

All-in-one Graph-based Indexing for Hybrid Search on GPUs

Zhonggen Li
Zhejiang University
zgli@zju.edu.cn

Yougen Li
Zhejiang University
yougen.24@intl.zju.edu.cn

Yifan Zhu
Zhejiang University
xtf_z@zju.edu.cn

Zhaoqiang Chen
Huawei Cloud
chenzhaoqiang1@huawei.com

Yunjun Gao
Zhejiang University
gaoyj@zju.edu.cn

Abstract

Hybrid search has emerged as a promising paradigm to overcome the limitations of single-path retrieval, enhancing accuracy for applications like recommendations, information retrieval, and Retrieval-Augmented Generation. However, existing methods are constrained by a trilemma: they sacrifice flexibility for efficiency, suffer from accuracy degradation due to separate retrievals, or incur prohibitive storage overhead for flexible combinations of retrieval paths.

This paper introduces Allan-Poe, a novel **All-in-one graph index** accelerated by **GPUs for efficient hybrid search**. We first analyze the limitations of existing retrieval paradigms and distill key design principles for an effective hybrid search index. Guided by these principles, we architect a unified graph-based index that flexibly integrates four retrieval paths—dense vector, sparse vector, full-text, and knowledge graph—within a single, cohesive structure. To enable efficient construction, we design a GPU-accelerated pipeline featuring a warp-level hybrid distance kernel, RNG-IP joint pruning, and keyword-aware neighbor recycling. For query processing, we introduce a dynamic fusion framework that supports any combination of retrieval paths and weights without index reconstruction, leveraging logical edges from the knowledge graph to resolve complex multi-hop queries. Extensive experiments on 6 real-world datasets demonstrate that Allan-Poe achieves superior end-to-end query accuracy and outperforms state-of-the-art methods by 1.5-186.4× in throughput, while significantly reducing storage overhead.

PVLDB Reference Format:

Zhonggen Li, Yougen Li, Yifan Zhu, Zhaoqiang Chen, and Yunjun Gao. All-in-one Graph-based Indexing for Hybrid Search on GPUs. PVLDB, 19(1): XXX-XXX, 2026. doi:XX.XX/XXX.XX

PVLDB Artifact Availability:

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

1 Introduction

Recent advancements in vector databases have substantially improved the accuracy and efficiency of dense vector retrieval [12, 30, 65]. State-of-the-art approximate nearest neighbor search algorithms now consistently achieve over 99% recall for top- k neighbor




Figure 1: Comparison of existing hybrid search paradigms.

searches. However, the end-to-end accuracy—that is, the accuracy of retrieved documents rather than vector similarity—remains limited. This is because the minimal distance between query and answer vectors in the embedding space does not guarantee their semantic relevance in natural language [53, 70, 98]. This fundamental discrepancy between geometric proximity and semantic meaning hinders the broader adoption of vector databases in critical areas such as search engines [32, 47, 49], recommendation systems [60, 66, 89], and Retrieval-Augmented Generation (RAG) [7, 15, 40, 94].

In addition to the popular use of dense vectors, alternative retrieval methods utilize statistics-based [58, 76] and learned-based [11, 24, 45] sparse vectors to improve the end-to-end accuracy. While these sparse vectors offer superior interpretability and cross-domain robustness, their semantic representation capabilities are generally weaker than dense vectors. Consequently, dependence on any single retrieval path is often insufficient for achieving high end-to-end retrieval relevance [56, 84]. To overcome this limitation, hybrid search has emerged as a promising solution [56, 63, 79]. As illustrated in Figure 1, two primary methodologies have been proposed: (1) *Fusion retrieval* integrates dense and sparse vectors by mapping them into a unified space via dimensionality reduction [10] or by constructing the graph-based index with weighted distance calculations [98]. However, these approaches often suffer from low efficiency and precision, or exhibit poor extensibility. Furthermore, they are typically limited to dense and sparse vectors, failing to accommodate the complex requirements of real-world applications. (2) *Multi-path retrieval* constructs separate indexes for various retrieval paths, including dense vector [26, 57, 91], sparse vector [11, 24, 45], and full-text search [58, 76]. For a given query, this paradigm retrieves the top- k' neighbors from each path independently. These intermediate results are subsequently fused using re-ranking methods—such as Reciprocal Rank Fusion (RRF) [9, 18], Weighted Sum [2, 4], or ColBERT [23, 44]—to produce the final top- k ($k \leq k'$) list. Due to its effectiveness and flexibility, this paradigm is widely adopted in modern databases [1–3, 5].

