $t = 1024$  and  $\psi = 2$  while CSR requires  $k = 32$ . The mapping complexity of CSR is dominated by the quadratic term  $O(d^2)$ , whereas that of IKE scales linearly with  $t$ . Given that  $d^2 \gg t$  in high-dimensional spaces, our method achieves a significant reduction in mapping overhead. Additionally, for similarity computation, our method leverages highly optimized bitwise operations, which are faster than the sparse matrix multiplication relied upon by CSR.

#### 4.4 Comparison with compression methods

In this section, we compare the IKE method with the learning-free vector compression tech-

niques, under **same code-length** scenario in the ANN search setting. Specifically, we study different code length constraints of  $\{1024, 512, 256, 128\}$  bytes. For a fixed code length, we plot QPS against the retrieval effectiveness measure (nDCG@10) for varying ANN search parameters. We compare our methods, IVF / HNSW (IKE), with several baselines: IVF / HNSW (rpLSH) (Charikar, 2002), IVF (PQfs) (André et al., 2015), and ScaNN (Guo et al., 2020). Detailed descriptions of these methods are provided in Appendix H.1. Note that the embedding length in IKE is determined by parameters  $t$  and  $\psi$ . Under a fixed code length, we find that smaller  $\psi$  and larger  $t$  yield better retrieval performance. Thus, we set  $\psi=2$  in this experiment.

Figure 6 illustrates QPS vs NDCG@10 curves for all methods on the HotpotQA (LLM2Vec) dataset, with the code lengths used indicated above each plot. Points closer to the upper right corner represent better overall performance. Our HNSW (IKE) method achieves the best balance between efficiency and effectiveness across all datasets. In terms of search efficiency, it achieves from  $2.9 \times$  to  $10 \times$  higher throughput than other learning-free compression methods in an ANN setting, while maintaining comparable retrieval accuracy. For instance, on the HotpotQA (LLM2Vec) dataset with a fixed code length of 512 bytes, when achieving approximately nDCG@10 of 60%, the IVF (PQfs) method attains around  $2.3 \times 10^4$  QPS, whereas the HNSW (IKE) method achieves around  $6.7 \times 10^4$  QPS, representing a 2.9 times improvement in efficiency. Similarly, when the code length is extended to 1024 bytes to reach 70% nDCG@10, HNSW (IKE) attains  $1.2 \times 10^4$  QPS, representing a  $10 \times$  efficiency improvement over IVF (rpLSH), which yields only  $1.1 \times 10^3$  QPS. Additional results are provided in Appendix H.2.




Figure 7: **Effect of parameter  $t$  (the number of iTrees) on Search Time and Retrieval Accuracy.** The gray dashed lines denote the corresponding retrieval performance in the LLM embedding space.

## 4.5 Extra analyses

### 4.5.1 Impact of hyperparameter $t$

The hyperparameter  $t$  denotes the number of iTrees in the iForest. Increasing  $t$  produces more fine-grained representations, which can enhance retrieval accuracy; conversely, reducing  $t$  trades off accuracy for lower query latency and reduced memory usage. To assess this trade-off, we evaluate the impact of varying  $t$  on the HotpotQA (LLM2Vec) and Istella22 (LLM2Vec) datasets under both exhaustive search and HNSW configurations. For the HNSW setup, we use a search parameter  $efSearch = 50$ . We draw  $t$  from 0.5K to 5K (where 1K = 1024), and at each value, we record the search time, retrieval accuracy (nDCG@10), and storage usage relative to LLM embeddings.

The results in Figure 7 show that search time increases approximately linearly with the number of iTrees  $t$ , whereas retrieval accuracy follows a logarithmic growth pattern. Specifically, while each increment in  $t$  below 3K yields noticeable gains in retrieval accuracy, the marginal improvement diminishes beyond this threshold. For instance, as illustrated in Figure 7b, reducing  $t$  from 4K to 3K decreases query latency from 0.42s to 0.34s, corresponding to a 19.05% improvement in search efficiency and a 25% reduction in memory footprint, at the expense of only a 0.36% drop in nDCG@10 (from 63.79% to 63.43%). Thus, tuning  $t$  enables substantial gains in search efficiency and memory savings with a minimal impact on retrieval accuracy. *In other words, IKE exhibits a desirable property similar to that of MRL, enabling flexible adaptation of code length to balance effectiveness and efficiency, while requiring no costly training.*

### 4.5.2 Impact of hyperparameter $\psi$

The hyperparameter  $\psi$  denotes the number of isolating partitions within an iTree. To investigate its influence, we conduct a sensitivity analysis on the validation set of the Istella22 (LLM2Vec) dataset using exhaustive search. Specifically, we fix the number of iTrees at  $t = 4096$  and vary  $\psi$  across the set  $\{2, 4, 8, 16\}$ . We report the resulting computational overhead and retrieval precision, as summarized in Table 2.

The experimental results indicate that, unlike the monotonic accuracy gains observed with parameter  $t$  in Figure 7, the impact of  $\psi$  on retrieval precision is relatively marginal and lacks a consistent monotonic trend. Notably, however, even with small values of  $\psi$ , IKE maintains remarkably high retrieval accuracy while reducing both mapping overhead and query latency. These findings provide critical engineering guidance: in latency-critical scenarios, selecting a small  $\psi$  enables practitioners to maximize computational efficiency with negligible compromises in retrieval performance.

Table 2: Sensitivity analysis for parameter  $\psi$

$t$	$\psi$	Time (s)		MRR@10	nDCG@10
		Other	Search		
4096	2	1.18	0.16	91.06	75.71
4096	4	1.69	0.25	90.99	75.83
4096	8	2.49	0.35	90.94	76.02
4096	16	3.70	0.36	91.43	76.35

### 4.5.3 Variations of IK implementations

In this section, we evaluate the performance of VDeH,  $\text{IKE}_{VD}$  ( $m \leq d$ ) and IKE. Here, we report retrieval performance under three configurations of  $m$ :  $m = 1$ ,  $m = d$  (equivalent to VDeH), and an

Table 3: Results of different implementations of the IKE in the exhaustive search setting.

Method	MRR@10 (%)	nDCG@10 (%)
FiQA2018 (LLM2Vec)		
IKE ( $\psi=6$ )	60.13	52.58
IKE <sub>VD</sub> ( $\psi=3, m=7$ )	60.23	52.89
IKE <sub>VD</sub> ( $\psi=3, m=1$ )	59.76	52.61
VDeH ( $\psi=3$ )	32.07	25.84
FiQA2018 (Qwen3)		
IKE ( $\psi=15$ )	71.14	63.25
IKE <sub>VD</sub> ( $\psi=4, m=5$ )	71.99	64.03
IKE <sub>VD</sub> ( $\psi=4, m=1$ )	69.87	62.16
VDeH ( $\psi=4$ )	58.91	49.78

optimally tuned  $m$ , obtained through hyperparameter search. The value of  $t$  is fixed at 4096, while  $\psi$  is tuned for best retrieval performance.

The results, presented in Table 3, reveal a substantial performance gap between IKE and IKE<sub>VD</sub> with  $m = d$ . For example, in terms of nDCG@10, IKE<sub>VD</sub> ( $m = d$ ) exhibits performance decreases of 50.86% and 21.30% on the two datasets, respectively. However, when  $m$  is reduced to smaller values such as  $m = 1$  or the optimally tuned  $m$  (which also tends to be small), the performance gap almost disappears. IKE<sub>VD</sub> ( $m = 1$ ) is slightly worse than IKE, which is aligned with the analysis in Appendix C.4, yet IKE<sub>VD</sub> with optimally tuned  $m$  can surpass IKE in MMR@10 and nDCG@10. Nevertheless, we recommend IKE over the optimally tuned IKE<sub>VD</sub> for two primary reasons. First, IKE<sub>VD</sub> requires tuning more hyperparameters (i.e., the extra parameter  $m$  beside  $t, \psi$ ). Second, the mapping process of IKE<sub>VD</sub> is computationally more expensive compared to IKE.

## 5 Conclusion

This paper presents IKE, a learning-free method based on the Isolation Kernel that transforms LLM embeddings into binary representations (in IKE space). IKE ensembles diverse partitions to reduce retrieval loss as the ensemble grows. It is lightweight, binary-encoded, and adaptable in code length, offering low memory use and fast mapping. Theoretical analysis highlights the importance of *high diversity* in IKE and its variant IKE<sub>VD</sub>.

