



# I/O Optimizations for Graph-Based Disk-Resident Approximate Nearest Neighbor Search: A Design Space Exploration

Liang Li  
China Telecom Cloud Computing  
Research Institute  
lil225@chinatelecom.cn

Shufeng Gong  
Northeastern University  
gongsf@mail.neu.edu.cn

Yanan Yang\*  
China Telecom Cloud Computing  
Research Institute  
yangyn11@chinatelecom.cn

Yiduo Wang  
China Telecom Cloud Computing  
Research Institute  
wangyd22@chinatelecom.cn

Jie Wu  
China Telecom Cloud Computing  
Research Institute  
Temple University  
jiewu@temple.edu

## ABSTRACT

Approximate nearest neighbor (ANN) search on SSD-backed indexes is increasingly I/O-bound (I/O accounts for 70–90% of query latency). We present an I/O-first framework for disk-based ANN that organizes techniques along three dimensions: memory layout, disk layout, and search algorithm. We introduce a page-level complexity model that explains how page locality and path length jointly determine page reads, and we validate the model empirically. Using consistent implementations across four public datasets, we quantify both single-factor effects and cross-dimensional synergies. We find that (i) memory-resident navigation and dynamic width provide the strongest standalone gains; (ii) page shuffle and page search are weak alone but complementary together; and (iii) a principled composition, OctopusANN, substantially reduces I/O and achieves 4.1–37.9% higher throughput than the state-of-the-art system Starling and 87.5–149.5% higher throughput than DiskANN at matched Recall@10=90%. Finally, we distill actionable guidelines for selecting storage-centric or hybrid designs across diverse concurrency levels and accuracy constraints, advocating systematic composition rather than isolated tweaks when pushing the performance frontier of disk-based ANN.

## PVLDB Reference Format:

Liang Li, Shufeng Gong, Yanan Yang, Yiduo Wang, and Jie Wu. I/O Optimizations for Graph-Based Disk-Resident Approximate Nearest Neighbor Search: A Design Space Exploration. PVLDB, 19(7): 1484 - 1498, 2026.

[doi:10.14778/3801059.3801064](https://doi.org/10.14778/3801059.3801064)

## PVLDB Artifact Availability:

The source code, data, and/or other artifacts have been made available at <https://github.com/LeonLee666/IObench4DiskANN>.

\*Corresponding author: Yanan Yang.

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](mailto:info@vldb.org). Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.

Proceedings of the VLDB Endowment, Vol. 19, No. 7 ISSN 2150-8097.  
[doi:10.14778/3801059.3801064](https://doi.org/10.14778/3801059.3801064)

## 1 INTRODUCTION

Approximate nearest neighbor (ANN) search is fundamental to modern retrieval-intensive applications [2, 3, 21, 57, 87], including information retrieval [18, 42], pattern recognition [15, 47], data mining [33, 37], machine learning [10, 14], recommendation systems [59, 71], retrieval-augmented generation (RAG) [49], and vector databases [31, 87, 90]. Among various approaches—hashing-based [27, 34], tree-based [3, 72], and quantization-based [23, 44, 66]—graph-based algorithms have emerged as state of the art for search quality and efficiency [7, 89, 94]. Representative methods such as HNSW [57], NSG [21], SSG [20], and Vamana [77] achieve exceptional performance with high recall and low latency. However, their substantial memory overhead renders them infeasible for billion-scale datasets on commodity hardware [13, 77, 88].

As datasets scale to billions of vectors, the prohibitive cost of DRAM necessitates a shift toward disk-resident solutions for vector search [26, 73, 77]. While in-memory indices have been extensively studied [7, 53, 89, 94], disk-resident graph indices face fundamentally different challenges driven by the I/O bottleneck. Unlike in-memory search, which is compute-bound, disk-based graph traversal becomes I/O-bound due to **high I/O latency** and **random I/O access patterns**: each hop may trigger fine-grained random reads whose per-request latency dominates distance computation. This fundamental shift leads to severe performance degradation and an optimization landscape that differs markedly from that of in-memory systems. Consequently, this work focuses exclusively on improving the I/O efficiency of disk-based graph indices.

Among existing disk-based solutions, DiskANN [77] has established itself as the *de facto* industrial standard [74, 76], widely adopted by major commercial vector databases including Microsoft Azure [62] and Huawei GaussDB [35], *etc.* It employs a three-layer architecture: a disk storage layer for the full graph and vectors, a memory layer for compressed navigation vectors, and a cache for high-degree nodes. Despite its robust design, our analysis reveals that I/O operations still account for 70–90% of total query latency. This bottleneck stems from four critical inefficiencies: (1) *poor data locality*, where scattered vertices cause massive I/O waste [88]; (2) *low bandwidth utilization* due to strictly sequential I/O-compute execution [29, 92]; (3) *long search paths* that require excessive disk

round-trips [29, 77, 88]; and (4) *frequent cache misses* resulting from unpredictable random access patterns [39, 40, 77, 92].

There has been extensive work addressing these issues, including caching strategies [39, 77, 92], quantization methods [77, 97], hierarchical structures [88, 97], and I/O pipelining approaches [29]. Although several studies [7, 53, 89, 94] have examined graph-based indexes, no survey systematically explores I/O optimizations. This gap motivates us to conduct an in-depth investigation of I/O efficiency. This paper explores I/O optimization methods for on-disk graph-based ANN indexes and focuses on answering the following three questions.

**Q1: How to classify I/O efficiency optimizations systematically to highlight the differences and connections among various methods?** Existing optimization techniques are studied in isolation without a unified classification framework. Although numerous approaches have been proposed—vector quantization [77, 97], hierarchical structures [13, 88], I/O pipelining [29, 92], and reorganization of data layout [80, 88]—they target different I/O bottlenecks without a systematic understanding of their orthogonal boundaries, making it difficult to navigate the design space.

**Q2: What is the effectiveness of each individual optimization and its contribution to I/O performance?** Some optimization methods, such as pipeline and dynamic width [29], have been shown to work synergistically, significantly improving I/O performance. However, the effectiveness of each of these optimizations when applied independently remains unclear.

**Q3: Do combinations of individual optimizations yield multiplicative performance improvements?** While state-of-the-art solutions such as Starling [88] and PipeANN [29] represent significant advances, systematic investigation into whether multidimensional optimization combinations can achieve breakthrough performance improvements remains lacking. The potential for carefully orchestrated combinations to deliver synergistic performance gains is largely unexplored.

To address these questions, this paper presents a systematic study of the I/O optimization design space for disk-based ANN search algorithms. We introduce a comprehensive three-dimensional taxonomy that organizes optimization techniques into orthogonal categories: (i) *memory layout optimization*, (ii) *disk layout optimization*, and (iii) *search algorithm optimization*. Using this framework and consistent experimental methodology across four public datasets, we conduct controlled evaluations that quantify both individual optimization impacts and reveal synergistic combinations. Our key insight is that carefully orchestrated multi-dimensional approaches yield synergistic improvements by jointly increasing page-level utility and shortening convergence paths.

Our work makes three key contributions:

**(1) Systematic taxonomy framework.** We establish a comprehensive taxonomy for I/O optimization algorithms, identifying three orthogonal dimensions that unify previously fragmented approaches and provide an objective function for the navigation optimization landscape.

**(2) Optimization evaluation and empirical findings.** We conduct a controlled, apples-to-apples evaluation across four massive

datasets to quantify the effectiveness and costs of individual I/O optimizations, and to extract a set of empirical findings that characterize when each technique is beneficial, marginal, or counterproductive.

**(3) Combination case study and actionable guidelines.** Building on the above findings, we further study how orthogonal optimizations can be composed in practice, and present *OctopusANN* as a representative effective combination. It substantially improves upon existing methods—achieving 4.1–37.9% higher throughput than Starling and 87.5–149.5% over DiskANN at 90% recall—while our findings provide evidence-based guidelines for selecting and combining techniques under different workload and resource regimes.

## 2 BACKGROUND

### 2.1 Nearest Neighbor Search

**Approximate Nearest Neighbor Search.** Given a dataset  $D \subset \mathbb{R}^d$  containing  $n$  points, the  $k$ -nearest neighbor ( $k$ -NN) search problem seeks the  $k$  closest points  $S^*$  to query point  $q \in \mathbb{R}^d$ , i.e.,  $p \in D \setminus S^*$ ,  $p' \in S^*$ ,  $\|p, q\| \geq \|p', q\|$ , where  $\|\cdot\|$  is a distance function, e.g., Euclidean or Cosine distance. While exact and robust search requires  $O(nd)$  time complexity [53], approximate nearest neighbor search (ANN) trades accuracy for efficiency by returning approximate near neighbors  $S$  at lower computational cost. The query accuracy of ANN methods is measured by the *recall* value,  $\text{Recall}@k = \frac{|S \cap S^*|}{k}$ .

**Graph-Based ANN methods.** Among various ANN search approaches, proximity-graph methods [58, 70] are widely used. Prior studies [4, 50, 53, 89] report strong performance for graph-based methods. Graph-based ANN search constructs graphs where vertices represent vectors and edges connect similar vectors. Main approaches include Delaunay graphs [19] (good geometric properties),  $k$ -nearest neighbor graphs [67] (predictable degree), relative neighborhood graphs [85] (sparse connectivity), minimum spanning trees [48] (minimal edges), and small-world graphs [56] (efficient for search). Modern algorithms such as HNSW [57], NSG [21], SSG [20], NGT [41], SPTAG [60], Vamana [77],  $\tau$ -MG [68], and ELPIS [6] have shown strong performance in practice.

**Best-first search.** Graph-based ANN search commonly employs best-first search [32] or beam search [82], maintaining a priority queue of candidate nodes sorted by distance to the query. The search iteratively selects the closest unvisited node (beam search is a width-limited variant in which each iteration selects the closest top- $\omega$  unvisited nodes, where  $\omega$  is the *beam width*), explores neighbors, and adds new candidates until identifying the required nearest neighbors. Search efficiency depends on entry point quality [7, 89], graph structure quality [9, 94], distance computation [1, 12, 16, 22, 91], and early termination strategies [51].