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment. Proceedings of the VLDB Endowment, Vol. 19, No. 1 ISSN 2150-8097. doi:XX.XX/XXX.XX

Despite the effectiveness, the paradigm of multi-path retrieval introduces two primary problems due to its decoupled architecture. (1) *Index storage overhead*. Each retrieval path necessitates the construction, maintenance, and storage of a dedicated index. This not only increases system complexity but also incurs significant storage and operational overhead¹ [51, 98]. (2) *Retrieval efficiency and accuracy*. The optimal results for a hybrid query may not be present within the top- k' results from any single path [77, 78, 101], leading to decreased accuracy. As a result, neither *Fusion retrieval* nor *Multi-path retrieval* can simultaneously deliver high accuracy, efficiency, and extensibility.

To overcome the aforementioned limitations, inspired by *Edgar Allan Poe's* literary theory of "*Unity of Effect*", this work explores unifying diverse multi-path indexes into a single, all-in-one hybrid index, leveraging the GPU to achieve efficient construction and real-time retrieval. The graph-based index has recently gained prominence as a leading approach for approximate nearest neighbor search (ANNS), owing to its superior efficiency and accuracy [39, 74, 80, 86]. This naturally raises a key research question: *how can we design a unified graph-based index for effective hybrid search?* However, such a design is non-trivial. Practical applications with real-time requirements—such as recommendation systems, search engines, and RAG—demand a hybrid index that is simultaneously efficient, effective, and flexible. We identify two primary challenges that must be addressed:

Challenge I: *How to design a unified and flexible graph-based hybrid index?* The fundamental heterogeneity of retrieval methods presents the primary obstacle: dense vectors utilize graph-based indexes, sparse vectors and full-text search rely on inverted indexes, and knowledge graphs employ entity-relationship structures. Fusing these disparate architectures into a single, coherent index without compromising performance is non-trivial. Furthermore, optimal fusion weights for the similarity from each retrieval path are inherently dynamic, varying by query context and user preference. Pre-computing and storing indexes for all possible weight combinations is infeasible. Finally, exhaustive use of all available paths for every query is suboptimal, unnecessarily compromising efficiency when fewer paths would suffice. Consequently, designing a flexible structure that supports arbitrary path combinations without index reconstruction remains a critical challenge.

Challenge II: *How to achieve efficient construction and effective search on the hybrid index?* While Inner Product (IP) serves as the primary similarity metric [10, 79], existing graph indexes optimized for it suffer from inefficient construction and redundant edge connectivity. The indexing and querying processes also involve massive, hybrid distance calculations. For example, dense vectors and sparse vectors comparisons require vector dot product and set intersection operations, respectively. These disparate computational characteristics create challenges for efficient memory access and effective GPU parallelization. Furthermore, the computational demands of edges from various retrieval paths in the hybrid index introduce multiple efficiency bottlenecks. Semantic matching alone proves inadequate for complex reasoning tasks. Although knowledge graphs provide complementary logical similarity [13, 100], a fundamental

granularity mismatch exists: our index operates on document-level representations while knowledge graphs model fine-grained entity relationships. This disparity complicates the effective integration of logical reasoning into the vector search process.