Experiments on multiple text retrieval datasets show that IKE cuts memory and search time by up to 16.7 $\times$  while maintaining or even surpassing LLM-level accuracy. Compared to the learning-based CSR, IKE achieves order-of-magnitude

speedups without retraining. When combined with HNSW, it outperforms other learning-free compression methods (rpLSH, ScaNN, PQfs) in the ANN setting, delivering from 2.9 $\times$  to 10 $\times$  higher throughput with comparable accuracy.

## Limitations

While our method demonstrates strong performance on the text retrieval task, its effectiveness on cross-modal retrieval tasks (e.g., Image-to-Text or Text-to-Image) remains limited. This may be attributed to the modality gap (Liang et al., 2022) between textual and visual representations, which poses a challenge for cross-modal alignment. This aspect has not been investigated in this paper and is left for future research.

## Acknowledgements

Kai Ming Ting is supported by the National Natural Science Foundation of China (Grant No. 92470116 & W2531050). This project is also supported by the National Natural Science Foundation of China (Grant No. W2532049) and the Fundamental and Interdisciplinary Disciplines Breakthrough Plan of the Ministry of Education of China (No. JYB2025XDXM118).

## References

- Fabien André, Anne-Marie Kermarrec, and Nicolas Le Scouarnec. 2015. Cache locality is not enough: high-performance nearest neighbor search with product quantization fast scan. *Proc. VLDB Endow.*, 9(4):288–299.
- Ilias Azizi, Karima Echihabi, and Themis Palpanas. 2025. Graph-based vector search: An experimental evaluation of the state-of-the-art. *Proc. ACM Manag. Data*, 3(1).
- Parishad BehnamGhader, Vaibhav Adlakha, Marius Mosbach, Dzmitry Bahdanau, Nicolas Chapados, and Siva Reddy. 2024. LLM2Vec: Large language models are secretly powerful text encoders. In *COLM*.
- Leo Breiman. 2001. Random forests. *Machine learning*, 45(1):5–32.
- Deng Cai. 2016. [A revisit of hashing algorithms for approximate nearest neighbor search](#). *IEEE TKDE*, PP.
- Deng Cai. 2019. [A revisit of hashing algorithms for approximate nearest neighbor search](#). *IEEE TKDE*, 33(6):2337–2348.

- Moses S. Charikar. 2002. Similarity estimation techniques from rounding algorithms. In *STOC*, page 380–388, New York, NY, USA.
- Nick Craswell. 2009. [Mean reciprocal rank](#). In *Encyclopedia of Database Systems*.
- Nick Craswell, Bhaskar Mitra, Emine Yilmaz, Hossein A. Rahmani, Daniel Campos, Jimmy Lin, Ellen M. Voorhees, and Ian Soboroff. 2023. Overview of the trec 2023 deep learning track. In *TREC*.
- Domenico Dato, Sean MacAvaney, Franco Maria Nardini, Raffaele Perego, and Nicola Tonellotto. 2022. The istella22 dataset: Bridging traditional and neural learning to rank evaluation. In *SIGIR, SIGIR '22*, page 3099–3107, New York, NY, USA.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Matthijs Douze, Alexandr Guzhva, Chengqi Deng, Jeff Johnson, Gergely Szilvasy, Pierre-Emmanuel Mazaré, Maria Lomeli, Lucas Hosseini, and Hervé Jégou. 2025. The faiss library. *IEEE TBD*, pages 1–17.
- Kenneth Enevoldsen, Isaac Chung, Imene Kerboua, Márton Kardos, Ashwin Mathur, David Stap, Jay Gala, Wissam Sibli, Dominik Krzemiński, Genta Indra Winata, Saba Sturua, Saiteja Utpala, Mathieu Ciancone, Marion Schaeffer, Gabriel Sequeira, Diganta Misra, Shreeya Dhakal, Jonathan Rystrom, Roman Solomatin, and 67 others. 2025. Mmteb: Massive multilingual text embedding benchmark. *arXiv preprint arXiv:2502.13595*.
- Yunchao Gong, Sanjiv Kumar, Henry A Rowley, and Svetlana Lazebnik. 2013a. Learning binary codes for high-dimensional data using bilinear projections. In *CVPR*, pages 484–491.
- Yunchao Gong, Svetlana Lazebnik, Albert Gordo, and Florent Perronnin. 2013b. Iterative quantization: A procrustean approach to learning binary codes for large-scale image retrieval. *IEEE Trans. Pattern Anal. Mach. Intell.*, 35:2916–29.
- Ruiqi Guo, Philip Sun, Erik Lindgren, Quan Gu, David Simcha, Felix Chern, and Sanjiv Kumar. 2020. Accelerating large-scale inference with anisotropic vector quantization. In *ICML*, pages 3887–3896.
- Kalervo Järvelin and Jaana Kekäläinen. 2002. Cumulated gain-based evaluation of IR techniques. *ACM Trans. Inf. Syst.*, 20(4):422–446.
- Zhongming Jin, Cheng Li, Yue Lin, and Deng Cai. 2013. Density sensitive hashing. *IEEE Trans. Cybern.*, 44(8):1362–1371.
- Jeff Johnson, Matthijs Douze, and Hervé Jégou. 2019. Billion-scale similarity search with GPUs. *IEEE TBD*, 7(3):535–547.
- Herve Jégou, Matthijs Douze, and Cordelia Schmid. 2011. Product quantization for nearest neighbor search. *IEEE Trans. Pattern Anal. Mach. Intell.*, 33(1):117–128.
- Aditya Kusupati, Gantavya Bhatt, Aniket Rege, Matthew Wallingford, Aditya Sinha, Vivek Ramanujan, William Howard-Snyder, Kaifeng Chen, Sham Kakade, Prateek Jain, and 1 others. 2022. Matryoshka representation learning. *NeurIPS*, 35:30233–30249.
- Yinhe Lan, Zhenyu Weng, and Yuesheng Zhu. 2017. A new spherical hashing method in a low-dimensional isotropic space. In *VCIP*, pages 1–4.
- Jinhyuk Lee, Zhuyun Dai, Xiaoqi Ren, Blair Chen, Daniel Cer, Jeremy R Cole, Kai Hui, Michael Boratko, Rajvi Kapadia, Wen Ding, and 1 others. 2024. Gecko: Versatile text embeddings distilled from large language models. *arXiv preprint arXiv:2403.20327*.
- Jinfeng Li, Xiao Yan, Jian Zhang, An Xu, James Cheng, Jie Liu, Kelvin K. W. Ng, and Ti-chung Cheng. 2018. A general and efficient querying method for learning to hash. In *SIGMOD*, page 1333–1347, New York, NY, USA.
- Zehan Li, Xin Zhang, Yanzhao Zhang, Dingkun Long, Pengjun Xie, and Meishan Zhang. 2023. Towards general text embeddings with multi-stage contrastive learning. *arXiv preprint arXiv:2308.03281*.
- Victor Weixin Liang, Yuhui Zhang, Yongchan Kwon, Serena Yeung, and James Y Zou. 2022. Mind the gap: Understanding the modality gap in multi-modal contrastive representation learning. *NeurIPS*, 35:17612–17625.
- Jimmy Lin. 2025. [Operational advice for dense and sparse retrievers: HNSW, flat, or inverted indexes?](#) In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 6: Industry Track)*, pages 865–872, Vienna, Austria. Association for Computational Linguistics.
- Fei Tony Liu, Kai Ming Ting, and Zhi-Hua Zhou. 2008. Isolation forest. In *IEEE ICDM*, pages 413–422.
- Huawen Liu, Wenhua Zhou, Zongda Wu, Shichao Zhang, Gang Li, and Xuelong Li. 2024a. Refining codes for locality sensitive hashing. *IEEE TKDE*, 36(3):1274–1284.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized bert pretraining approach](#). *Preprint*, arXiv:1907.11692.