**Limitations of graph-based ANN methods.** While in-memory graph-based methods achieve excellent search performance, they face significant memory limitations with large-scale (e.g., billion-scale) datasets. Memory overhead increases substantially with dataset size—for example, HNSW on a billion-scale dataset in 128 dimensions can consume hundreds of gigabytes of memory (e.g., 800 GB under common settings), often exceeding typical workstation RAM capacity. When dataset sizes exceed available memory,




Figure 1: DiskANN data layout. Each record is page-aligned and packed into 4 KB pages.

traditional in-memory approaches become infeasible, prompting disk-based ANN search algorithms.

## 2.2 Disk-Based ANN

To address the memory constraints that necessitate disk-resident search, we adopt DiskANN as a canonical model to illustrate how a three-layer architecture organizes on-disk vectors and graph traversal, setting the stage for our I/O-centric analysis in §3.

**DiskANN overview and architecture.** DiskANN [26, 38, 73, 77] is a *de facto* standard disk-based ANN system adopted by products such as Milvus [100], Azure Database [62], Timescale [84], and GaussDB [35], and serves as a standard baseline in public benchmarks [74, 76]. As shown in Figure 1, DiskANN follows a three-layer architecture: (1) a **disk storage layer** that stores full-precision vectors and neighbor lists in page-aligned records; (2) a **memory quantization layer** (e.g., Product Quantization, PQ) that keeps compressed codes for fast, memory-resident filtering; and (3) a **cache management layer** that retains hot vertices. This design implements an “approximate guidance, precise refinement” workflow: memory-resident PQ codes guide candidate selection, while disk-resident full vectors enable final re-ranking.

**Representative disk-based systems.** Beyond DiskANN, representative systems can be categorized into three orthogonal routes. *Starling/DiskANN++* [63, 88] are layout-centered, employing an in-memory navigation graph and locality-aware page shuffling (optionally in-page search) to raise per-page utility and shorten search paths. *PipeANN* [29] is algorithm-centric, using I/O-compute pipelining with dynamic width; it may trigger speculative reads under concurrency. *AiSAQ/LM-DiskANN* [65, 80] are storage-centric, placing PQ codes on SSD to minimize memory usage but risking read amplification and a larger on-disk footprint. Other variants include *GRIP/Zoom* [96, 97] (memory-first: PQ with hierarchical navigation/MemGraph under moderate memory budgets) and *SPANN/BBANN* [13, 99] (partitioned/hierarchical navigators with storage-aware locality). This organization directly impacts per-page utility and traversal hops, which we analyze through an I/O-centric lens in §3. We use DiskANN as a running reference to anchor later analysis of I/O efficiency challenges and to provide a consistent baseline for design-space comparisons.

## 2.3 I/O is the Main Bottleneck

Although disk-based graph indices address the memory bottleneck of in-memory indices, they still face significant I/O efficiency




Figure 2: Latency breakdown across four datasets.

challenges [29, 88]. Figure 2 shows the percentage breakdown of compute and I/O latency for DiskANN queries across four datasets. The figure reveals that I/O latency dominates query latency, accounting for approximately 70%–90% of the total (note that dataset descriptions are provided in §5; the GIST dataset has higher dimensionality and requires more search hops). This dominance stems from the fundamental nature of graph-based vector search: (1) **High I/O latency**, where fetching nodes from disk is orders of magnitude slower than in-memory distance computation; (2) **Random I/O access patterns**, as graph traversal inherently induces fine-grained random I/Os that fail to exploit the internal parallelism and throughput of modern SSDs.

## 3 I/O CHALLENGES ANALYSIS

### 3.1 I/O Complexity for Best-First Search

To establish a formal foundation for analyzing I/O efficiency, we first introduce key terminology specific to graph-structured disk-resident indices, then present the I/O complexity model.

**Overlap Ratio.** To ground the I/O complexity analysis, we first formalize page-level utility via an overlap ratio. Specifically, we define a local, page-utility-oriented overlap ratio at the vertex level and then aggregate it to obtain the global  $OR(G)$ . The overlap ratio for individual vertex  $u$  quantifies the proportion of its graph neighbors that reside on the same disk page:  $OR(u) = \frac{|B(u) \cap N(u)|}{n_p - 1}$ , where  $B(u)$  represents the set of vertices stored in the same disk page as vertex  $u$ , and  $N(u)$  denotes the set of graph neighbors of vertex  $u$ . The denominator  $n_p - 1$  excludes vertex  $u$  itself, focusing on the co-location of its neighbors. For example, if a page stores  $n_p=16$  records and  $|B(u) \cap N(u)|=2$ , then  $OR(u) = 2/(16-1) = 13.3\%$ . For a graph index  $G = (V, E)$ , aggregating local overlap ratios gives the global overlap ratio as the vertex-wise average:  $OR(G) = \frac{\sum_{u \in V} OR(u)}{|V|}$ .

**I/O complexity model.** We adopt the disk-resident best-first search model [88] to illustrate the I/O challenge. Let the disk-resident graph index be  $G = (V, E)$ . Let  $H$  denote the expected hop count per query,  $\bar{R}$  the average out-degree, and  $n_p$  the number of records per page. Under a uniform page size  $P$  (e.g., 4 KB) and record size  $s_{rec}$ , we have  $n_p = \lfloor P/s_{rec} \rfloor$ . Based on the terminology and overlap ratio defined above, the page I/O complexity for disk-resident best-first search satisfies<sup>1</sup>

$$\text{Page reads per query} = O\left(\frac{\bar{R} \cdot H}{OR(G) \cdot n_p}\right) \quad (1)$$

<sup>1</sup>This model does not account for optimization techniques such as quantization, caching, and MemGraph, etc.

**Intuition.** A query performs  $O(\bar{R} \cdot H)$  effective neighbor expansions; each page read contributes an expected  $OR(G) \cdot n_p$  useful neighbors, so the number of unique page reads scales as  $O(\bar{R} \cdot H / (OR(G) \cdot n_p))$ . Intuitively, shorter convergence paths ( $H \downarrow$ ) and higher page-level utility ( $OR(G) \uparrow$ ,  $n_p \uparrow$ ) reduce I/Os, while larger  $\bar{R}$  expands the frontier and increases I/Os.

### 3.2 Objectives

Grounded in the theoretical model above and the profiling results in Figure 2, we systematically identify four root causes of I/O inefficiency: (1) *I/O waste* occurs when graph traversal reads disk pages containing vertices that are never explored during best-first search, resulting in substantial wasted I/O operations [88]; (2) *low bandwidth utilization* arises from sequential I/O–compute execution, where disk I/O remains idle during computation phases, leaving bandwidth underutilized [29, 92]; (3) *long search paths* emerge when entry points are far from query targets, requiring extensive graph traversal with numerous hops, each demanding independent disk access and dramatically increasing I/O operations [29, 77, 88]; and (4) *frequent cache misses* occur during best-first search due to unpredictable access patterns, necessitating effective cache strategies to reduce miss rates [39, 40, 77, 92].

Through a systematic survey of existing work [11, 13, 23, 29, 39, 40, 44, 45, 63, 65, 69, 77, 78, 80, 83, 88, 92, 96, 97, 99], we organize mainstream I/O optimization techniques into an orthogonal three-dimensional taxonomy: **Memory Layout Optimization** (§4.1), **Disk Layout Optimization** (§4.2), and **Search Algorithm Optimization** (§4.3), encompassing eight representative techniques. Table 1 maps each technique to the four bottlenecks above, revealing how different optimizations address specific I/O challenges. §4 provides detailed descriptions of each technique.

**Scope.** We target general-purpose SSDs and fix DiskANN’s Vamana-based logical graph structure; our study focuses on physical data layout and search scheduling under realistic memory budgets. Beyond traditional SSD-based settings, specialized devices—e.g., SmartSSD [45, 78, 83], Persistent Memory [11, 69], and CXL [39, 40, 46, 79]—are out of scope. While alternative logical graph structures can influence performance [94], this work confines its scope to the physical graph structure under DiskANN’s Vamana-based logic.

## 4 DESIGN SPACE


### 4.1 Memory Layout Optimization

Memory layout optimization strategies reduce the frequency of disk I/O by maintaining carefully designed auxiliary data structures in memory. This dimension encompasses three approaches: (1) *vector quantization* for compressing high-dimensional data [23, 44, 77, 96, 97]; (2) *cache management* for keeping frequently accessed vertices in memory [39, 40, 77]; and (3) *hierarchical graphs* for reducing search path length through coarse-grained navigation [13, 63, 99].

**4.1.1 Product Quantization (PQ).** Product quantization [24, 25, 44] is a critical technique for achieving memory efficiency, aiming to reduce memory footprint and disk I/O by compressing high-dimensional vector representations. As shown in Figure 3, the core

**Table 1: Eight techniques versus four I/O challenges (I/O waste, bandwidth utilization, long search paths, high cache misses).  $\uparrow$ : Improve;  $\downarrow$ : cost;  $-/-$ : Reduce.**

Method	I/O waste	Bandwidth utilization	Long search path	High cache misses
PQ	$\uparrow$	$-/-$	$-/-$	$-/-$
Cache	$-/-$	$-/-$	$-/-$	$\uparrow$
MemGraph	$-/-$	$-/-$	$\uparrow$	$-/-$
PageShuffle	$\uparrow$	$-/-$	$-/-$	$-/-$
All-in-Storage	$\downarrow$	$-/-$	$-/-$	$-/-$
DynamicWidth	$\uparrow$	$-/-$	$\uparrow$	$-/-$
Pipeline	$\downarrow$	$\uparrow$	$-/-$	$-/-$
PageSearch	$\uparrow$	$\downarrow$	$-/-$	$-/-$



**Figure 3: Quantization reduces disk reads by moving most distance calculation to memory.**

benefit of quantization lies in maintaining compressed vector representations in memory (with compression rates of tens of times) while storing full-precision vectors on disk. This enables a two-stage search process: ① *approximate filtering* using memory-resident quantized vectors to identify promising candidates, followed by ② *precise refinement* with disk-based full vectors only for top candidates. The key insight is that distance calculations can be performed directly using quantized coordinates in memory, eliminating the need to read full vectors from disk for initial candidate ranking. In the illustrated example (Figure 3), traditional DiskANN requires 7 disk I/O operations for vector distance computations, while the quantization-optimized approach reduces this to only 3 I/O operations primarily for neighbor topology retrieval, achieving significant I/O efficiency improvement. Formally, with PQ optimization, the I/O complexity from Equation 1 is reduced to:

$$\text{Optimized page reads} = O\left(\frac{H}{OR(G) \cdot n_p}\right) \quad (2)$$

This improvement eliminates the  $\bar{R}$  factor in the numerator because distance computations are performed in memory and no longer require disk reads. Theoretically, the lower bound for page reads is  $O(H/n_p)$  (i.e., all required vectors are co-located on the same page). However, this ideal page locality is difficult to attain in practice due to the inherent randomness of graph traversals and the complexity of optimizing physical layouts for dynamic query patterns.