To address these challenges, we propose Allan-Poe, a unified and flexible graph-based hybrid index accelerated by GPUs. Our solution integrates dense, sparse, full-text, and knowledge graph-based retrieval through an isolated heterogeneous edge storage mechanism. This design supports any combination of retrieval paths without requiring index reconstruction or sacrificing efficiency. We create a Unified Semantic Metric Space that fuses multiple vector representations from diverse retrieval paths into a single similarity metric, with theoretical proof demonstrating its effectiveness for arbitrary fusion weights. To achieve efficient hybrid index construction, we develop a GPU-accelerated indexing pipeline featuring: (1) a warp-level hybrid distance computation kernel optimizing both dense and sparse operations parallelized on GPUs; (2) RNG-IP joint pruning that maintains search quality while reducing index complexity by combining Relative Neighborhood Graph (RNG) and Inner Product (IP) neighbor pruning; (3) keyword-aware neighbor recycling that preserves keyword search functionality by efficiently recycling the pruned neighbors to ensure keyword-based navigation on the index; and (4) logical edge augmentation that integrates the entity-level knowledge graph edges into the document-level hybrid index. To deliver a high-accuracy and high-throughput search service, we design a dynamic query framework on GPUs, incorporating: (1) dynamic heterogeneous edges loading for efficient traversal on the hybrid index; and (2) entity-document joint traversal for knowledge graph integration.

In summary, this paper makes the following contributions.

- We analyze the limitations of existing retrieval paradigms and derive a set of design principles for effective and flexible hybrid indexes (§ 2).
- We integrate multi-path retrieval into a unified semantic metric space, demonstrating its capability to handle fused distances with arbitrary weights. Based on this foundation, we design a hybrid index that supports any combination of retrieval paths without requiring reconstruction (§ 3).
- We propose an efficient GPU-accelerated framework for index construction and query processing. This framework enhances the inner product neighbors search, optimizes the hybrid distance calculation, and seamlessly integrates semantic and logical similarities (§ 4).
- We conduct comprehensive experiments on 6 real-world datasets, demonstrating that Allan-Poe achieves superior performance compared to state-of-the-art methods (§ 5).

The paper is organized as follows. Section 2 reviews the related work and illustrates the motivation. Section 3 introduces the structure of the hybrid graph-based index. Section 4 describes the index construction and query framework of Allan-Poe. Section 5 presents the experimental results. We conclude this paper in Section 6.

2 Background and Motivation

In this section, we first review related work on single-path and hybrid retrieval methods. We then establish the motivation for

¹For instance, Infinity [2] requires 5GB to store the indexes of three retrieval paths for a dataset with 1M documents, where the HNSW index size for dense vector retrieval only accounts for 23%.

designing an all-in-one hybrid index by analyzing the limitations of existing approaches.

2.1 Related Work

2.1.1 Single-path Retrieval. In the field of information retrieval, there are 4 distinct mainstream retrieval strategies.

(1) **Full-text search** is a lexical search method based on exact keyword matching. It evaluates the term importance through frequency-based models such as TF-IDF [64, 85, 87] and BM25 [8, 55, 69]. The inverted index is always used to achieve full-text search, employing retrieval algorithms such as WAND [43] and Block-Max WAND [58]. However, the exact term matching limits the recall of semantically relevant documents that lack the specific query keywords.

(2) **Sparse vector search** is another modern lexical approach that retrieves documents based on learned semantic representations of keywords. It utilizes models such as SPLADE [25] to encode documents into high-dimensional sparse vectors, where each dimension corresponds to the importance of a term from an expanded vocabulary. The inverted index and various pruning strategies are used to retrieve similar documents via vector similarity [11, 59]. While the sparse vector addresses the issues caused by exact term matching, it still lacks comprehensive semantic understanding.

(3) **Dense vector search** constitutes a semantic retrieval paradigm that employs deep language models like BERT [19] to generate dense vector representations capturing overall document semantics. Similarity is evaluated using hash-based [54, 97], tree-based [21, 96], or graph-based indexes [26, 57]. Despite its popularity, this approach is limited by embedding space constraints and the absence of explicit term matching, which can compromise retrieval accuracy.

(4) **Knowledge graph search** implements logical and semantic retrieval by converting queries into subgraphs and identifying similar structures through subgraph matching algorithms [37, 71, 95]. Subsequent enhancements incorporate entity and relation embeddings for improved efficiency [81, 99]. Recent approaches like GraphRAG leverage dense vector search and community detection for document retrieval of global questions [13, 22, 48, 100]. While GraphRAG excels at summarization tasks, it typically underperforms vanilla RAG for question answering [35]. Unlike GraphRAG, our work selectively integrates logical information from knowledge graphs to enhance vector search, establishing a distinct paradigm applicable to more general scenarios beyond RAG.