- Zheng Liu, Chaofan Li, Shitao Xiao, Yingxia Shao, and Defu Lian. 2024b. [Llama2Vec: Unsupervised adaptation of large language models for dense retrieval](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3490–3500, Bangkok, Thailand. Association for Computational Linguistics.
- Yu A Malkov and Dmitry A Yashunin. 2018. Efficient and robust approximate nearest neighbor search using hierarchical navigable small world graphs. *IEEE Trans. Pattern Anal. Mach. Intell.*, 42(4):824–836.
- Niklas Muennighoff, Nouamane Tazi, Loic Magne, and Nils Reimers. 2023. [MTEB: Massive text embedding benchmark](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 2014–2037, Dubrovnik, Croatia. Association for Computational Linguistics.
- Zhijie Nie, Zhangchi Feng, Mingxin Li, Cunwang Zhang, Yanzhao Zhang, Dingkun Long, and Richong Zhang. 2024. [When text embedding meets large language model: A comprehensive survey](#). *CoRR*, abs/2412.09165.
- Zach Nussbaum, John X. Morris, Brandon Duderstadt, and Andriy Mulyar. 2025. Nomic embed: Training a reproducible long context text embedder. *TMLR*.
- OpenAI. 2024a. New embedding models and api updates.
- OpenAI. 2024b. New embedding models and api updates. <https://openai.com/index/new-embedding-models-and-api-updates>.
- Yun Peng, Byron Choi, Tsz Nam Chan, Jianye Yang, and Jianliang Xu. 2023. Efficient approximate nearest neighbor search in multi-dimensional databases. *Proc. ACM Manag. Data*, 1(1).
- Zexuan Qiu, Qinliang Su, Jianxing Yu, and Shijing Si. 2022. [Efficient document retrieval by end-to-end refining and quantizing BERT embedding with contrastive product quantization](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 853–863, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *JMLR*, 21(140):1–67.
- Suhas Jayaram Subramanya, Devvrit, Rohan Kadekodi, Ravishankar Krishaswamy, and Harsha Vardhan Simhadri. 2019. *DiskANN: fast accurate billion-point nearest neighbor search on a single node*. Curran Associates Inc., Red Hook, NY, USA.
- Nandan Thakur, Luiz Bonifacio, Maik Fröbe, Alexander Bondarenko, Ehsan Kamalloo, Martin Potthast, Matthias Hagen, and Jimmy Lin. 2024. Systematic evaluation of neural retrieval models on the Touché 2020 argument retrieval subset of BEIR. In *SIGIR*.
- 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 *NeurIPS*.
- 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, Volume 1 (Long Papers)*, pages 809–819, New Orleans, Louisiana. Association for Computational Linguistics.
- Kai Ming Ting, Takashi Washio, Ye Zhu, Yang Xu, and Kaifeng Zhang. 2024. Is it possible to find the single nearest neighbor of a query in high dimensions? *AIJ*, 336:104206.
- Kai Ming Ting, Yue Zhu, and Zhi-Hua Zhou. 2018. Isolation kernel and its effect on svm. In *SIGKDD*, pages 2329–2337.
- L. Van Der Maaten, E. Postma, J. Vanden Herik, and et al. 2009. Dimensionality reduction: a comparative. *JMLR*, 10:66–71.
- Mengzhao Wang, Weizhi Xu, Xiaomeng Yi, Songlin Wu, Zhangyang Peng, Xiangyu Ke, Yunjun Gao, Xiaoliang Xu, Rentong Guo, and Charles Xie. 2024. Starling: An i/o-efficient disk-resident graph index framework for high-dimensional vector similarity search on data segment. *Proc. ACM Manag. Data*, 2(1).
- Jiuqi Wei, Xiaodong Lee, Zhenyu Liao, Themis Palpanas, and Botao Peng. 2025. Subspace collision: an efficient and accurate framework for high-dimensional approximate nearest neighbor search. *Proc. ACM Manag. Data*, 3(1):1–29.
- Yair Weiss, Antonio Torralba, and Rob Fergus. 2008. Spectral hashing. In *NeurIPS*, volume 21.
- Tiansheng Wen, Yifei Wang, Zequn Zeng, Zhong Peng, Yudi Su, Xinyang Liu, Bo Chen, Hongwei Liu, Stefanie Jegelka, and Chenyu You. 2025. Beyond matryoshka: Revisiting sparse coding for adaptive representation. In *ICML*.
- Yan Xia, Kaiming He, Pushmeet Kohli, and Jian Sun. 2015. Sparse projections for high-dimensional binary codes. In *CVPR*, pages 3332–3339.
- Yang Xu and Kai Ming Ting. 2025. Voronoi diagram encoded hashing. In *ECML PKDD*, page 87–103, Berlin, Heidelberg.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. 2018. [HotpotQA: A dataset for diverse, explainable multi-hop question answering](#).

In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2369–2380, Brussels, Belgium. Association for Computational Linguistics.

Felix Yu, Sanjiv Kumar, Yunchao Gong, and Shih-Fu Chang. 2014. Circulant binary embedding. In *ICML*, volume 32 of *Proceedings of Machine Learning Research*, pages 946–954, Beijing, China.

Peng Yu, Luke Merrick, Giovanni Nuti, and Daniel Campos. 2024. Arctic-embed 2.0: Multilingual retrieval without compromise. *arXiv preprint arXiv:2412.04506*.

Hailin Zhang, Xiaodong Ji, Yilin Chen, Fangcheng Fu, Xupeng Miao, Xiaonan Nie, Weipeng Chen, and Bin Cui. 2025a. Pqcache: Product quantization-based kvcache for long context llm inference. *Proc. ACM Manag. Data*, 3(3).

Yanzhao Zhang, Mingxin Li, Dingkun Long, Xin Zhang, Huan Lin, Baosong Yang, Pengjun Xie, An Yang, Dayiheng Liu, Junyang Lin, Fei Huang, and Jingren Zhou. 2025b. Qwen3 embedding: Advancing text embedding and reranking through foundation models. *arXiv preprint arXiv:2506.05176*.

Guang-Tong Zhou, Kai Ming Ting, Fei Tony Liu, and Yilong Yin. 2012. Relevance feature mapping for content-based multimedia information retrieval. *Pattern Recogn.*, 45(4):1707–1720.

Xiaofeng Zhu, Lei Zhang, and Zi Huang. 2014. A sparse embedding and least variance encoding approach to hashing. *IEEE Trans. Image Process.*, 23(9):3737–3750.

## A Related Work

### A.1 LLM Embedded Representation

Text representation plays a crucial role in tasks such as semantic matching, clustering, and information retrieval (Nie et al., 2024). In recent years, pre-trained language models with millions or even billions of parameters have been proposed for a wide range of NLP tasks, including text representation. Early approaches relied on encoder-based models such as BERT (Devlin et al., 2019) and RoBERTa (Liu et al., 2019). More recently, there has been a paradigm shift toward leveraging large language models (LLMs) for text embeddings, with notable examples including LLM2Vec (BehnamGhader et al., 2024), Llama2Vec (Liu et al., 2024b), and Qwen Embedding (Li et al., 2023; Zhang et al., 2025b). These LLM-based methods now represent the state of the art in text representation.

However, language model-based embeddings typically have high dimension: for example, the hidden states of T5-11B (Raffel et al., 2020) have 1024 dimensions, while mainstream 7B decoder-only LLMs use 4096 dimensions. Such high-dimensional embeddings substantially increase storage and inference costs. To address this, Kusupati et al. (2022) introduced a flexible technique that enables fine-grained control over embedding size, balancing latency and accuracy. This approach has since been extended to various applications (OpenAI, 2024b; Nussbaum et al., 2025; Yu et al., 2024). Recently, Wen et al. (2025) proposed CSR, which combines sparse coding with contrastive learning to produce efficient sparse embeddings with minimal performance loss.

Similar to (Kusupati et al., 2022; Wen et al., 2025), our study also targets representation efficiency and is complementary to the aforementioned LLM-based representation. However, while these studies rely on supervised learning to reduce performance degradation relative to LLM embeddings, our method does not require learning and instead leverages IK for transformation.

### A.2 Efficient Retrieval

For efficient retrieval in high-dimensional spaces, techniques such as dimensionality reduction (Van Der Maaten et al., 2009), product quantization (Guo et al., 2020; Jégou et al., 2011; André et al., 2015), and hashing (Charikar, 2002; Xu and Ting, 2025) are commonly applied, though often at the expense of accuracy. Our study falls within the

hashing paradigm, yet demonstrates strong potential to preserve—and in some cases even improve—retrieval accuracy.

Hashing techniques can be broadly categorized into training-based approaches (Learning to Hash, L2H) (Liu et al., 2024a; Weiss et al., 2008; Yu et al., 2014; Gong et al., 2013b,a; Xia et al., 2015; Jin et al., 2013; Zhu et al., 2014; Lan et al., 2017) and learning-free approaches (Charikar, 2002; Cai, 2019; Xu and Ting, 2025). L2H aims to map high-dimensional vectors into a binary coding space for efficient processing and retrieval. In general, learning is to ensure three key criteria (Xu and Ting, 2025) for an effective hashing: full space coverage, entropy maximization, and bit independence.