However, vector quantization introduces precision loss, creating a trade-off between memory efficiency and search accuracy. To mitigate this limitation, typical approaches expand the candidate set during search and perform precise re-ranking using full vectors from disk for final Top-k results. DiskANN’s co-location design stores vector data and neighbor information together, ensuring that reading neighbor IDs for candidate expansion simultaneously retrieves associated vector data, avoiding additional disk I/Os for the re-ranking phase.






Figure 7: AiSAQ layout. Each record stores the node vector and neighbor IDs, plus PQ codes of neighbors.




Figure 8: Dynamic search width. The beam width increases progressively across the approach and converge phases.

Product Quantization (AiSAQ) [80] exemplifies this design by moving full-precision vectors and PQ codes from memory to storage. As shown in Figure 7, AiSAQ co-locates vector data, neighbor IDs, and the PQ codes of neighboring vertices within the same storage units. This design reduces memory requirements from gigabytes to merely megabytes for billion-scale datasets while maintaining competitive search performance. The approach offers significant advantages for memory-constrained environments and enables rapid switching between multiple large-scale indexes—a crucial capability for applications like retrieval-augmented generation (RAG) that require dynamic switching between multiple indices [43]. However, we do not include AiSAQ in our experimental evaluation, as it targets memory-constrained scenarios with fundamentally different design trade-offs. Comparing AiSAQ against methods that utilize more substantial memory budgets (as in our evaluation setup) would be unfair, since AiSAQ’s primary optimization goal—minimizing memory footprint—conflicts with our focus on I/O efficiency under realistic memory budgets.

### 4.3 Search Algorithm Optimization

Beyond optimizing data layout, search algorithm optimization focuses on improving the execution efficiency of graph traversal itself. This dimension addresses two critical algorithmic bottlenecks: (1) *vertex utilization inefficiency* when accessing disk pages, and (2) *bandwidth underutilization* caused by sequential I/O–compute execution. These strategies complement layout improvements by maximizing the effectiveness of each disk access operation.

**4.3.1 Dynamic Width (DW).** The best-first search or beam search process can be naturally divided into two distinct phases: the *approach phase* and the *converge phase*. As illustrated in Figure 8, during the ① *approach phase*, the search rapidly navigates toward the vicinity of the target query vector. In the ② *converge phase*, the search gradually converges to the target query vector. This two-phase phenomenon has been discussed extensively in multiple studies [29, 95, 98]. PipeANN [29] leverages the progressive convergence characteristics of beam search by adopting a smaller search width during the *approach phase* to avoid unnecessary I/O operations, since larger search widths do not improve recall effectiveness in this phase. Subsequently, the search width is gradually increased




Figure 9: Pipelined processing: (a) overlap compute and I/O to improve utilization; (b) continuous I/O pipelining to maximize bandwidth efficiency.

during the *converge phase* to accelerate the convergence process and improve search efficiency. To make this intuition precise, we define the *I/O utilization* as

$$U_{io} \triangleq \frac{N_{\text{eff}}}{N_{\text{read}}} = \frac{N_{\text{eff}}}{N_{\text{eff}} + N_{\text{rbu}}}, \quad (3)$$

where  $N_{\text{read}}$  is the number of nodes retrieved from SSD,  $N_{\text{eff}}$  is the number of retrieved nodes that are actually explored (expanded) by the search, and  $N_{\text{rbu}}$  denotes “read-but-unexplored” nodes. Increasing beam width  $\omega$  tends to inflate  $N_{\text{rbu}}$  (especially in the approach phase where extra reads rarely improve recall), thus lowering  $U_{io}$  and wasting I/O. Dynamic width mitigates this by keeping  $\omega$  small early to suppress wasted reads, and gradually increasing  $W$  later when more retrieved candidates are likely to be explored, improving both I/O utilization and convergence.

**4.3.2 Pipeline Search (Pipeline).** Traditional disk-based graph search suffers from severe bandwidth underutilization due to sequential execution patterns. As illustrated in Figure 9a, beam search’s alternating I/O and computation phases leave storage devices idle during processing, resulting in actual bandwidth utilization far below hardware capabilities.

Pipeline search addresses this limitation through two progressive optimizations. First, asynchronous I/O–compute overlap [92] enables simultaneous execution of distance calculations and disk reads, improving resource utilization. Second, continuous I/O pipelining [29] eliminates batch processing delays by immediately issuing new I/O requests upon completion of individual operations, as demonstrated in Figure 9b. This approach maintains near-continuous disk activity, achieving high bandwidth utilization and significantly reducing query latency.

**4.3.3 Page Search (PSe).** While pipeline search optimizes bandwidth utilization, page search addresses vertex utilization inefficiency within loaded disk pages. Given that distance computation overhead is orders of magnitude lower than disk I/O latency, page search exploits this asymmetry by computing distances to all vertices within each accessed page, not just the target neighbors.

This approach transforms the fundamental trade-off in disk-based search: instead of minimizing distance calculations, page search maximizes the utility of expensive disk accesses. When retrieving a neighbor during graph traversal, the algorithm evaluates all colocated vertices within the same page, potentially discovering closer candidates without additional I/O operations. This strategy effectively increases the number of search nodes without proportional I/O overhead, improving search quality while maintaining

Table 2: Coverage of representative systems across the three optimization dimensions.

Algorithms	Memory Layout Optimization			Disk Layout Optimization		Search Algorithm Optimization		
	PQ	Cache	MemGraph	PageShuffle	All-in-Storage	DynamicWidth	Pipeline	PageSearch
DiskANN [77]	✓	✓	✗	✗	✗	✗	✗	✗
Starling [88]/DiskANN++ [63]	✓	✗	✓	✓	✗	✗	✗	✓
AiSAQ [80]/LM-DiskANN [65]	✓	✗	✗	✗	✓	✗	✗	✗
PipeANN [29]	✓	✗	✓	✗	✗	✓	✓	✗
GRIP [97]/Zoom [96]	✓	✗	✓	✗	✗	✗	✗	✗
SPANN [13]/BBANN [99]	✗	✗	✓	✓	✗	✗	✗	✗

Table 3: Datasets and parameters used in the evaluation.  $B$  and  $M$  are memory budgets (GB).

Dataset	Dataset Statistics				Index-Build Param.				
	Dim	Type	#Vec	#Q	R	L	$\alpha$	B	M
SIFT [36]	128	uint8	100M	10K	64	125	1.2	4.0	50
DEEP [8]	96	float	100M	10K	64	125	1.2	4.0	50
SPACEV [61]	100	int8	100M	29K	64	125	1.2	3.0	40
GIST [36]	960	float	1M	10K	64	125	1.2	0.2	3

I/O efficiency. The key insight is that leveraging the full content of each disk page can enhance search effectiveness at minimal computational cost.

#### 4.4 Algorithmic Landscape Analysis

Grounded in our three-dimensional taxonomy, we position representative systems against the eight techniques summarized in Table 2. This mapping clarifies each system’s design center and expected I/O behaviors within the design space. In short, recent progress has moved from single-axis tweaks to multi-dimensional, synergy-driven designs that directly address the four I/O bottlenecks identified earlier.

### 5 EXPERIMENTAL TESTBED

#### 5.1 Experimental Settings

**Hardware.** We run experiments on a server with an Intel Xeon w7-3455 (24 cores), 128 GB DDR4 memory, and a 4 TB NVMe SSD, on Ubuntu 22.04. Using `fio` [5], we measure 4KB random read IOPS of **819 K** and random read bandwidth of **3,200 MB/s**, representative of high-performance SSDs. We also measure 16KB random read IOPS of **318 K** and random read bandwidth of **4,962 MB/s**.

**Datasets.** We evaluate on four widely adopted public datasets, SIFT [36], DEEP [8], SPACEV [61], and GIST [36], covering natural clustering, deep embeddings, and production-scale data. Dataset statistics and index build parameters are summarized in Table 3.

**Performance indicators.** We report search efficiency (queries per second, QPS; mean per-query latency), accuracy (Recall@10), space footprint (disk and memory), and I/O metrics (I/O operations per query, I/O operations per second, IOPS; bandwidth in MB/s). Unless otherwise stated, comparisons are made at matched Recall@10. We do not separately report Eq. 3, as it is strongly correlated with the number of I/O operations per query.

**Implementation and parameters.** We build on DiskANN [75] and integrate representative optimizations from Starling [81] and PipeANN [30]. As `io_uring` is a system-level optimization that all algorithms can be applied but currently only PipeANN use it. To ensure fairness, we disable `io_uring` [55] in PipeANN so that all




Figure 10: Performance comparison of SPANN and DiskANN on four datasets.

methods use the same asynchronous I/O engine (`libaio`). Note that all the I/O operations are performed in `DIRECT_IO` mode, *i.e.*, without using the file system page cache. We run with 48 workers; device utilization (IOPS and bandwidth) is sampled at 1 s intervals via `iostat`. For index construction, we adopt product quantization under a memory budget  $B$ , partition data via overlapping clustering, and build per-shard indexes under budget  $M$  using Vamana parameters: maximum out-degree  $R = 64$ , candidate set size  $L = 125$ , and pruning with  $\alpha = 1.2$ ; shards are then merged and stored with a 4 KB page size (we use 8 KB or 16 KB pages for GIST due to its larger per-record size). During query processing, we use a beam width  $\omega = 8$  with dynamic-width expansion up to  $\omega = 32$ . For MemGraph, we sample 0.1% vertices to construct an in-memory navigation graph ( $R = 48$ ,  $L = 128$ ) and employ a Single-Source Shortest Path (SSSP)-based cache with 0.1% records. Results are averaged over five independent runs. Dataset-specific budgets and settings are summarized in Table 3.