2.1.2 Hybrid Retrieval. Given the individual limitations of single-path retrieval methods, hybrid retrieval has emerged as a prominent search paradigm. Existing hybrid retrieval approaches can be categorized into two primary types.

(1) **Fusion retrieval** integrates multiple retrieval methods within a unified index structure. Current fusion methods are primarily restricted to two-path combinations [79]. For instance, DS-ANN [98] employs pre-defined fusion weight to combine dense and sparse vectors, constructing an HNSW index [57] for efficient querying. While efficient, this method requires complete index reconstruction if the fusion weights change. IVF-Fusion [10] addresses this by reducing the dimensionality of sparse vectors before combining them with dense vectors, then using an IVF index for retrieval. Although this eliminates weight-dependent reconstruction, it limits the flexibility to select different retrieval paths for varying scenarios.




Figure 2: Gaps between the vector similarity and end-to-end document similarity of two graph-based indexes.

(2) **Multi-path retrieval** represents a more flexible paradigm that executes searches separately across different indexes and subsequently fuses the results [1–3]. However, as discussed in Section 1, this flexibility adversely affects query efficiency and accuracy while complicating index management.

2.2 Motivation

2.2.1 Limitations of Single-path Search. Recently, dense vector search has become the most popular paradigm in vector databases among single-path retrieval methods, distinguished by its capacity for comprehensive semantic representation [65, 80]. Numerous specialized indexes have been developed to enhance their query efficiency and accuracy [29, 33, 68, 93]. However, vector similarity alone does not guarantee end-to-end relevance between queries and documents. To investigate this limitation empirically, we conduct experiments using two established real-world QA datasets: **NaturalQuestions (NQ)** [82] for simple question answering and **2WikiMultiHopQA (WM)** [36] for multi-hop question answering. For each dataset, we evaluate the first 1,000 queries, each associated with 1–3 ground-truth documents. We employ the BGE-M3 model [14] for the embedding of both documents and queries. The ground-truth baselines of vector similarity are established by computing the top-10 nearest neighbors via brute-force vector similarity search. We then assess the retrieval accuracy of two state-of-the-art graph-based indexes: HNSW (CPU-based) [57] and CAGRA (GPU-based) [62].

As shown in Figure 2, while vector-similarity accuracy can easily reach 99% within 4ms on the CPU and 0.1ms on the GPU, the corresponding end-to-end accuracy is substantially lower. In practical document retrieval systems, it is this end-to-end accuracy—not vector-similarity accuracy—that dictates performance in downstream tasks. Furthermore, Figure 2 reveals that the end-to-end accuracy for the WM dataset is lower than that for NQ, whereas their vector-based accuracies are comparable. This discrepancy underscores the limitations of dense vector search in handling complex queries. Consequently, reliance on any single-path retrieval method is insufficient for achieving satisfactory end-to-end performance, thereby restricting its utility in downstream applications.

2.2.2 Effectiveness of hybrid search. Hybrid search has emerged as a powerful strategy to mitigate the limitations of single-path retrieval and improve end-to-end accuracy [2, 79]. To evaluate its




Figure 3: Comparison of various retrieval paths using Infinity [2]. DVS, SVS, and FTS denote dense vector, sparse vector, and full-text search, respectively.




Figure 4: Example of retrieval in separate paths on NQ. The ground truth documents are doc₂ and doc₄.

effectiveness, we measure retrieval quality using Infinity [2], a modern database featuring efficient hybrid search. To better assess the quality of retrieved documents, we employ Normalized Discounted Cumulative Gain at rank k (nDCG@ k) [83] with $k = 5$, which evaluates both the recall and positional ranking of relevant documents in the retrieved results.

As shown in Figure 3, multi-path retrieval methods such as DVS+STS and DVS+STS+FTS generally exhibit higher nDCG than single-path retrieval. This demonstrates that multi-path retrieval can leverage the complementary strengths of individual paths to enhance result quality.