Recently, Xu and Ting (2025) showed that Voronoi Diagrams (VD) offer a natural alternative to L2H, as they inherently satisfy these three criteria. Building on this insight, they proposed VDeH, a learning-free method that leverages a VD-based implementation of the Isolation Kernel for effective hashing. Though VDeH has been shown to outperform many L2H methods for non-LLM-related databases, our early examination has revealed that VDeH performs poorly in LLM-embedded databases (see Figure 2). To address this limitation, we propose a generalized version of VDeH that is amenable to LLM-embedded databases in this paper. It is worth noting that prior work (Zhou et al., 2012) also explored the application of Isolation Forest in the field of retrieval. Their method employs the path length of data points on each Isolation Tree as the mapped feature representation, along with a weighted-sum similarity calculation, rather than adopting a hashing paradigm. This, however, results in high computational overhead and difficulty in low-bit compression.

## B Extension of section 2

### B.1 Detailed Description of Isolation Kernel

In practice, Isolation Kernel  $K_\psi$  is usually estimated from a finite number of  $t$  partitions  $H_i$  ( $i = 1, \dots, t$ ) and each  $H_i$  is derived from a data subset

$\mathcal{D}_i \subset D$ :

$$\begin{aligned} K_\psi(\mathbf{x}, \mathbf{y} \mid D) &\simeq \frac{1}{t} \sum_{i=1}^t \mathbb{I}(\mathbf{x}, \mathbf{y} \in \theta \mid \theta \in H_i) \\ &= \frac{1}{t} \sum_{i=1}^t \sum_{\theta_j \in H_i} \mathbb{I}(\mathbf{x} \in \theta_j) \mathbb{I}(\mathbf{y} \in \theta_j). \end{aligned} \quad (1)$$

Given  $H_i$ , there exist  $\psi$  isolating partitions  $\theta_j \in H_i$  ( $j=1, \dots, \psi$ ). For any data point  $\mathbf{x}$ , we define

$$\Phi_{ij}(\mathbf{x} \mid D) = \mathbb{I}(\mathbf{x} \in \theta_j \mid \theta_j \in H_i).$$

Since  $\mathbf{x}$  can belong to exactly one of these  $\psi$  isolating partitions, the vector

$$\begin{aligned} \Phi_i(\mathbf{x} \mid D) \\ = [\Phi_{i1}(\mathbf{x} \mid D), \Phi_{i2}(\mathbf{x} \mid D), \dots, \Phi_{i\psi}(\mathbf{x} \mid D)] \end{aligned}$$

is a  $\psi$ -dimensional one-hot encoding. Given  $t$  partitions,  $\Phi(\mathbf{x} \mid D)$  is the concatenation of  $\Phi_1(\mathbf{x} \mid D), \dots, \Phi_t(\mathbf{x} \mid D)$ .

In practice, each data point  $\mathbf{x}$  is mapped into a binary feature space of dimension  $t \times \psi$  via the mapping  $\Phi : \mathbf{x} \mapsto \{0, 1\}^{t \times \psi}$ . Under this representation, the Isolation Kernel  $K_\psi$  can be approximated by the normalized inner product:

$$K_\psi(\mathbf{x}, \mathbf{y} \mid D) \simeq \frac{1}{t} \langle \Phi(\mathbf{x} \mid D), \Phi(\mathbf{y} \mid D) \rangle. \quad (2)$$

For notational brevity, we henceforth omit the conditioning on  $D$  whenever the meaning is clear.

Since the mapped data point  $\Phi(\mathbf{x})$  is highly sparse, with only  $t$  out of the  $t \times \psi$  entries equal to 1, we store it in an index-based sparse representation  $\Phi^{\text{idx}}(\mathbf{x}) \in \mathbb{N}^t$ . Specifically, the  $i$ -th component of  $\Phi^{\text{idx}}(\mathbf{x})$  indicates the position of the sole ‘‘1’’ in the one-hot vector  $\Phi_i(\mathbf{x})$  for all  $i \in [1, t]$ . Formally, we define

$$\begin{aligned} \Phi^{\text{idx}}(\mathbf{x}) &= [\Phi_1^{\text{idx}}(\mathbf{x}), \Phi_2^{\text{idx}}(\mathbf{x}), \dots, \Phi_t^{\text{idx}}(\mathbf{x})] \\ \Phi_i^{\text{idx}}(\mathbf{x}) &= \arg \max_{1 \leq j \leq \psi} \Phi_{ij}(\mathbf{x}). \end{aligned}$$

By storing each element of  $\Phi^{\text{idx}}(\mathbf{x})$  using  $\lceil \log_2(\psi) \rceil$  bits, the code length of each mapped data point becomes  $t \lceil \log_2(\psi) \rceil$  bits, significantly reducing the memory requirements. Thus, we can approximate the Isolation Kernel  $K_\psi$  as:

$$K_\psi(\mathbf{x}, \mathbf{y} \mid D) \simeq \frac{1}{t} \sum_{i=1}^t \mathbb{I}(\Phi_i^{\text{idx}}(\mathbf{x}) = \Phi_i^{\text{idx}}(\mathbf{y})). \quad (3)$$

## C Theoretical Guarantees for IKE

This section formally defines the Voronoi Diagram-based Isolation Kernel Embedding ( $\text{IKE}_{VD}$ ) operating on a randomly selected subspace (see Appendix C.1) and establishes that it preserves the three key criteria demonstrated for VDeH (Xu and Ting, 2025): **full space coverage**, **entropy maximization**, and **bit independence**. The proofs follow directly from the theoretical framework of (Xu and Ting, 2025); since the core mechanisms of random sampling and partitioning remain unchanged, we show that adding a random subspace projection does not affect the original guarantees. For brevity, we refer to the corresponding proofs and streamline the arguments where the reasoning is identical (see Appendix C.2).

Furthermore, under suitable distributional assumptions, we prove that the iForest-based IKE also satisfies these three criteria (see Appendix C.3). Finally, by viewing both VDeH and IKE as ensemble methods, we introduce **diverse partitioning** as a fourth criterion. We demonstrate that, compared to VDeH, IKE produces more robust hash representations due to its ability to generate a more diverse set of base partitioners (see Appendix C.4).

### C.1 Formal Definition of $\text{IKE}_{VD}$

Let  $\mathcal{X} = [x_1, \dots, x_N] \in \mathbb{R}^{d \times N}$  be the input dataset. An  $\text{IKE}_{VD}$  partition is constructed via a three-step process: (1) We randomly select  $m$  dimensions from the original  $d$ -dimensional embedding, which is equivalent to applying a projection function  $P_J : \mathbb{R}^d \rightarrow \mathbb{R}^m$  that maps points to a random  $m$ -dimensional subspace; (2) Randomly sampling  $\psi$  anchor points  $\mathcal{D} = \{s_1, \dots, s_\psi\} \subset \mathcal{X}$ ; (3) Constructing a Voronoi Diagram  $H_J$  in the  $\mathbb{R}^m$  subspace using the projected anchor points  $P_J(\mathcal{D})$ . For any point  $x \in \mathbb{R}^d$ , its hash value is determined by which of the  $\psi$  Voronoi cells its projection  $P_J$  falls into. An  $\text{IKE}_{VD}$  model consists of  $t$  such partitions, each constructed independently.

### C.2 Proof of Criteria of $\text{IKE}_{VD}$

**Proposition 1** (Full Space Coverage). *The hashing scheme of  $\text{IKE}_{VD}$  covers the entire input space, ensuring that every point can be assigned a binary code.*

*Proof.* The  $\text{IKE}_{VD}$  scheme inherently provides full space coverage. For any point  $x \in \mathbb{R}^d$ , the projection  $P_J$  maps it to a unique point  $P_J(x)$  in the subspace  $\mathbb{R}^m$ . By definition, a Voronoi Diagram

forms a complete partition of the space in which it is constructed. Therefore,  $P_J(x)$  is guaranteed to fall into exactly one Voronoi cell. This ensures that every point  $x$  in the original  $d$ -dimensional space is assigned a hash code.  $\square$

**Proposition 2** (Entropy Maximization). *For a given dataset, every hash function in  $\text{IKE}_{VD}$  covers approximately the same number of data points.*

*Proof.* Based on Lemma 1 of the existing work (Xu and Ting, 2025), for a partition constructed from  $\psi$  anchor points drawn *i.i.d.* from the data distribution, the probability of an arbitrary point falling into any given Voronoi cell is uniformly  $1/\psi$ .