#### 5.2 Graph-based vs. Inverted Indices

Although this paper primarily focuses on evaluating graph-based indexing techniques, we first investigate whether graph indices consistently outperform inverted indices. We evaluated SPANN [13] following SPFresh [93] parameter settings; note that we enable the replica feature and set the replication factor to 8. As shown in Figure 10, for low-dimensional data, SPANN’s performance degrades much more rapidly than DiskANN’s as the recall target increases. This is because achieving higher recall requires SPANN to read coarse-grained posting lists, leading to more I/O operations and consequently faster performance degradation. Furthermore, on high-dimensional datasets, SPANN exhibits poor performance due to the curse of dimensionality, which significantly increases the difficulty of effective clustering. Consequently, SPANN outperforms DiskANN only on low-dimensional datasets under low recall requirements. More evaluation details could be found in our technique report [52]. Therefore, we decided to exclude SPANN from our subsequent detailed evaluation.




Figure 11: Recall-QPS trade-off: seven optimizations across four datasets.




Figure 12: Latency-Recall@10 trade-off across seven optimizations and four datasets.




Figure 13: I/O operations per query across four datasets (Recall@10 matched).

**Finding 1:** SPANN outperforms DiskANN only on low-dimensional datasets under low recall requirements. On high-dimensional datasets, SPANN exhibits poor performance due to the curse of dimensionality, which significantly increases the difficulty of effective clustering. Moreover, SPANN’s index size is 1.5× to 3.4× larger than DiskANN’s.

## 6 INDIVIDUAL OPTIMIZATION EVALUATION

In this section, we evaluate seven representative methods introduced in §4: Baseline (PQ), Cache, MemGraph, PageShuffle, DynamicWidth, Pipeline, and PageSearch.

### 6.1 Search Performance Analysis

On top of the PQ baseline, we ablate each of the six techniques to quantify their standalone effects. For completeness, these techniques span all three dimensions of our taxonomy: memory layout (Cache, MemGraph), disk layout (PageShuffle), and search algorithm optimization (Pipeline, DynamicWidth, PageSearch). We report Pareto fronts across multiple recall targets in Figures 11, 12, and 13, covering throughput (QPS), average latency, and I/O per query, respectively. Intuitively, Figures 12 and 13 exhibit highly consistent trends, reinforcing the latency breakdown in Figure 2. Because end-to-end latency scales with the number of I/O operations, improving I/O yields proportional latency reductions, validating our I/O-centric optimization focus.

**Finding 2:** I/O operations dominate query latency across datasets. Latency-recall and I/O-per-query curves align closely, indicating that improving I/O translates into proportional latency reductions in disk-based ANN systems.

Next, we group all those algorithms into three classes: (1) techniques that deliver obvious performance improvements; (2) techniques with negligible or marginal gains; and (3) techniques that are counterproductive when used alone.

**Effective single-factor optimizations.** Optimizations targeting memory layout and search-width adaptivity provide the strongest standalone gains across all four datasets. With SIFT, MemGraph reduces I/O by about 45% and increases QPS by about 90% at relaxed accuracy (e.g.,  $L = 10$  with  $\text{Recall}@10 \approx 0.7$ ); even at high accuracy (e.g.,  $L = 100$  with  $\text{Recall}@10 \geq 0.95$ ), it still reduces I/O by about 19% and boosts throughput by about 25%. Its overhead is modest: sampling 0.1% vertices for an in-memory navigation graph adds only ~30-50 MB (cf. Table 5). The Cache policy accelerates early-hop I/O by retaining nodes near entry points in memory. Because this SSSP-based caching is sensitive to graph quality, it delivers especially strong gains on SPACEV, as SPACEV’s graph index contains many low-degree nodes that concentrate early-hop paths near the entry points.

Finally, DynamicWidth improves I/O utilization by adapting the search width between the approach and converge phases: at relaxed

Table 4: Disk Metrics Comparison.

Method	SIFT		DEEP		SPACEV		GIST	
	IOPS	BW	IOPS	BW	IOPS	BW	IOPS	BW
	(K)	(MB/s)	(K)	(MB/s)	(K)	(MB/s)	(K)	(MB/s)
Baseline	791	3091	819	3200	768	3001	313	4889
Cache	808	3156	813	3177	815	3182	322	5029
MemGraph	818	3194	814	3181	813	3177	318	4973
PageShuffle	803	3136	803	3136	741	2895	309	4832
Pipeline	806	3197	808	3158	776	3030	323	5050
DynamicWidth	796	3109	783	3060	765	2987	315	4930
PageSearch	690	2698	796	3111	470	1834	315	4917

recall requirements, it reduces I/O operations by about 25% and raises QPS by about 50%. This adaptivity can, however, slightly reduce accuracy compared with a fixed width  $\omega$  configuration; e.g., on GIST with  $L = 10$ , recall decreases from 66% to 50%. On high-dimensional datasets with high recall requirements, DynamicWidth exhibits severe performance degradation. This occurs because the number of search iterations increases significantly, causing the search width to grow progressively larger and consequently triggering substantially more I/O operations. Therefore, in such scenarios, we do not recommend using DynamicWidth in isolation; instead, it should be combined with other techniques that reduce hop counts.

**Finding 3: The strongest single-factor gains come from memory layout and search width adaptivity.** MemGraph and Cache reduce I/O by providing better entry points that shorten disk traversal paths, while DynamicWidth reduces I/O by eliminating wasteful exploration in early search phases. These optimizations target the root causes of I/O inefficiency: poor starting positions and over-expansion during approach phases.

**Optimizations with limited standalone impact.** PageShuffle increases inter page neighbor locality, yet average I/O drops by only ~9% on SIFT. In beam search, most I/O is incurred by pages along the multi-hop expansion path; improving one-hop neighbor locality seldom alters that path and therefore does little to reduce page I/O operations.

PageSearch leaves I/O counts essentially unchanged in our setup but increases in-page computation by computing all records on each fetched page; For uint8 datasets such as SIFT and SPACEV, where each page contains more than 10 vector records, PageSearch introduces substantial per-page computational overhead, which sharply degrades SSD performance (cf. Table 4). The extra computation prolongs SSD idle time and reduces IOPS (cf. Fig. 9a). Prior work (e.g., Starling) reports strong gains when PageShuffle and PageSearch are combined; we evaluate this claim in the combination study below.

**Finding 4: PageShuffle and PageSearch underperform in isolation.** Locality gains and per-page computing do not translate into fewer I/Os; PageSearch may even introduce performance degradation on SIFT and SPACEV. Their value may emerge when used together.

**Counterintuitive degradations from advanced techniques.** As shown in Table 4, Pipeline fails to improve disk I/O because baseline algorithms already saturate disk performance under high concurrency. Instead, Pipeline widens the exploration frontier




Figure 14: Performance impacts of memory budget.

by issuing I/O requests before the best candidate is confirmed, leading to speculative reads. This increases the total number of I/O operations (cf. Figure 13), which paradoxically degrades end-to-end performance rather than improving it.

**Finding 5: Pipeline optimization can be counterproductive.** Pipeline widens the exploration frontier by speculative reading, which introduces more IO operations and reduces the overall performance.

## 6.2 Memory Budget Analysis

As we know, increasing memory budgets yields substantial performance improvements: larger MemGraph sampling graphs provide better performance gains, and allocating more memory to PQ enables higher-dimensional quantized vectors, thereby improving search accuracy. However, when memory headroom exists, the question of how to optimally allocate budgets between PQ and MemGraph warrants investigation.

Quantized coordinates require storing compressed representations for every vector in the dataset, achieving compression ratios of tens-to-one compared to the original data, yet the resulting memory footprint remains substantial. In contrast, MemGraph samples a small fraction of the original graph and stores only topological information (no coordinate data), resulting in memory overhead that is typically much smaller than PQ compressed coordinates.

Figure 14 compares performance under different PQ coordinate and MemGraph memory budgets on the SIFT dataset. Two key observations emerge. First, under identical search conditions, increasing PQ coordinate dimensions improves recall accuracy but has minimal impact on throughput. Second, increasing MemGraph sampling ratios (the figure shows ratios of 1/10,000, 1/1,000, and 1/100) substantially improves throughput but has little effect on recall. At high sampling ratios, further increasing the sampling frequency yields diminishing performance returns.

**Finding 6: When memory headroom exists, prioritize allocating a modest budget to MemGraph sampling, then incrementally increase PQ coordinate memory to achieve optimal performance.**

## 6.3 Graph Index Construction Overhead

As shown in Table 5, we evaluate the index construction overhead, including build time, peak memory consumption during construction, and index structure size (covering both on-disk index

**Table 5: Index Construction Overhead. Columns: BT = total build time (seconds), MaxM = peak memory consumption during construction (GB), Disk = total on-disk index size (GB), Mem = total in-memory auxiliary structure size (GB).**

Method	SIFT				DEEP				SPACEV				GIST			
	BT	MaxM	Disk	Mem	BT	MaxM	Disk	Mem	BT	MaxM	Disk	Mem	BT	MaxM	Disk	Mem
Baseline	8814	31.9	39.06	3.81	14960	41.5	65.11	3.82	9393	36.4	39.31	3.81	1135	14.25	5.1	0.2
Baseline + MemGraph	8817	31.9	39.06	3.84	14981	41.5	65.11	3.87	9398	36.4	39.31	3.84	1135	14.25	5.1	0.204
Baseline + MemGraph + PageShuffle	11095	87.7	39.06	4.61	17303	90.8	65.11	4.70	11760	109.66	39.31	4.61	1158	14.25	5.1	0.209

storage and in-memory auxiliary structures). The Baseline configuration represents DiskANN’s original disk structure construction. Among the evaluated techniques, MemGraph and PageShuffle require additional offline preprocessing, while the other methods are query-time optimizations that incur no extra construction overhead. MemGraph constructs a small in-memory navigation graph by sampling a subset of vertices; its total overhead is modest. However, PageShuffle imposes substantial costs. First, the shuffling phase is time-consuming because it involves multiple iterations to solve an NP-hard optimization problem. Second, its memory footprint is large: the algorithm must load both the full graph and its reverse graph into memory, significantly increasing peak memory requirements—this constraint explains why our evaluation is limited to datasets up to 100 million vectors. Additionally, since records are reordered, the system must maintain an in-memory mapping structure between record IDs and page IDs, which further enlarging the memory index footprint.