However, no single configuration is optimal for all scenarios. Retrieval with more paths does not consistently outperform fewer or single paths. For example, in Figure 3(a), the two-path combination of dense vector and full-text search yields lower nDCG than the single-path methods using either dense or sparse vectors alone on dataset NQ. And the three-path combination shows lower nDCG than the two-path retrieval of dense and sparse vector search. On other datasets, the accuracy ranking may be completely different. Furthermore, different path combinations present a distinct trade-off between accuracy and efficiency. Although the single-path methods typically achieve lower accuracy than the multi-path approaches, they are more efficient and incur less overhead (Figure 3(b)), sometimes making it preferable for a few real-time applications. For instance, the sparse vector retrieval method achieves comparable accuracy with the three-path combination while maintaining lower latency. These findings underscore the necessity for flexible path combination in hybrid search systems.




Figure 5: Trilemma of existing retrieval methods.

2.2.3 Limitations of Separate Multi-path Search. Multi-path retrieval is widely adopted for flexible hybrid search in modern vector databases [1–3]. Although it achieves superior accuracy compared to single-path approaches, this comes at the cost of increased time overhead. Moreover, this paradigm performs separate retrievals across individual indexes before fusing the results. The optimal results for a hybrid query may not be present within the top- k results from any single path or may be excluded after fusing the results [77, 78]. To illustrate this, we examine the retrieval of top-5 documents from the NQ dataset using dense and sparse vectors independently. Figure 4 presents a representative example. Using dense vectors, ground-truth documents doc₂ and doc₄ are ranked 2nd and 30th, respectively, while sparse vectors rank them 3rd and 12th. If we fuse the top-5 results from both paths using equal weights (i.e., $0.5 \times \text{dense similarity} + 0.5 \times \text{sparse similarity}$), doc₄ will be excluded because it only involves the similarity score from the sparse path. To include doc₄, more candidates should be included for each retrieval path (e.g., top-30). However, to ensure the accuracy of more candidates, more time is required for each retrieval path, leading to high end-to-end retrieval overhead.

2.3 Design Principles of Hybrid Index

The preceding analysis reveals that existing retrieval methods face a fundamental trilemma, being unable to simultaneously achieve high accuracy, efficiency, and flexibility. As illustrated in Figure 5, single-path search methods (e.g., HNSW [57], BM25 [69], and SPLADE [25]) achieve high efficiency but suffer from the semantic gap between vector similarity and end-to-end relevance, resulting in limited accuracy. Multi-path retrieval methods (e.g., Infinity [2] and Milvus [3]) enable flexible hybrid search through separate-then-fuse strategies but incur efficiency and accuracy costs, as illustrated previously. Fusion retrieval methods (e.g., IVF-Fusion [10] and DS-ANN [98]) improve accuracy over single-path approaches while maintaining intermediate efficiency. However, existing approaches are restricted to two-path combinations of dense and sparse vector search, limiting the potential accuracy gains. Moreover, as the combination of dense and sparse vectors is not always effective, they also face concerns about the adaptability to diverse scenarios.

Consequently, an effective hybrid index should satisfy the following three key requirements:

- **Flexibility:** Supporting arbitrary combinations of retrieval paths and fusion weights without index reconstruction to accommodate diverse application needs.

- **Accuracy:** Leveraging complementary information from multiple paths to maximize end-to-end relevance.
- **Efficiency:** Maintaining low-latency retrieval despite multiple paths and heterogeneous distance computations.

Guided by these principles, we propose Allan-Poe, which provides flexible integration of four retrieval paradigms via a well-designed isolated heterogeneous edges mechanism. Unlike separate-then-fuse approaches, Allan-Poe performs path fusion during query processing within a unified index, thereby avoiding the associated accuracy limitations. Additionally, we leverage massive GPU parallelism to achieve both efficient index construction and real-time query performance, meeting the requirement of efficiency.

3 Hybrid Index Structure

This section presents the structure of our proposed hybrid index, demonstrating how it resolves the trilemma by simultaneously achieving accuracy, efficiency, and flexibility.