This criterion remains valid for  $\text{IKE}_{VD}$ . The anchor points are still sampled *i.i.d.* from the data distribution. When these points and the data are projected onto the same random subspace  $\mathbb{R}^m$ , their underlying distributional relationship is preserved. The entropy maximization criterion of VDeH is based on the random sampling process and is agnostic to the space’s dimensionality. Thus, its conclusion holds true within the subspace  $\mathbb{R}^m$ :

$$\Pr(P_J(x) \in V_i) = \frac{1}{\psi},$$

for every cell  $i = 1, \dots, \psi$ . This uniform probability ensures entropy maximization.  $\square$

**Proposition 3** (Bit Independence). *All bits in the final concatenated binary code generated by  $\text{IKE}_{VD}$  are mutually independent.*

*Proof.* The bit independence criterion in VDeH relies on two aspects, both of which are maintained in  $\text{IKE}_{VD}$ . We assume the standard encoded hashing mechanism, which maps a cell index from a partition of size  $\psi = 2^w$  to a  $w$ -bit binary code.

First,  $\text{IKE}_{VD}$  holds the independence between partitions. As with VDeH, each of the  $t$  partitions in  $\text{IKE}_{VD}$  is constructed independently. This independence is further reinforced in our case by the fact that each partition is constructed not only with a new set of random anchor points but also in a new, independently chosen random subspace. This ensures the statistical independence of the  $w$ -bit codes generated from different partitions. Second,  $\text{IKE}_{VD}$  holds the independence within a partition. The proof for this case, detailed in Theorem 1 of the existing work (Xu and Ting, 2025), is conditional on the entropy maximization criterion holding. Specifically, it requires the probability of a

point falling into any of the  $\psi$  cells to be  $1/\psi$ , which is also satisfied in  $\text{IKE}_{VD}$ . Since the prerequisite holds and the same encoding logic is applied, the conclusion of Theorem 1 in VDeH directly applies: any two bits generated from the same encoded partition are mutually independent.  $\square$

### C.3 Proof of Criteria of IKE (iForest)

For iForest-based IKE, as with  $\text{IKE}_{VD}$ , the full space coverage criterion is naturally satisfied, while the bit independence criterion depends on the entropy maximization.

First, an iTree partitions space using axis-aligned hyperplanes. The set of all leaf nodes forms a collection of disjoint hyper-rectangles whose union completely tiles the  $\mathbb{R}^d$  space. Any point  $x \in \mathbb{R}^d$  can be deterministically routed through the tree’s splits to a unique leaf node. Thus, full space coverage is satisfied, irrespective of the data distribution.

**Theorem 1.** *Under the assumption of a uniform data distribution, the hashing scheme of IKE satisfies the entropy maximization and the bit independence criteria.*

*Proof.* Let an iTree be constructed to have  $\psi$  leaf nodes  $\{V_1, \dots, V_\psi\}$ . We need to show that for an arbitrary point  $x$ , the probability of it landing in any leaf node  $V_i$  is uniform, i.e.,  $\Pr(x \in V_i) = 1/\psi$ . Consider a single split at an arbitrary node in the tree. The split value  $p$  is chosen uniformly at random from the range  $[m_q, M_q]$ , where  $m_q$  and  $M_q$  are the minimum and maximum values of the data points within that node along the chosen dimension  $q$ . Under the uniform data distribution, the probability mass within any sub-region is proportional to its volume. When a split point  $p$  is chosen uniformly within the range  $[m_q, M_q]$ , the probability that a point  $x$  within this region falls on the left side of the split ( $x_q < p$ ) is given by  $\frac{p-m_q}{M_q-m_q}$ . The expected probability of going left, over the random choice of  $p$ , is:

$$\begin{aligned} E_p [\Pr(x \text{ goes left})] &= E_p \left[ \frac{p - m_q}{M_q - m_q} \right] \\ &= \frac{E_p[p] - m_q}{M_q - m_q}. \end{aligned}$$

Since  $p$  is uniform in  $[m_q, M_q]$ , its expectation is

$E_p[p] = \frac{m_q + M_q}{2}$ . Substituting this in, we get:

$$\begin{aligned} E_p [\Pr(x \text{ goes left})] &= \frac{\frac{m_q + M_q}{2} - m_q}{M_q - m_q} \\ &= \frac{\frac{M_q - m_q}{2}}{M_q - m_q} = \frac{1}{2}. \end{aligned}$$

Thus, at each node, a data point has an equal probability of  $1/2$  of being routed to either child node. For a point to reach a specific leaf node, it must follow a unique path of splits from the root. The probability of any specific path is therefore uniform across all leaves, leading to  $\Pr(x \in V_i) = 1/\psi$ , which establishes the entropy maximization criterion.

The bit independence criterion follows from the establishment of entropy maximization. Assuming the standard encoded hashing mechanism where a leaf index from a partition of size  $\psi = 2^w$  is mapped to a  $w$ -bit binary code. The proof considers two cases: (i) The independence between partitions is guaranteed because each of the  $t$  iTrees in the ensemble is constructed independently. This ensures the statistical independence of the  $w$ -bit codes generated from different trees. (ii) The independence within a single partition. Let  $\alpha$  and  $\beta$  be two distinct bits in the  $w$ -bit vector generated from one iTREE. For any values  $v_\alpha, v_\beta \in \{0, 1\}$ , the joint probability is the sum of probabilities of landing in all leaf nodes  $V_i$  whose code satisfies the condition:

$$\begin{aligned} \Pr(\alpha = v_\alpha, \beta = v_\beta) &= \sum_{i=1}^{\psi} \Pr(x \in V_i) \cdot \mathbb{I}(\alpha = v_\alpha, \beta = v_\beta \text{ for cell } i) \\ &= \frac{1}{\psi} \sum_{i=1}^{\psi} \mathbb{I}(\alpha = v_\alpha, \beta = v_\beta). \end{aligned}$$

Since the  $\psi = 2^w$  binary codes exhaust all possible  $w$ -bit vectors, exactly  $\psi/4$  of them will satisfy the condition that two specific bits have the values  $v_\alpha$  and  $v_\beta$ . Therefore:

$$\Pr(\alpha = v_\alpha, \beta = v_\beta) = \frac{1}{\psi} \left( \frac{\psi}{4} \right) = \frac{1}{4}.$$

Similarly, the marginal probabilities are:

$$\Pr(\alpha = v_\alpha) = \frac{1}{\psi} \sum_{i=1}^{\psi} \mathbb{I}(\alpha = v_\alpha) = \frac{1}{\psi} \left( \frac{\psi}{2} \right) = \frac{1}{2},$$

$$\Pr(\beta = v_\beta) = \frac{1}{\psi} \sum_{i=1}^{\psi} \mathbb{I}(\beta = v_\beta) = \frac{1}{\psi} \left( \frac{\psi}{2} \right) = \frac{1}{2}.$$

Since  $\Pr(\alpha = v_\alpha, \beta = v_\beta) = \Pr(\alpha = v_\alpha) \cdot \Pr(\beta = v_\beta)$ , the bits are mutually independent. This completes the proof.  $\square$

Therefore, under the uniform distribution, IKE satisfies three key criteria, which is to implement effective binary codes.

In fact, LLM embedding is non-uniform, which makes IKE appear to perform poorly due to not strictly satisfying the two criteria. However, IKE's strength lies in its use of a large ensemble of  $t$  iTrees. While a single tree may produce biased partitions and correlated bits, the final similarity score is an aggregation over thousands of independent random processes. The biases and errors of individual trees are averaged out, yielding a robust model where the strong independence between trees compensates for the lack of strict independence within them. In our experience, IKE is even better than  $\text{IKE}_{VD}$  when  $m = 1$ , which satisfies all three criteria.

#### C.4 Diverse Partitioning

It is noteworthy that VDeH satisfies the three criteria discussed above, yet it performs poorly on LLM embeddings. This indicates that these criteria alone are insufficient for effective retrieval. In the following, we demonstrate that diverse partitioning accounts for the superior retrieval performance of IKE compared to VDeH.

**Diverse partitioning as the fourth criterion.** As VDeH and IKE can be expressed as ensemble methods of a similarity kernel function  $K(\mathbf{x}, \mathbf{y})$  (Ting et al., 2018, 2024), the superiority of IKE stems from a key principle in ensemble learning: diverse, simple base learners often outperform ensembles of complex, correlated ones.

**Definition 1 (Base Partitioner).** A base partitioner  $h(\mathbf{x}; \Theta)$  maps  $\mathbf{x} \in \mathbb{R}^d$  to a discrete partition, with  $\Theta$  as a random variable capturing all randomness:

- For VDeH,  $\Theta_{\text{VDeH}}$  consists of  $\psi$  randomly sampled anchor points from the corpus.
- For IKE,  $\Theta_{\text{IKE}}$  includes random dimension selections, and value splits during iTREE construction.