**Finding 7: PageShuffle is both time-intensive and memory-intensive.** *If practitioners adopt PageShuffle to mitigate DiskANN’s read amplification by reorganizing page layouts, they must allocate additional build time and carefully evaluate memory capacity constraints, particularly for large-scale deployments.*

## 7 COMBINATION EFFECT EVALUATION

This section evaluates representative multi-factor combinations across memory layout, disk layout, and search scheduling, reports their cumulative effects and interactions, and compares the best combination against state-of-the-art systems.

### 7.1 Combination Evaluation

**7.1.1 Combination Design.** To keep the study tractable and, more importantly, interpretable, we select a small set of combinations guided by three principles: (i) complementarity—prioritizing pairs whose goals are orthogonal and therefore likely to yield synergy, (ii) minimal stepwise deltas—adding at most one new factor at a time so that marginal effects remain attributable.

Concretely, we exclude two techniques from the combination space: AiS, which incurs substantial on-disk expansion and read amplification; and SSSP-based Caching, which is less flexible in our setting due to static entry-point dependency and is rarely adopted in recent systems compared with MemGraph (cf. Table 2). As references, we keep three single-factor baselines from §5: Baseline (PQ), MemGraph, and DynamicWidth.

Based on these criteria, we instantiate five progressive combinations beyond the PQ baseline to reveal cumulative effects and potential conflicts/synergies. Specifically, we consider C1 (PS + PSe), which enables PageShuffle and PageSearch on top of Baseline;

**Table 6: Disk Metrics for five combination algorithms.**

Method	SIFT		DEEP		SPACEV		GIST	
	IOPS (K)	BW (MB/s)	IOPS (K)	BW (MB/s)	IOPS (K)	BW (MB/s)	IOPS (K)	BW (MB/s)
Baseline	791	3091	819	3200	768	3001	313	4889
MemGraph	818	3194	814	3181	813	3177	318	4973
DynamicWidth	796	3109	783	3063	765	2987	315	4930
PS+PSe	756	2954	809	3160	705	2755	309	4828
Pipe+DW	829	3238	827	3230	811	3166	321	5020
MemG+PS+PSe	751	2934	793	3099	683	2668	312	4872
MemG+Pipe+DW	818	3202	816	3189	826	3226	311	4951
MemG+PS+PSe+DW	704	2748	745	2910	643	2512	314	4909

C2 (Pipe + DW), which enables Pipeline and DynamicWidth on top of Baseline; C3 (MemG + PS + PSe), which adds MemGraph to C1; C4 (MemG + Pipe + DW), which adds MemGraph to C2; and C5 (MemG + PS + PSe + DW), which combines MemGraph with PS, PSe, and DynamicWidth. Pilot sweeps over additional permutations (e.g., PS+DW, PSe+DW, PS+Pipe) did not change conclusions and are omitted for brevity and space constraints.

**7.1.2 Combination Optimization Performance Analysis.** By examining Figures 15–16 and Table 6, we obtain three main observations:

**PageShuffle + PageSearch exhibit strong complementarity.** PS increases the likelihood that useful candidates are colocated within fetched pages; PSe then fully exploits each page by computing all in-page records. As shown in Figures 15 and 16, under the same  $L$  constraint, C1 attains higher recall (e.g., about a 10% gain on SIFT); moreover, even at high-accuracy settings (e.g.,  $L = 100$  with Recall  $\approx 98\%$ ), C1 still delivers roughly a 27% QPS improvement. Adding MemGraph (a.k.a, C3) further reduces the number of required hops by improving entry-point quality, compounding the benefit.

**Finding 8: PS + PSe deliver robust synergy: PS raises locality while PSe raises per-page utility.** *Their goals are complementary, jointly mitigating locality waste and underutilized pages. With MemGraph (C3), both throughput and latency improve further (vs. PQ baseline at matched Recall@10).*

**Pipeline and DynamicWidth are complementary but not co-equal.** DynamicWidth substantially mitigates the useless page reads induced by speculative I/O in pipeline search, making it a natural complement to Pipeline. In the single-factor results, Pipeline alone underperforms the PQ baseline (see Figure 11); in the combination study, however, C2 improves over the baseline by roughly 38% at matched accuracy (Figure 15). Nevertheless, C2 still trails the DynamicWidth-only configuration, which is consistent with Finding 5: under common high-concurrency conditions, Pipeline is often counterproductive and should not be enabled. For high-dimensional datasets, practitioners should tune the growth rate of DynamicWidth to mitigate potential adverse effects.




Figure 15: Recall–QPS trade-off: five combinations (C1–C5) across four datasets.




Figure 16: Latency–Recall@10 trade-off: five combinations (C1–C5) across four datasets.

**Finding 9: Pipeline and DynamicWidth are complementary, but DynamicWidth dominates. C2 (Pipe + DW) improves over the PQ baseline yet remains below DynamicWidth-only. Due to speculative reads and contention, Pipeline is unsuitable under concurrency; enable it only for single-thread, device-exclusive cases.**

**Full combination (C5) and definition of OctopusANN.** Integrating MemGraph with PageShuffle, PageSearch, and DynamicWidth achieves the best overall results among the evaluated combinations (throughput, latency, and I/O performance) under our settings, while keeping resource overhead modest (see Figures 15–16 and Table 6). The gains arise from complementary effects—higher page-level utilization (PS + PSe) and shorter convergence paths (MemGraph + DynamicWidth)—that together amplify I/O effectiveness. We refer to this full combination (C5) as **OctopusANN**. On SPACEV, C5 attains up to a 133% improvement. The additional cost is minor: small memory and disk overheads, plus extra offline building time for page shuffle during index construction.

**Finding 10: The multi-positive combination PQ + MemGraph + PS + PSe + DynamicWidth (C5) delivers the strongest overall results with modest cost—small memory/disk overheads and extra offline page shuffle time—while the online benefits dominate via better page utilization and faster convergence (vs. PQ baseline at matched Recall@10).**

## 7.2 State-of-the-Art Comparison

We compare our best combination (OctopusANN, C5) against state-of-the-art disk-based ANN systems at **Recall@10=90%** and **Recall@10=95%**. The competitors include DiskANN [77], Starling [88], and PipeANN [29]. At **Recall@10=90%**, OctopusANN delivers 4.1–37.9% higher QPS than Starling, and 87.5–149.5% over DiskANN (Figure 17). At the more stringent **Recall@10=95%**, the throughput advantage over Starling narrows to 1.6–14.6% (Figure 18), indicating diminishing gains for DynamicWidth at higher accuracy.




Figure 17: QPS comparison of state-of-the-art algorithms at 90% recall across 4 datasets; OctopusANN vs. Starling (+16.8–32.3%) and vs. DiskANN (+34.4–86.6%).




Figure 18: QPS comparison of state-of-the-art algorithms at 95% recall across 4 datasets; OctopusANN shows higher throughput.

Figure 19 shows the QPS comparison of four algorithms on the top-100k dataset. The relative performance trends among algorithms are similar to those observed for top-10 queries. However, for high-recall requirements (e.g., Recall@100 > 90%), OctopusANN shows less pronounced improvements. This is because the effectiveness of DynamicWidth diminishes as the number of search iterations increases, and top-100 queries typically require substantially more iterations than top-10 queries.

**OctopusANN breakdown analysis.** To clarify the cumulative effects of our systematic optimizations, Figure 20 presents a breakdown of the OctopusANN strategy and quantifies each technique’s incremental contribution to overall performance on SIFT dataset. The leftmost bar denotes the baseline; subsequent bars add optimizations cumulatively from left to right.




Figure 19: QPS comparison of top 100 across 4 systems.




Figure 21: The impact of page sizes.




Figure 20: OctopusANN optimization breakdown analysis.

The breakdown analysis reveals critical insights into our systematic optimization approach. The QPS improvements (left) show that MemGraph provides the foundation with 54.2% improvement over baseline, followed by PageShuffle&PageSearch (28.9%), and DynamicWidth (12.5%). The I/O performance analysis (right) demonstrates dramatic reductions in read pages per query: MemGraph reduces pages by 32.5%, PageShuffle&PageSearch by 28.3%, and DynamicWidth by 25.2%. This dual perspective validates that performance gains stem from I/O optimization.

**Finding 11: Systematic multi-dimensional optimization is most effective at moderate accuracy:** at 90% recall, OctopusANN outperforms Starling by 4.1–37.9% and DiskANN by 87.5–149.5%, whereas at 95% recall the gains are limited.

### 7.3 Discussion

**High-Dimensional Data and Page Size Trade-offs.** For high-dimensional data (e.g., approaching or exceeding 1,000 dimensions), the standard 4 KB page size is insufficient to store a single vector record, necessitating larger page sizes. Figure 21 compares the QPS of Baseline and PS+PSe on the GIST dataset under 8 KB and 16 KB page configurations, revealing three key trade-offs. First, optimization effectiveness varies significantly: with 8 KB pages, PS+PSe provides no performance improvement because each page accommodates only a single record, rendering intra-page shuffling ineffective; conversely, 16 KB pages enable effective shuffling by holding multiple records per page. Second, index sizes differ substantially—7.7 GB for 8 KB pages versus 5.1 GB for 16 KB pages—because the smaller page size incurs substantial storage waste from unused space, while 16 KB pages accommodate three records per page, reducing overhead. Third, despite the storage inefficiency, the 8 KB configuration achieves better overall query performance than 16 KB because smaller pages reduce per-I/O latency, even though they increase the total number of I/O operations.




Figure 22: A practical decision guide for disk-based ANN.

**Finding 12: For high-dimensional datasets, layout optimizations become less effective.** When each page holds fewer vector records, PageShuffle+PageSearch may not provide significant performance improvements. The I/O amplification problem for high-dimensional data warrants further investigation by the research community.

## 8 CONCLUSION

In this paper, we present an I/O-first, three-dimensional taxonomy for disk-based ANN and show on 100M-scale datasets that composing techniques across memory layout, disk layout, and search enhances both QPS and end-to-end latency. To translate the findings into practice, Figure 22 provides a concise decision guide. The workflow proceeds as follows: (i) when memory is sufficient, prefer mature in-memory indexes (e.g., HNSW, Vamana); (ii) when memory is tight and fast cold-start is required, All-in-Storage is a convenient choice despite higher on-disk expansion; (iii) for low-dimensional datasets under relaxed accuracy targets, inverted-index approaches such as SPANN can be a pragmatic choice; (iv) otherwise, decide by concurrency—under low concurrency, PipeANN can be competitive; under high concurrency, choose OctopusANN for high Recall@10 or deploy a lightweight MemGraph when moderate Recall@10 suffices. This matches our empirical results and clarifies when each method is the best fit.