3.1 Overview

The hybrid index of Allan-Poe integrates the dense vectors, sparse vectors, full-text, and knowledge graph retrieval within a Unified Semantic Metric Space (USMS).

DEFINITION 1 (UNIFIED SEMANTIC METRIC SPACE - USMS). A USMS is a tuple $H = \{D, F, M_w\}$ defined as follows:

- D : The set of all documents in the corpus.
- F : A set of feature extractors $\{f_{\text{dense}}, f_{\text{sparse}}, f_{\text{full}}, f_{\text{kg}}\}$ that maps each document $d \in D$ to its respective feature representation. For example, $f_{\text{dense}}(d) \in \mathbb{R}^m$ and $f_{\text{kg}}(d) \subseteq E$, where m is the dense vector dimension and E denotes the entity set.
- M_w : A composite similarity metric $M_w : D \times D \rightarrow \mathbb{R}$, defined for any weight vector $w = [w_d, w_s, w_f, w_k] \in \mathbb{R}^4$ as $M_w(q, d) = w_d \cdot \text{sim}_d(q, d) + w_s \cdot \text{sim}_s(q, d) + w_f \cdot \text{sim}_f(q, d) + w_k \cdot \text{sim}_k(q, d)$, where sim_d , sim_s , and sim_f denote the inner product similarities, and sim_k represents the path length between query and document entities in the knowledge graph.

Notably, sim_k employs path length as its measurement, contrasting with the inner product metrics used by other similarities in the composite similarity metric M_w . Furthermore, the knowledge graph retrieval operates on fine-grained entities, while other paths function at the document level, which reflects a fundamental granularity difference. Additionally, sim_k captures logical relationships, whereas the other represents semantic similarity. Direct integration of these dissimilar metrics could compromise the structural integrity of semantic edges. Consequently, we maintain separate semantic and logical edges within our graph-based index.

Figure 6 illustrates the hybrid index architecture. As shown in Figure 6(b), each node in the graph index (representing a corpus document) stores four data types corresponding to USMS features: dense vector, sparse vector, full-text vector, and entities. We classify the heterogeneous edges connecting these nodes into two categories based on the aforementioned incompatibility: semantic edges and logical edges. Semantic edges are further categorized into: (1) base semantic edges connecting nodes with similar fused vector semantics (detailed in Section 3.2), and (2) keyword edges connecting nodes sharing common keywords (detailed in Section 3.3). These




Figure 6: Overview of the hybrid index in Allan-Poe.

edges guide the traversal toward nearest neighbors in vector space. Logical edges, established from knowledge graph relations (detailed in Section 3.4), complement semantic edges by connecting nodes that are distant in vector space but logically related. The isolated heterogeneous edge storage guarantees the flexibility of Allan-Poe when dealing with any combination of retrieval paths.

3.2 Hybrid Vector Representation

To integrate dense vector, sparse vector, and full-text retrieval, we employ a vector fusion technique [77] that maps these representations into the USMS. Specifically, dense and sparse vectors are naturally represented in vector form, while keywords used in full-text search can be encoded as sparse vectors where each dimension represents the importance of a term in the vocabulary. For clarity, the sparse vectors generated by full-text search models such as BM25 [69] are denoted as statistical sparse vectors, while those produced by learning models such as SPLADE [25] are called learned sparse vectors. The vector fusion process concatenates these three vector types into a unified high-dimensional representation. Formally, for a document d , the concatenated vector is defined as $f_{\text{concat}}(d) = [f_{\text{dense}}(d), f_{\text{sparse}}(d), f_{\text{full}}(d)]$.

THEOREM 1. Given the RNG-based index constructed from the fused vectors $f_{\text{concat}}(d)$ in UMUS, for any weight vector $w = [w_d, w_s, w_f] \in \mathbb{R}^3$ applied to query vectors $\{f_{\text{dense}}(q), f_{\text{sparse}}(q), f_{\text{full}}(q)\}$, the nearest neighbors can always be retrieved from the index.