**Definition 2 (Isolation Kernel).** Given a base partitioner  $h(\cdot; \Theta)$  (Definition 1), the Isolation Kernel  $K_\psi(\mathbf{x}, \mathbf{y})$  is the expectation that  $\mathbf{x}$  and  $\mathbf{y}$  are mapped to the same partition:

$$K_\psi(\mathbf{x}, \mathbf{y}) = \mathbb{E}_\Theta[\mathbb{I}(h(\mathbf{x}; \Theta) = h(\mathbf{y}; \Theta))], \quad (4)$$

This is consistent with the definition in Section 2, re-expressed using the base partitioner notation.

Intuitively, if two points  $\mathbf{x}$  and  $\mathbf{y}$  are close in the data space, they are more likely to fall into the same isolating partition, yielding a larger kernel value  $K_\psi(\mathbf{x}, \mathbf{y})$ . Conversely, distant points are less likely to co-occur in a partition, yielding a smaller value.

In practice, both IKE and VDeH estimate  $K_\psi(\mathbf{x}, \mathbf{y})$  using a finite ensemble of  $t$  base partitioners,  $\{h_1, \dots, h_t\}$ . Each partitioner  $h_i$  is an instance  $h(\cdot; \Theta_i)$ , generated from an i.i.d. realization  $\Theta_i$  of the random parameter  $\Theta$ :

$$K(\mathbf{x}, \mathbf{y}; t) = \frac{1}{t} \sum_{i=1}^t \mathbb{I}(h(\mathbf{x}; \Theta_i) = h(\mathbf{y}; \Theta_i)). \quad (5)$$

The quality of the estimator  $K(\mathbf{x}, \mathbf{y}; t)$  is measured as follows:

$$\begin{aligned} \text{MSE}(K(\mathbf{x}, \mathbf{y}; t)) &= \mathbb{E}[(K(\mathbf{x}, \mathbf{y}; t) - K_\psi(\mathbf{x}, \mathbf{y}))^2] \\ &= \text{Bias}^2 + \text{Variance}. \end{aligned} \quad (6)$$

Since each  $Z_i = \mathbb{I}(h(\mathbf{x}; \Theta_i) = h(\mathbf{y}; \Theta_i))$  is i.i.d. with  $\mathbb{E}[Z_i] = K_\psi(\mathbf{x}, \mathbf{y})$ , the sample mean  $K(\mathbf{x}, \mathbf{y}; t) = \frac{1}{t} \sum_i Z_i$  is an unbiased estimator of  $K_\psi(\mathbf{x}, \mathbf{y})$ . Thus,  $\text{MSE}(K(\mathbf{x}, \mathbf{y}; t))$  is determined entirely by the variance term. Let  $p = K_\psi(\mathbf{x}, \mathbf{y})$  and  $\sigma^2 = p(1-p)$ . The variance of the estimator is

$$\begin{aligned} \text{Var}(K(\mathbf{x}, \mathbf{y}; t)) \\ = \frac{1}{t^2} \left( \sum_{i=1}^t \text{Var}(Z_i) + \sum_{i \neq j} \text{Cov}(Z_i, Z_j) \right). \end{aligned} \quad (7)$$

Letting  $\rho = \text{Cov}(Z_i, Z_j)/\sigma^2$  for  $i \neq j$  gives the following standard identity for the variance of the mean of identically distributed random variables (Breiman, 2001):

$$\text{Var}(K(\mathbf{x}, \mathbf{y}; t)) = \sigma^2 \left( \frac{1-\rho}{t} + \rho \right). \quad (8)$$

**Corollary 1** (Asymptotic Variance Limit). *As the ensemble size  $t \rightarrow \infty$ , the variance of  $K(\mathbf{x}, \mathbf{y}; t)$  converges to:*

$$\lim_{t \rightarrow \infty} \text{Var}(K(\mathbf{x}, \mathbf{y}; t)) = \rho \sigma^2 \quad (9)$$

Thus, the quality of the estimator  $K(\mathbf{x}, \mathbf{y}; t)$  is ultimately constrained by the correlation  $\rho$ .

**Theorem 2.** *Assuming that the dimensions in the input space are independent, the correlations of the VDeH,  $\text{IKE}_{VD}$ , and IKE partitioning schemes satisfy the following inequality for dimension  $d > 1$ :*

$$\rho_{\text{VDeH}} > \rho_{\text{IKE}_{VD}(m=1)} > \rho_{\text{IKE}} \geq 0. \quad (10)$$

*Proof.* For VDeH, the randomness is driven solely by the data sample  $D_\psi$ . Therefore, its correlation is entirely determined by the data, yielding  $\rho_{\text{VDeH}} = \rho_{\text{data}}$ . For  $\text{IKE}_{VD}(m=1)$ , randomness comes from  $D_\psi$  and a random dimension choice  $q \in \{1, \dots, d\}$ . We analyze its covariance  $\text{Cov}_{VD,1}(Z_i, Z_j)$  using the law of total covariance, conditioning on the dimension choices  $q_i, q_j$ :

$$\begin{aligned} \mathbb{E}_{q_i, q_j}[\text{Cov}(Z_i, Z_j | q_i, q_j)] \\ + \text{Cov}_{q_i, q_j}(\mathbb{E}[Z_i | q_i], \mathbb{E}[Z_j | q_j]). \end{aligned} \quad (11)$$

By dimension independence,  $\mathbb{E}[Z | q]$  is constant, making the second term zero. The covariance is thus equal to the first term:

- Case A ( $q_i = q_j$ ): Occurs with probability  $1/d$ , giving  $\text{Cov}(Z_i, Z_j | q_i = q_j) = \sigma^2 \rho_{\text{data}}$ .
- Case B ( $q_i \neq q_j$ ): Occurs with probability  $1 - 1/d$ , yielding zero conditional covariance due to dimension independence.

Combining these cases yields:

$$\rho_{\text{IKE}_{VD}(m=1)} = \frac{\text{Cov}_{VD,1}(Z_i, Z_j)}{\sigma^2} = \frac{\rho_{\text{data}}}{d}. \quad (12)$$

For  $d > 1$ , we have the following inequality

$$\rho_{\text{VDeH}} > \rho_{\text{IKE}_{VD}(m=1)}. \quad (13)$$

The randomness in IKE,  $\Theta_{\text{IKE}} = (D_\psi, \mathcal{R})$ , involves the data sample  $D_\psi$  and a rich set of subsequent random choices  $\mathcal{R}$  to construct the entire iTree (i.e., random dimension selections and value splits). This randomness further reduces the correlation  $\rho$  between partitioners. Thus, we have

$$\rho_{\text{IKE}_{VD}(m=1)} > \rho_{\text{IKE}} \geq 0. \quad (14)$$

Combing Eq.(13) and Eq.(14) completes the proof.  $\square$

Theorem 2 shows that the randomization process employed by IKE makes it superior in generating diverse base partitioners, leading to a robust binary hash representation.

## D Method Details

This section provides details of our method. The code is publicly available at: [https://github.com/Zed-Zed-b/IK\\_Emb](https://github.com/Zed-Zed-b/IK_Emb).

### D.1 Construction Details of Isolation Tree (iTree)

Given a dataset  $\mathcal{D} = \{\mathbf{x}_1, \dots, \mathbf{x}_\psi\} \subset \mathbb{R}^d$ , one first chooses a feature  $q$  uniformly at random from the set of all features of  $\mathcal{D}$ . Next, a split value  $p$  is drawn uniformly from the interval  $[\min_q(\mathcal{D}), \max_q(\mathcal{D})]$ , where  $\min_q(\mathcal{D})$  and  $\max_q(\mathcal{D})$  denote the minimum and maximum values of feature  $q$  over all points in  $\mathcal{D}$ . The dataset  $\mathcal{D}$  is then partitioned into two disjoint subsets:  $\mathcal{D}_\ell = \{\mathbf{x} \in \mathcal{D} \mid x_q < p\}$ ,  $\mathcal{D}_r = \{\mathbf{x} \in \mathcal{D} \mid x_q \geq p\}$ . This splitting process is applied recursively to  $\mathcal{D}_\ell$  and  $\mathcal{D}_r$ , with the current tree height  $e$  incremented by one at each recursive call. Recursion terminates when any of the following conditions is satisfied: (i)  $|\mathcal{D}| = 1$ , (ii) the current tree height  $e$  reaches the specified height limit  $l$ , or (iii) the chosen split  $(q, p)$  fails to produce two nonempty subsets (i.e.,  $\mathcal{D}_\ell = \emptyset$  or  $\mathcal{D}_r = \emptyset$ ). The partitioning in iTree is demonstrated in Figure 3a. It is noteworthy that the entire construction process of iTree does not require learning.