## ACKNOWLEDGMENTS

We would like to express our sincere gratitude to the anonymous reviewers for their valuable feedback and insightful comments, which greatly contributed to the improvement of this manuscript. This work is supported in part by the National Natural Science Foundation of China under Grant No. U25B2020 and U25B2021.

## REFERENCES

- [1] Fabien André, Anne-Marie Kermarrec, and Nicolas Le Scouarnec. 2017. Accelerated Nearest Neighbor Search with Quick ADC. In *Proceedings of the 2017 ACM on International Conference on Multimedia Retrieval*. <https://doi.org/10.1145/3078971.3078992>
- [2] Kazuo Aoyama, Kazumi Saito, Hiroshi Sawada, and Naonori Ueda. 2011. Fast approximate similarity search based on degree-reduced neighborhood graphs. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*. <https://doi.org/10.1145/2020408.2020576>
- [3] Akhil Arora, Sakshi Sinha, Piyush Kumar, and Arnab Bhattacharya. 2018. HD-Index: Pushing the Scalability-Accuracy Boundary for Approximate kNN Search in High-Dimensional Spaces. *Proc. VLDB Endow.* 11 (2018), 906–919. <https://api.semanticscholar.org/CorpusID:5045956>
- [4] Martin Aumüller, Erik Bernhardsson, and Alexander Faithfull. 2020. ANN-Benchmarks: A benchmarking tool for approximate nearest neighbor algorithms. *Information Systems* (Jan 2020), 101374. <https://doi.org/10.1016/j.is.2019.02.006>
- [5] Jens Axboe. 2025. fio - Flexible I/O Tester. <https://github.com/axboe/fio>. Accessed: 2025-06-24.
- [6] Ilias Azizi, Karima Echihiabi, and Themis Palpanas. 2023. ELPIS: Graph-Based Similarity Search for Scalable Data Science. *Proc. VLDB Endow.* 16, 6 (Feb. 2023), 1548–1559. <https://doi.org/10.14778/3583140.3583166>
- [7] Ilias Azizi, Karima Echihiabi, and Themis Palpanas. 2025. Graph-Based Vector Search: An Experimental Evaluation of the State-of-the-Art. *Proceedings of the ACM on Management of Data* 3 (2025), 1 – 31. <https://api.semanticscholar.org/CorpusID:276249205>
- [8] 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*. 2055–2063.
- [9] Antoine Boutet, Anne-Marie Kermarrec, Nupur Mittal, and Francois Taiani. 2016. Being prepared in a sparse world: The case of KNN graph construction. In *2016 IEEE 32nd International Conference on Data Engineering (ICDE)*. 241–252. <https://doi.org/10.1109/ICDE.2016.7498244>
- [10] Yuan Cao, Heng Qi, Wenrui Zhou, Jien Kato, Keqiu Li, Xiulong Liu, and Jie Gui. 2018. Binary Hashing for Approximate Nearest Neighbor Search on Big Data: A Survey. *IEEE Access* (Jan 2018), 2039–2054. <https://doi.org/10.1109/access.2017.2781360>
- [11] Mingkai Chen, Tianhua Han, Cheng Liu, Shengwen Liang, Kuai Yu, Lei Dai, Ziming Yuan, Ying Wang, Lei Zhang, Huawei Li, and Xiaowei Li. 2024. DRIM-ANN: An Approximate Nearest Neighbor Search Engine based on Commercial DRAM-PIMs. *ArXiv abs/2410.15621* (2024). <https://api.semanticscholar.org/CorpusID:273501609>
- [12] Patrick Chen, Wei-Cheng Chang, Jyun-Yu Jjiang, Hsiang-Fu Yu, Inderjit Dhillon, and Cho-Jui Hsieh. 2023. FINGER: Fast Inference for Graph-based Approximate Nearest Neighbor Search. In *Proceedings of the ACM Web Conference 2023* (Austin, TX, USA) (WWW '23). Association for Computing Machinery, New York, NY, USA, 3225–3235. <https://doi.org/10.1145/3543507.3583318>
- [13] Qi Chen, Bing Zhao, Haidong Wang, Mingqin Li, Chuanjie Liu, Zengzhong Li, Mao Yang, and Jingdong Wang. 2021. SPANN: Highly-efficient Billion-scale Approximate Nearest Neighbor Search. *NeurIPS* (Nov 2021).
- [14] Scott Cost and Steven Salzberg. 1993. A Weighted Nearest Neighbor Algorithm for Learning with Symbolic Features. *Machine Learning* 10, 1 (Jan 1993), 57–78. <https://doi.org/10.1023/a:1022664626993>
- [15] T. Cover and P. Hart. [n.d.]. Nearest neighbor pattern classification. *IEEE Transactions on Information Theory* (n. d.), 21–27. <https://doi.org/10.1109/tit.1967.1053964>
- [16] Liwei Deng, Penghao Chen, Ximu Zeng, Tianfu Wang, Yan Zhao, and Kai Zheng. 2024. Efficient Data-Aware Distance Comparison Operations for High-Dimensional Approximate Nearest Neighbor Search. *Proc. VLDB Endow.* 18, 3 (Nov. 2024), 812–821. <https://doi.org/10.14778/3712221.3712244>
- [17] Edsger W. Dijkstra. 1959. A note on two problems in connexion with graphs. *Numer. Math.* 1, 1 (1959), 269–271. <https://doi.org/10.1007/BF01386390>
- [18] Myron Flickner, Harpreet Sawhney, W. Niblack, Jonathan Ashley, Qian Huang, Byron Dom, M. Gorkani, JamesLee Hafner, Denis Lee, Dragutin Petkovic, DavidA. Steele, and P. Yanker. 2002. Query by Image and Video Content: The QBIC System. (Jan 2002).
- [19] STEVEN FORTUNE. 1992. VORONOI DIAGRAMS and DELAUNAY TRIANGULATIONS. 193–233. [https://doi.org/10.1142/9789814355858\\_0006](https://doi.org/10.1142/9789814355858_0006)
- [20] Cong Fu, Changxu Wang, and Deng Cai. 2022. High Dimensional Similarity Search With Satellite System Graph: Efficiency, Scalability, and Unindexed Query Compatibility. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 44, 8 (2022), 4139–4150. <https://doi.org/10.1109/TPAMI.2021.3067706>
- [21] Cong Fu, Chao Xiang, Changxu Wang, and Deng Cai. 2017. Fast Approximate Nearest Neighbor Search With The Navigating Spreading-out Graph. *Proc. VLDB Endow.* 12 (2017), 461–474. <https://api.semanticscholar.org/CorpusID:44175941>
- [22] Jianyang Gao and Cheng Long. 2023. High-Dimensional Approximate Nearest Neighbor Search: with Reliable and Efficient Distance Comparison Operations. *Proc. ACM Manag. Data* 1, 2, Article 137 (June 2023), 27 pages. <https://doi.org/10.1145/3589282>
- [23] Jianyang Gao and Cheng Long. 2024. RaBitQ: Quantizing High-Dimensional Vectors with a Theoretical Error Bound for Approximate Nearest Neighbor Search. *Proceedings of the ACM on Management of Data* 2 (2024), 1 – 27. <https://api.semanticscholar.org/CorpusID:269929866>
- [24] Tiezheng Ge, Kaiming He, Qifa Ke, and Jian Sun. 2013. Optimized product quantization. *IEEE transactions on pattern analysis and machine intelligence* 36, 4 (2013), 744–755.
- [25] Allen Gersho and Robert M Gray. 2012. *Vector quantization and signal compression*. Vol. 159. Springer Science & Business Media.
- [26] Siddharth Gollapudi, Neel Karia, Varun Sivashankar, Ravishankar Krishnaswamy, Nikit Begwani, Swapnil Raz, Yiyong Lin, Yin Zhang, Neelam Mahapatro, Premkumar Srinivasan, Amit Singh, and Harsha Vardhan Simhadri. 2023. Filtered-DiskANN: Graph Algorithms for Approximate Nearest Neighbor Search with Filters. In *Proceedings of the ACM Web Conference 2023* (Austin, TX, USA) (WWW '23). Association for Computing Machinery, New York, NY, USA, 3406–3416. <https://doi.org/10.1145/3543507.3583552>
- [27] Long Gong, Huayi Wang, Mitsunori Ogihara, and Jun Xu. 2020. iDEC: Indexable Distance Estimating Codes for Approximate Nearest Neighbor Search. *Proceedings of the VLDB Endowment* (May 2020), 1483–1497. <https://doi.org/10.14778/3397230.3397243>
- [28] Goetz Graefe. 2011. Modern B-Tree Techniques. *Foundations and Trends® in Databases* 3, 4 (2011), 203–402. <https://doi.org/10.1561/19000000028>
- [29] Hao Guo and Youyou Lu. 2025. Achieving Low-Latency Graph-Based Vector Search via Aligning Best-First Search Algorithm with SSD. In *Proceedings of the 19th USENIX Symposium on Operating Systems Design and Implementation (OSDI '25)*.
- [30] Hao Guo and Youyou Lu. 2025. PipeANN: A Low Latency, Billion-Scale, and Updatable Graph-Based Vector Store on SSD. <https://github.com/thustorage/PipeANN>. <https://github.com/thustorage/PipeANN> GitHub repository.
- [31] Rentong Guo, Xiaofan Luan, Long Xiang, Xiao Yan, Xiaomeng Yi, Jigao Luo, Qianya Cheng, Weizhi Xu, Jiarui Luo, Frank Liu, Zhenshan Cao, Yanliang Qiao, Ting Wang, Bo Tang, and Charles Xie. 2022. Manu: a cloud native vector database management system. *Proc. VLDB Endow.* 15, 12 (Aug. 2022), 3548–3561. <https://doi.org/10.14778/3554821.3554843>
- [32] Peter E. Hart, Nils J. Nilsson, and Bertram Raphael. 1968. A Formal Basis for the Heuristic Determination of Minimum Cost Paths. *IEEE Transactions on Systems Science and Cybernetics* 4, 2 (1968), 100–107. <https://doi.org/10.1109/TSSC.1968.300136>
- [33] Qiang Huang, Jianlin Feng, Qiong Fang, Wilfred Ng, and Wei Wang. 2017. Query-aware locality-sensitive hashing scheme for  $l_p$  norm. *The VLDB Journal* 26, 5 (Oct 2017), 683–708. <https://doi.org/10.1007/s00778-017-0472-7>
- [34] Qiang Huang, Jianlin Feng, Yikai Zhang, Qiong Fang, and Wilfred Ng. 2015. Query-aware locality-sensitive hashing for approximate nearest neighbor search. *Proceedings of the VLDB Endowment* (Sep 2015), 1–12. <https://doi.org/10.14778/2850469.2850470>
- [35] Huawei Cloud. 2024. GaussDB(for openGauss) Vector Database User Guide (for Centralized/Distributed Mode). <https://doc.hcs.huawei.com/en-us/gaussdb/doc/download/pdf/gaussdb-vector-cent.pdf>. Accessed: 2025-05-27.
- [36] IRISA. 2025. SIFT: Datasets for ANN search. <http://corpus-texmex.irisa.fr/>. Accessed: 2025-06-25.
- [37] Masajiro Iwasaki. 2016. *Pruned Bi-directed K-nearest Neighbor Graph for Proximity Search*. 20–33. [https://doi.org/10.1007/978-3-319-46759-7\\_2](https://doi.org/10.1007/978-3-319-46759-7_2)
- [38] Shikhar Jaiswal, Ravishankar Krishnaswamy, Ankit Garg, Harsha Vardhan Simhadri, and Sheshansh Agrawal. 2022. OOD-DiskANN: Efficient and Scalable Graph ANNS for Out-of-Distribution Queries. *arXiv:2211.12850* [cs.LG] <https://arxiv.org/abs/2211.12850>
- [39] Junhyeok Jang, Hanjin Choi, Hanyeoreum Bae, Seungjun Lee, Miryeong Kwon, and Myoungsoo Jung. 2023. CXL-ANNS: Software-Hardware Collaborative Memory Disaggregation and Computation for Billion-Scale Approximate Nearest Neighbor Search. (Feb 2023).
- [40] Junhyeok Jang, Hanjin Choi, Hanyeoreum Bae, Seungjun Lee, Miryeong Kwon, and Myoungsoo Jung. 2024. Bridging Software-Hardware for CXL Memory Disaggregation in Billion-Scale Nearest Neighbor Search. *ACM Transactions on Storage* 20, 2 (Feb. 2024), 1–30. <https://doi.org/10.1145/3639471>
- [41] Yahoo Japan. 2015. NGT: Nearest Neighbor Search with Neighborhood Graph and Tree for High-dimensional Data. <https://github.com/yahoojapan/NGT>. Accessed: 2025-06-18.
- [42] ZhuChun Jiang, Tan Zhu, Haining Li, Jinbo Bi, and Minghu Song. 2019. Accelerating Large-Scale Molecular Similarity Search through Exploiting High Performance Computing. *IEEE Conference Proceedings, IEEE Conference Proceedings* (Jan 2019).
- [43] Chao Jin, Zili Zhang, Xuanlin Jiang, Fangyue Liu, Xin Liu, Xuanzhe Liu, and Xin Jin. 2024. RAGCache: Efficient Knowledge Caching for Retrieval-Augmented Generation. *arXiv:2404.12457* [cs.DC] <https://arxiv.org/abs/2404.12457>
- [44] H Jégou, M Douze, and C Schmid. 2011. Product Quantization for Nearest Neighbor Search. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (Jan 2011), 117–128. <https://doi.org/10.1109/tpami.2010.57>

- [45] Ji-Hoon Kim, Yeo-Reum Park, Jaeyoung Do, Soo-Young Ji, and Joo-Young Kim. 2023. Accelerating Large-Scale Graph-Based Nearest Neighbor Search on a Computational Storage Platform. *IEEE Trans. Comput.* (Jan 2023), 278–290. <https://doi.org/10.1109/tc.2022.3155956>
- [46] Seoyoung Ko, Hyunjeong Shim, Wanju Doh, Sungmin Yun, Jinin So, Yong-suk Kwon, Sang-Soo Park, Si-Dong Roh, Minyong Yoon, Taeksang Song, and Jung Ho Ahn. 2025. Cosmos: A CXL-Based Full In-Memory System for Approximate Nearest Neighbor Search. *IEEE Computer Architecture Letters* 24, 1 (Jan. 2025), 173–176. <https://doi.org/10.1109/lca.2025.3570235>
- [47] Atsutake Kosuge and Takashi Oshima. 2019. An Object-Pose Estimation Acceleration Technique for Picking Robot Applications by Using Graph-Reusing k-NN Search. In *2019 First International Conference on Graph Computing (GC)*. <https://doi.org/10.1109/gc46384.2019.00018>
- [48] Joseph B. Kruskal. 2010. On the shortest spanning subtree of a graph and the traveling salesman problem. *Proc. Amer. Math. Soc.* (Jul 2010), 48–50. <https://doi.org/10.1090/s0002-9939-1956-0078686-7>
- [49] Patrick Lewis, Ethan Perez, Aleksandara Piktus, Filippo 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. (May 2020).
- [50] Conglong Li, Minjia Zhang, David G. Andersen, and Yuxiong He. 2020. Improving Approximate Nearest Neighbor Search through Learned Adaptive Early Termination. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data*. <https://doi.org/10.1145/3318464.3380600>
- [51] Conglong Li, Minjia Zhang, David G. Andersen, and Yuxiong He. 2020. Improving Approximate Nearest Neighbor Search through Learned Adaptive Early Termination. *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data* (2020). <https://api.semanticscholar.org/CorpusID:218914310>
- [52] Liang Li, Shufeng Gong, Yanan Yang, Yiduo Wang, and Jie Wu. 2026. I/O Optimizations for Graph-Based Disk-Resident Approximate Nearest Neighbor Search: A Design Space Exploration. *arXiv:2602.21514* [cs.DB] <https://arxiv.org/abs/2602.21514>
- [53] Wen Li, Ying Zhang, Yifang Sun, Wei Wang, Mingjie Li, Wenjie Zhang, and Xuemin Lin. 2016. Approximate Nearest Neighbor Search on High Dimensional Data – Experiments, Analyses, and Improvement. *IEEE Transactions on Knowledge and Data Engineering* (Aug 2016), 1475–1488. <https://doi.org/10.1109/tkde.2019.2909204>
- [54] Shengwen Liang, Ying Wang, Youyou Lu, Zhe Yang, Huawei Li, and Xiaowei Li. 2019. Cognitive {SSD}: A deep learning engine for {In-Storage} data retrieval. In *2019 USENIX Annual Technical Conference (USENIX ATC 19)*. 395–410.
- [55] Linux Man Pages Project. 2025. *io\_uring(7) – Linux manual page*. Linux Programmer’s Manual. Accessed: 2025-06-24.
- [56] Yury Malkov, Alexander Ponomarenko, Andrey Logvinov, and Vladimir Krylov. 2014. Approximate nearest neighbor algorithm based on navigable small world graphs. *Inf. Syst.* 45 (2014), 61–68. <https://api.semanticscholar.org/CorpusID:9896397>
- [57] 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* (Apr 2020), 824–836. <https://doi.org/10.1109/tpami.2018.2889473>
- [58] Luke Mathieson and Pablo Moscato. 2019. *An Introduction to Proximity Graphs*. Springer International Publishing, Cham, 213–233. [https://doi.org/10.1007/978-3-030-06222-4\\_4](https://doi.org/10.1007/978-3-030-06222-4_4)
- [59] Yitong Meng, Xiaohui Dai, Yan Xiao, James Cheng, Weiwen Liu, Benben Liao, Jun Guo, and Guangyong Chen. 2019. PMD: An Optimal Transportation-based User Distance for Recommender Systems. *Cornell University - arXiv,Cornell University - arXiv* (Sep 2019).
- [60] Microsoft. 2018. SPTAG: A library for fast approximate nearest neighbor search. <https://github.com/microsoft/SPTAG/>. Accessed: 2025-06-18.
- [61] Microsoft. 2024. SPACEV1B: A Billion-scale Dataset for Vector Search. <https://github.com/microsoft/SPTAG/tree/main/datasets/SPACEV1B>. Part of Microsoft SPTAG project, O-UDA license.
- [62] Microsoft. 2025. Enable and use DiskANN extension. <https://learn.microsoft.com/en-us/azure/postgresql/flexible-server/how-to-use-pgdiskann>. Accessed: 2025-05-27.
- [63] Jiongfeng Ni, Xiaoliang Xu, Yuxiang Wang, Can Li, Jiajie Yao, Shihai Xiao, and Xuecang Zhang. 2023. DiskANN++: Efficient Page-based Search over Isomorphic Mapped Graph Index using Query-sensitivity Entry Vertex. *arXiv:2310.00402* [cs.IR] <https://arxiv.org/abs/2310.00402>
- [64] Salman Niazi, Mahmoud Ismail, Seif Haridi, Jim Dowling, Steffen Grohshchmiedt, and Mikael Ronström. 2017. HopsFS: Scaling Hierarchical File System Metadata Using NewSQL Databases. In *15th USENIX Conference on File and Storage Technologies (FAST 17)*. USENIX Association, Santa Clara, CA, 89–104. <https://www.usenix.org/conference/fast17/technical-sessions/presentation/niazi>
- [65] Yu Pan, Jianxin Sun, and Hongfeng Yu. 2023. LM-DiskANN: Low Memory Footprint in Disk-Native Dynamic Graph-Based ANN Indexing. In *2023 IEEE International Conference on Big Data (BigData)*. 5987–5996. <https://doi.org/10.1109/BigData59044.2023.10386517>
- [66] Zhibin Pan, Liangzhuang Wang, Yang Wang, and Yuchen Liu. 2020. Product quantization with dual codebooks for approximate nearest neighbor search. *Neurocomputing* (Aug 2020), 59–68. <https://doi.org/10.1016/j.neucom.2020.03.016>
- [67] Rodrigo Paredes and Edgar Chávez. 2005. *Using the k-Nearest Neighbor Graph for Proximity Searching in Metric Spaces*. 127–138. [https://doi.org/10.1007/11575832\\_14](https://doi.org/10.1007/11575832_14)
- [68] 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, Article 54 (May 2023), 27 pages. <https://doi.org/10.1145/3588908>
- [69] Jie Ren, Minjia Zhang, and Dong Liu. 2020. HM-ANN: Efficient Billion-Point Nearest Neighbor Search on Heterogeneous Memory. *NeurIPS* (Nov 2020).
- [70] Philipp M. Riegger. 2010. Literature Survey on Nearest Neighbor Search and Search in Graphs. (Jan 2010).
- [71] Badrul Sarwar, George Karypis, Joseph Konstan, and John Riedl. 2001. Item-based collaborative filtering recommendation algorithms. In *Proceedings of the 10th international conference on World Wide Web*. <https://doi.org/10.1145/371920.372071>
- [72] Chanop Silpa-Anan and Richard Hartley. 2008. Optimised KD-trees for fast image descriptor matching. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*. <https://doi.org/10.1109/cvpr.2008.4587638>
- [73] Harsha Simhadri and Ravishankar Krishnaswamy. 2021. FreshDiskANN: A Fast and Accurate Graph-Based ANN Index for Streaming Similarity Search. (May 2021). <https://www.microsoft.com/en-us/research/publication/freshdiskann-a-fast-and-accurate-graph-based-ann-index-for-streaming-similarity-search/>
- [74] Harsha Vardhan Simhadri, Martin Aumüller, Amir Ingber, Matthijs Douze, George Williams, Magdalen Dobson Manohar, Dmitry Baranchuk, Edo Liberty, Frank Liu, Ben Landrum, et al. 2024. Results of the Big ANN: NeurIPS’23 competition. *arXiv preprint arXiv:2409.17424* (2024).
- [75] Harsha Vardhan Simhadri, Ravishankar Krishnaswamy, Gopal Srinivasa, Suhas Jayaram Subramanya, Andrija Antonijevic, Dax Pryce, David Kaczynski, Shane Williams, Siddarth Gollapudi, Varun Sivashankar, Neel Karia, Aditi Singh, Shikhar Jaiswal, Neelam Mahapatro, Philip Adams, Bryan Tower, and Yash Patel. 2025. DiskANN: Graph-structured Indices for Scalable, Fast, Fresh and Filtered Approximate Nearest Neighbor Search. <https://github.com/Microsoft/DiskANN> Microsoft Open Source Project.
- [76] Harsha Vardhan Simhadri, George Williams, Martin Aumüller, Matthijs Douze, Artem Babenko, Dmitry Baranchuk, Qi Chen, Lucas Hosseini, Ravishankar Krishnaswamy, Gopal Srinivasa, et al. 2022. Results of the NeurIPS’21 challenge on billion-scale approximate nearest neighbor search. In *NeurIPS 2021 Competitions and Demonstrations Track*. PMLR, 177–189.
- [77] Suhas Jayaram Subramanya, Devvrit, Rohan Kadekodi, Ravishankar Krishnaswamy, and Harsha Vardhan Simhadri. 2019. DiskANN : Fast Accurate Billion-point Nearest Neighbor Search on a Single Node. <https://api.semanticscholar.org/CorpusID:209392043>
- [78] Gongjin Sun and Sang-Woo Jun. 2020. Bandwidth Efficient Near-Storage Accelerator for High-Dimensional Similarity Search. In *2020 International Conference on Field-Programmable Technology (ICFPT)*. 129–138. <https://doi.org/10.1109/icfpt51103.2020.00026>
- [79] Wenda Tang, Ying Han, Tianxiang Ai, Guanghui Li, Bin Yu, and Xin Yang. 2024. Yggdrasil: Reducing Network I/O Tax with (CXL-Based) Distributed Shared Memory. In *Proceedings of the 53rd International Conference on Parallel Processing, ICPP 2024, Gotland, Sweden, August 12-15, 2024*. ACM, 597–606. <https://doi.org/10.1145/3673038.3673138>
- [80] Kento Tatsuno, Daisuke Miyashita, Taiga Ikeda, Kiyoshi Ishiyama, Kazunari Sumiyoshi, and Jun Deguchi. 2025. AISAQ: All-in-Storage ANNS with Product Quantization for DRAM-free Information Retrieval. *arXiv:2404.06004* [cs.IR] <https://arxiv.org/abs/2404.06004>
- [81] Zilliz Tech. 2025. Starling: An I/O-Efficient Disk-Resident Graph Index Framework for High-Dimensional Vector Similarity Search on Data Segment. <https://github.com/zilliztech/starling>. <https://github.com/zilliztech/starling> GitHub repository.
- [82] Eric Sadit Tellez, Guillermo Ruiz, Edgar Chávez, and Mario Graff. 2017. A scalable solution to the nearest neighbor search problem through local-search methods on neighbor graphs. *Pattern Analysis and Applications* 24 (2017), 763 – 777. <https://api.semanticscholar.org/CorpusID:233466017>
- [83] Bing Tian, Haikun Liu, Zhuohui Duan, Xiaofei Liao, Hai Jin, and Yu Zhang. 2024. Scalable Billion-point Approximate Nearest Neighbor Search Using SmartSSDs. *ATC* (Jan 2024).
- [84] Timescale. 2025. pgvector: A complement to pgvector for high performance, cost efficient vector search on large workloads. <https://github.com/timescale/pgvector>. Accessed: 2025-05-27.
- [85] Godfried T Toussaint. 1980. The relative neighbourhood graph of a finite planar set. *Pattern recognition* 12, 4 (1980), 261–268.
- [86] T. Tsao, D. Comeau, and H. Margolin. 1972. *A Multifactor Paging Experiment: I. The Experiment and Conclusions*. Academic Press.