### D.2 Formal Description for $\text{IKE}_{VD}$

Let  $D = [x_1, \dots, x_N] \in \mathbb{R}^{d \times N}$  be the input dataset. An  $\text{IKE}_{VD}$  partition is constructed via a three-step process: (1) Randomly selecting  $m$  dimensions ( $m \leq d$ ); (2) Randomly sampling  $\psi$  anchor points  $\mathcal{D} = \{s_1, \dots, s_\psi\} \subset D$ ; (3) Constructing a Voronoi Diagram  $H_J$  using the only selected dimensions of anchor points. For any point  $x \in \mathbb{R}^d$ , its transformed value is determined by which of the  $\psi$  Voronoi cells it falls into (see the demonstration in Figure 3b).

### D.3 An Example for Efficient Similarity Computation

This section provides a step-by-step example illustrating the efficient computation of matching elements between two bit strings. Let  $D_x$  and  $D_y$  be two 8-bit strings, each segmented into  $n_b = 2$  bit element. Thus, each string consists of  $8/2 = 4$  elements. The goal is to count the number of positions where the corresponding elements match exactly.

Assume the values of  $D_x$  and  $D_y$  are:

$$D_x = 00\ 01\ 10\ 01$$

$$D_y = 00\ 10\ 11\ 01$$

The bitwise XOR operation identifies positions where the two bits differ. A segment of all zeros after XOR indicates a matching element.

$$D = D_x \oplus D_y = 00\ 11\ 01\ 00$$

Checking each segment for all zeros directly is inefficient. Instead, a logical right shift followed by a bitwise OR is used to propagate any '1' within a segment to its rightmost bit. If the rightmost bit is '0' after this operation, it indicates that all bits in this segment are 0 (meaning the element matches):

$$\begin{aligned} M &= D \mid (D \gg 1) \\ &= 00\ 11\ 01\ 00 \mid 00\ 01\ 10\ 10 \\ &= \mathbf{00\ 11\ 01\ 10} \end{aligned}$$

The bold bits represent the rightmost bit of each segment.

To count segments where the rightmost bit is '0', all other bits (non-rightmosts) in  $M$  are set to '1' by OR-ing with a pre-defined mask. The mask for  $n_b = 2$  is 10 10 10 10. For  $n_b = 4$ , it would be 1110 1110.

$$\begin{aligned} M &= M \mid \text{mask} \\ &= \mathbf{00\ 11\ 01\ 10} \mid 10\ 10\ 10\ 10 \\ &= \mathbf{10\ 11\ 11\ 10} \end{aligned}$$

The number of '1' bits in  $M$  is counted using the population count (popcnt) function. The number of matching elements equals the length of  $M$  minus this count,

$$\begin{aligned} n_1 &= \text{popcnt}(M) = 6 \\ n_{\text{match}} &= \text{length}(M) - n_1 = 2 \end{aligned}$$

Thus, between  $D_x$  and  $D_y$ , two out of four elements match.

### D.4 Discussion on Database Scalability and Distribution Shifts

IKE natively supports efficient incremental updates without requiring a full system rebuild. During the initialization phase, IKE utilizes a representative sub-sample set to construct a fixed mapping

function  $f(x) : \mathbb{R}^d \rightarrow \{0, 1\}^{t \times \psi}$  based on the iTree structure. Once this mapping function is determined, it functions as a static, data-dependent hashing mechanism. For newly arriving data points, the system only needs to perform a single forward pass to convert the high-dimensional vectors into binary codes using the existing iTrees. These codes are then directly inserted into the existing indexing structures (e.g., HNSW or inverted lists). This property ensures that IKE can seamlessly scale to expanding databases with minimal computational overhead.

A common challenge in industrial-scale retrieval is the evolution of document corpora over time. While the mapping function remains effective as long as the underlying data distribution is relatively stable, a significant distribution shift may necessitate an update to the model to maintain optimal retrieval precision. Unlike learning-based quantization or hashing methods that require expensive retraining on massive datasets, IKE maintains a distinct advantage due to its learning-free nature and extremely low reconstruction cost. In the event of a detected distribution shift, IKE can rapidly reconstruct the mapping function through efficient re-sampling.

To quantify this efficiency, we conducted benchmarks on the Istella22 dataset (1 million data points) using the LLM2Vec model ( $t = 4096, \psi = 16$ ). On the Linux server reported in the Section 4, a complete round of sample re-sampling and model reconstruction takes only 0.045 s. Furthermore, remapping the entire 1-million-sample corpus with the updated model is completed within 3.21 s. Such high responsiveness makes IKE a highly cost-effective solution for dynamic, industrial-scale retrieval scenarios compared to traditional learning-based approaches.

## E Experiment Details

In this section, we provide the details of the datasets and hyperparameter tuning procedure employed in our study.

### E.1 Dataset Details

The Massive Text Embedding Benchmark (MTEB) (Muennighoff et al., 2023) is a widely used benchmark for validating LLMs’ ability to extract text embeddings. We exploit 3 retrieval datasets from the MTEB (eng, v2) benchmark: FiQA2018 (Thakur et al.,

2021), FEVER-HN (Thorne et al., 2018), and Touche2020.V3 (Thakur et al., 2024). To assess scalability with large corpora, we also include an additional dataset from the MTEB (eng, v1) benchmark: HotpotQA (Yang et al., 2018), which has a corpus size of 5.2 million. To further validate the effectiveness of our method on multi-level retrieval datasets, we use the Istella22 (Dato et al., 2022) and TREC DL 23 (Craswell et al., 2023) datasets (with a 1 million subset extracted from their original corpora). All datasets provide human-annotated relevance labels for each query. In multi-level datasets like Istella22 and TREC DL, documents are rated on a graded relevance scale, ranging from 1 to 5 in Istella22 and from 1 to 4 in TREC DL. Detailed dataset statistics are presented in Table 4.

Table 4: **Dataset Statistics.** Datasets are categorized by source. “HN” in the dataset name indicates “Hard Negatives”. Relevance indicates the relationship between query and document: binary score (relevant, non-relevant) or multi-level score, in which each document is assigned the level of relevance w.r.t the query.

Dataset	Relevance	Corpus Size	Query Size
<i>MTEB</i>			
HotpotQA	Binary	5,233,329	7,405
FiQA2018	Binary	57,638	648
FEVER-HN	Binary	163,698	1,000
Touche2020.V3	3-Level	303,732	49
<i>Extra Datasets</i>			
Istella22 (subset)	5-Level	1,000,000	234
TREC DL 23 (subset)	4-Level	1,000,000	82

### E.2 Hyperparameter Tuning

The parameter  $\psi$ , which controls the number of partitions, is determined for each dataset through a grid search on a validation set, aimed at optimizing retrieval accuracy under an exhaustive search setting. This process is efficient, requiring only a single grid search per dataset. Specifically, we tune over the range  $[2, 16]$  and select the best value.

Notably, the FEVER-HN and Touche2020.V3 datasets do not provide an official validation split. We construct their validation sets as follows:

- **FEVER-HN:** This dataset is a subset derived from the FEVER dataset (Thorne et al., 2018). We therefore adopt the official validation split of the original FEVER dataset as the validation set for FEVER-HN. The official test set of FEVER-HN is subsequently used for final evaluation.

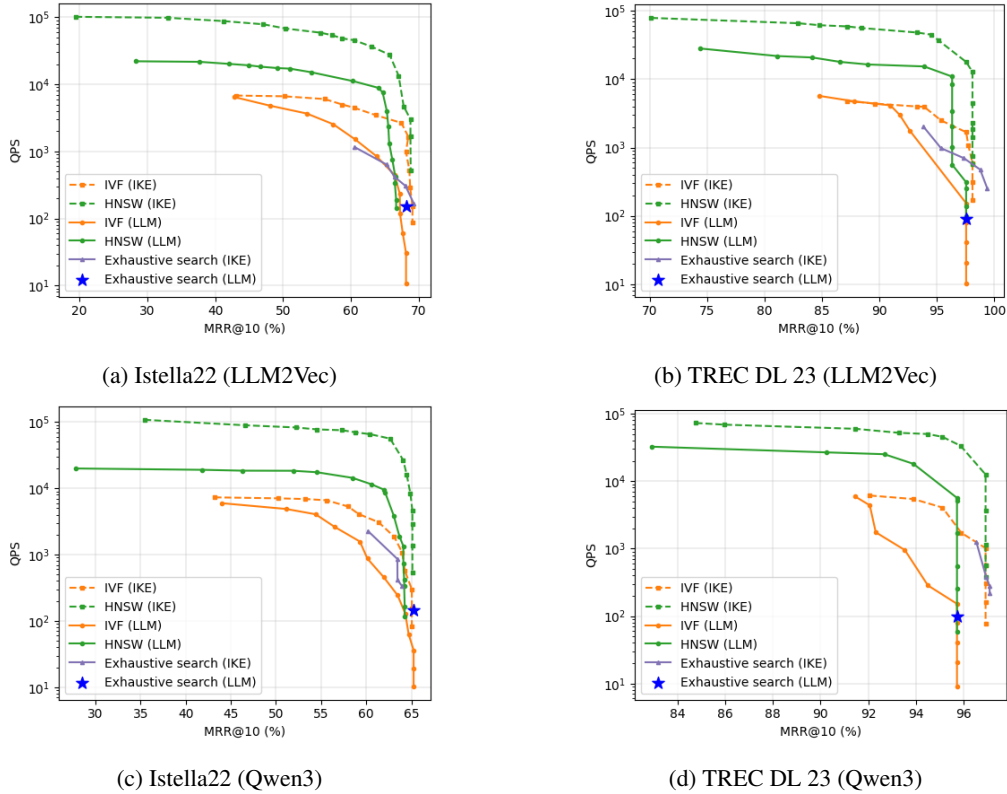


Figure 8: QPS vs. MRR@10 for Four ANN Methods on Istella22 and TREC DL 23.

- **Touche2020.V3**: We randomly sample 30% of the official test set to serve as the validation set, using the remaining 70% for final testing.

## F Extension of section 4.2

### F.1 Implementation details of combining two ANN methods with IKE

**IVF (IKE)** We use  $k$ -means clustering to partition the data and construct the IVF index in the LLM embedding space. During retrieval, the query’s LLM embedding is used to identify the closest clusters, and then the similarity between the query and the data points within these clusters is computed in the IKE space.

**HNSW (IKE)** The index construction and search processes follow the methodology proposed in the original paper, with the only modification being the replacement of the cosine similarity defined in the LLM embedding space with the similarity defined in the IKE space.

### F.2 Detailed Experimental Configuration

For the two ANN algorithms with IKE, we retain the same value of  $\psi$  in the exhaustive search setting for each dataset. For the IVF index, the

number of clusters  $nlist$  is set to 4,096, following the recommendation in Faiss (Douce et al., 2025) for million-scale datasets. The hyperparameters for the HNSW index are set as follows:  $M = 32$  and  $efConstruction = 500$ . In each plot, dotted lines represent retrieval performed in the IKE space, whereas solid lines represent retrieval in the LLM embedding space. The same color indicates the same underlying index structure for direct comparison. Additionally, the blue pentagram markers denote the retrieval performance achieved through exhaustive search in the LLM embedding space. The purple curve corresponds to the retrieval performance of IKE under the exhaustive search setting, which was obtained by tuning the hyperparameter  $t$  while keeping the parameter  $\psi$  fixed.

### F.3 Additional experimental results in the ANN search setting

Additional experimental results for the ANN search setting are provided in Figure 8, which shows the QPS-MRR@10 trade-off curves on the Istella22 and TREC DL 23 datasets. Consistent with the results presented in the main paper, performing retrieval with IKE using the HNSW index achieves multiple times higher throughput com-

pared to retrieval directly with LLM embeddings. For instance, on the Istella22 (Qwen3) dataset under the around 60% MRR@10 condition, QPS improves from  $1.1 \times 10^4$  (LLM) to  $6.5 \times 10^4$  (IKE), representing an nearly  $6\times$  acceleration in retrieval speed. Additionally, on some datasets, the peak performance on IKE exceeds that of the LLM, which aligns with the findings presented in Table 1. These results demonstrate that retrieval in the IKE space provides substantial efficiency gains while achieving comparable or even better accuracy.

## G Extension of section 4.3

Figure 9 presents the additional experimental results in the exhaustive search setting between CSR and IKE on the Istella22 (LLM2Vec) and TREC DL 23 (LLM2Vec) datasets. The results show that IKE significantly outperforms CSR in terms of search efficiency across all datasets, while achieving comparable or even superior retrieval accuracy.

## H Extension of section 4.4

### H.1 Description of Compared ANN methods

The compared ANN methods are described in detail as follows:

- **IVF (IKE) / HNSW (IKE):** This refers to two primary ANN methods, operated in the IKE space. The detailed implementations of these methods are the same as those in the previous section.
- **IVF (rpLSH) / HNSW (rpLSH):** rpLSH is a short name for Random-Projection-based Locality Sensitive Hashing (Charikar, 2002), which is a simple yet powerful LSH algorithm. Previous research has demonstrated its superior effectiveness, ranking it first among eleven popular hashing algorithms (Cai, 2016). It should be noted that Faiss only provides an exhaustive search implementation of rpLSH. To ensure a fair comparison, we integrate rpLSH with the IVF and HNSW indexing structures to perform an ANN search.
- **IVF (PQfs):** Product Quantization (PQ) (Jégou et al., 2011) is widely adopted in practice due to its effective balance between retrieval accuracy and compression capability (Subramanya et al., 2019; Johnson et al., 2019;


Zhang et al., 2025a; Qiu et al., 2022). However, conventional PQ implementations can suffer from slower search speeds due to table lookup operations. A recent approach, Product Quantization Fast Scan (PQfs) (André et al., 2015), utilizes a SIMD-based fast implementation that significantly accelerates the process, achieving a 4 to 6 times speedup while maintaining the same level of accuracy as standard PQ. IVF (PQfs) integrates PQfs with the IVF indexing structure. We utilize the Faiss implementation of IVF (PQfs) as a baseline for comparison.

- **ScaNN:** ScaNN (Guo et al., 2020), developed by Google, employs a novel anisotropic quantization loss function for constructing the quantization codebook. We use the official ScaNN implementation in its ANN search mode for evaluation.


### H.2 Additional experimental results in the ANN search setting

Figure 10 illustrates the QPS versus nDCG@10 trade-off curves for all methods on the HotpotQA (Qwen3), Istella22, and TREC DL 23 datasets, with the code lengths used indicated above each plot. In terms of search efficiency, HNSW (IKE) achieves up to  $10\times$  higher throughput than other learning-free compression methods in ANN setting, while maintaining comparable retrieval accuracy. For instance, on the Istella22 (Qwen3) dataset with a 512-byte code length, HNSW (IKE) improves throughput from  $1.1 \times 10^4$  to  $1.5 \times 10^5$  QPS at around 51% nDCG@10 compared to ScaNN, resulting in an order-of-magnitude improvement. Furthermore, the IVF (IKE) method attains search efficiency on par with the IVF (PQfs) approach.

Regarding accuracy, all methods achieve similar top-line performance across most datasets. Notably, our method attains significantly higher retrieval accuracy on the Istella22 (Qwen3) dataset, as shown in Figure 10c, where it outperforms other methods by 3% at a code length of 512 bytes. It is also noteworthy that the rpLSH method achieves competitive accuracy on several datasets. However, its retrieval speed is substantially lower than that of our approach, primarily due to its embedding mapping process involving high-dimensional matrix multiplications with high computational overhead. As the code length increases, the QPS gap between HNSW (rpLSH) and HNSW (IKE) gradually be-




(a) Istella22 (LLM2Vec)




(b) TREC DL 23 (LLM2Vec)

Figure 9: **Comparison of Retrieval Accuracy and Search Time between CSR and IKE in the exhaustive search setting on the Istella22 (LLM2Vec) and TREC DL 23 (LLM2Vec) datasets.** The search time includes both the mapping time for all queries and the similarity computation time, rather than solely the similarity calculation time as reported in the original CSR paper.


comes large. This occurs because the mapping time begins to dominate the total retrieval process of rpLSH, thereby diminishing the efficiency advantage typically offered by the HNSW index.




(a) HotpotQA (Qwen3)




(b) Istella22 (LLM2Vec)



(c) Istella22 (Qwen3)



(d) TREC DL 23 (LLM2Vec)



(e) TREC DL 23 (Qwen3)

Figure 10: Performance Comparison of Compression Methods on ANN Retrieval: QPS vs. nDCG@10 on HotpotQA (Qwen3), Istella22, and TREC DL 23. The curves are obtained by adjusting the search parameter of the index structures.