- [87] Jianguo Wang, Xiaomeng Yi, Rentong Guo, Hai Jin, Peng Xu, Shengjun Li, Xiangyu Wang, Xiangzhou Guo, Chengming Li, Xiaohai Xu, Kun Yu, Yuxing Yuan, Yinghao Zou, Jiquan Long, Yudong Cai, Zhenxiang Li, Zhifeng Zhang, Yihua Mo, Jun Gu, Ruiyi Jiang, Yi Wei, and Charles Xie. 2021. Milvus: A Purpose-Built Vector Data Management System. *Proceedings of the 2021 International Conference on Management of Data* (2021). <https://api.semanticscholar.org/CorpusID:235474148>
- [88] 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. *Proceedings of the ACM on Management of Data* (Mar 2024), 1–27. <https://doi.org/10.1145/3639269>
- [89] Mengzhao Wang, Xiaoliang Xu, Qiang Yue, and Yuxiang Wang. 2021. A Comprehensive Survey and Experimental Comparison of Graph-Based Approximate Nearest Neighbor Search. *Proc. VLDB Endow.* 14 (2021), 1964–1978. <https://api.semanticscholar.org/CorpusID:231728434>
- [90] Chuangxian Wei, Bin Wu, Sheng Wang, Renjie Lou, Chaoqun Zhan, Feifei Li, and Yuanzhe Cai. 2020. AnalyticDB-V: a hybrid analytical engine towards query fusion for structured and unstructured data. *Proc. VLDB Endow.* 13, 12 (Aug. 2020), 3152–3165. <https://doi.org/10.14778/3415478.3415541>
- [91] 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, Article 79 (Feb. 2025), 29 pages. <https://doi.org/10.1145/3709729>
- [92] Songlin Wu. 2024. An Asynchronous I/O-Based Disk Vector Retrieval Algorithm. *Computer Science and Application* 14, 01 (2024), 68–77. <https://doi.org/10.12677/csa.2024.141008>
- [93] Yuming Xu, Hengyu Liang, Jin Li, Shuotao Xu, Qi Chen, Qianxi Zhang, Cheng Li, Ziyue Yang, Fan Yang, Yuqing Yang, et al. 2023. Sffresh: Incremental in-place update for billion-scale vector search. In *Proceedings of the 29th Symposium on Operating Systems Principles*. 545–561.
- [94] Shuo Yang, Jiadong Xie, Yingfan Liu, Jeffrey Xu Yu, Xiyue Gao, Qianru Wang, Yanguo Peng, and Jiangtao Cui. 2024. Revisiting the Index Construction of Proximity Graph-Based Approximate Nearest Neighbor Search. *ArXiv abs/2410.01231* (2024). <https://api.semanticscholar.org/CorpusID:273025855>
- [95] Xizhe Yin, Chao Gao, Zhijia Zhao, and Rajiv Gupta. 2025. PANNs: Enhancing Graph-based Approximate Nearest Neighbor Search through Recency-aware Construction and Parameterized Search. In *Proceedings of the 30th ACM SIGPLAN Annual Symposium on Principles and Practice of Parallel Programming* (Las Vegas, NV, USA) (PPoPP '25). Association for Computing Machinery, New York, NY, USA, 369–381. <https://doi.org/10.1145/3710848.3710867>
- [96] Minjia Zhang and Yuxiong He. 2018. Zoom: SSD-based Vector Search for Optimizing Accuracy, Latency and Memory. *arXiv: Computer Vision and Pattern Recognition*, arXiv: Computer Vision and Pattern Recognition (Sep 2018).
- [97] Minjia Zhang and Yuxiong He. 2019. GRIP: Multi-Store Capacity-Optimized High-Performance Nearest Neighbor Search for Vector Search Engine. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management (CIKM '19)*. ACM, 1673–1682. <https://doi.org/10.1145/3357384.3357938>
- [98] Qianxi Zhang, Shuotao Xu, Qi Chen, Guoxin Sui, Jiadong Xie, Zhizhen Cai, Yaoqi Chen, Yinxuan He, Yuqing Yang, Fan Yang, Mao Yang, and Lidong Zhou. 2023. VBASE: Unifying Online Vector Similarity Search and Relational Queries via Relaxed Monotonicity. In *17th USENIX Symposium on Operating Systems Design and Implementation (OSDI 23)*. USENIX Association, Boston, MA, 377–395. <https://www.usenix.org/conference/osdi23/presentation/zhang-qianxi>
- [99] Zilliz. 2024. *BBAnn: Block-based Approximate Nearest Neighbor*. <https://github.com/zilliztech/BBAnn> An algorithm optimized for SSD storage that organizes data aligned with SSD block size.
- [100] Zilliz. 2025. On-disk Index – Milvus Documentation. [https://milvus.io/docs/disk\\_index.md](https://milvus.io/docs/disk_index.md). Accessed: 2025-05-27.