PROOF. We construct a weighted query vector by concatenating the component vectors with their respective weights: $f_{\text{concat}}(q) = [w_d * f_{\text{dense}}(q), w_s * f_{\text{sparse}}(q), w_f * f_{\text{full}}(q)]$. With the Relative Neighborhood Graph (RNG) [73] as the index, the nearest neighbors of the vector $f_{\text{concat}}(q)$ can always be found using the greedy search algorithm [26, 80, 93]. Let $f_{\text{concat}}(d^*)$ be one of the nearest neighbors. The inner product between the query vector $f_{\text{concat}}(q)$ and candidate vector $f_{\text{concat}}(d^*)$ expands as $[w_d * f_{\text{dense}}(q), w_s * f_{\text{sparse}}(q), w_f * f_{\text{full}}(q)] \cdot [f_{\text{dense}}(d^*), f_{\text{sparse}}(d^*), f_{\text{full}}(d^*)] = w_d * f_{\text{dense}}(q) \cdot f_{\text{dense}}(d^*) + w_s * f_{\text{sparse}}(q) \cdot f_{\text{sparse}}(d^*) + w_f * f_{\text{full}}(q) \cdot f_{\text{full}}(d^*)$, which corresponds exactly to the weighted combination of similarities from the three retrieval paths. Consequently, for any weight vector $w = [w_d, w_s, w_f] \in \mathbb{R}^3$, the nearest neighbors of the three types of vectors can always be retrieved. $\square$

Theorem 1 establishes that the index supports arbitrary weight combinations without requiring reconstruction. As the overhead of constructing an exact RNG is significant, we construct an approximate RNG utilizing the strategy in NGT [38] and CAGRA [62], which has been proven effective and efficient for approximate neighbor retrieval on GPUs [62].

3.3 Keyword Edges Supplement

While the vector fusion approach described in Section 3.2 enables flexible hybrid retrieval within a unified index, it inherently compromises the keyword-based search function in the original full-text search. In many applications, users explicitly require certain keywords to appear in retrieved documents to enhance accuracy [46, 52]. Although the function of keyword search can be easily achieved using the traditional inverted index for full-text search, it is non-trivial for graph-based indexes. Recent research has explored graph-based indexes with attribute filtering capabilities [6, 31, 67], which can be employed to achieve keyword search in our hybrid index. However, these approaches typically integrate attribute constraints directly into the primary graph structure, limiting flexibility. To address this limitation, we propose a dual-assessment mechanism that selectively preserves pruned edges as dedicated keyword edges.

During construction of the hybrid index, we first prune the graph using the strategy in CAGRA [62] to leverage GPU computational power. Recall that the pruning strategy used in CAGRA prunes edges according to the number of detourable routes, where edges with more detourable routes will be pruned. Specifically, for a node A and its neighbors X , if there exists another neighbor Y such that $\max(\text{dis}(A, X), \text{dis}(X, Y)) < \text{dis}(A, Y)$, then the path $A \rightarrow X \rightarrow Y$ constitutes a detourable route for the edge $A \rightarrow Y$. CAGRA retains the d neighbors with the fewest such detourable routes. Our dual-assessment mechanism operates during the above pruning phase of CAGRA by evaluating keyword overlap between nodes. For each neighbor X of node A that would normally be pruned by CAGRA, if there exists another neighbor Y of A such that $K(A) \cap K(X) \subseteq K(Y)$, then X is pruned. This is because the navigation from A to X can be replaced by Y to X for any keywords. However, if a neighbor scheduled for pruning does not satisfy this condition, it is preserved as a keyword edge. As depicted in Figure 6 previously, keyword edges are maintained separately from base semantic edges to ensure clear separation. This distinct edge organization guarantees the pluggable nature of keyword functionality. The incorporation of keyword edges facilitates efficient traversal to semantically relevant neighbors sharing common keywords, significantly enhancing keyword-aware search performance.

3.4 Logical Edges Augmentation

While previous sections integrated dense vector, sparse vector, and full-text search through semantic edges in our hybrid index, semantic-based graph search still faces two fundamental challenges: (1) *Semantic search retrieves semantically similar but logically unrelated documents.* For example, for a query "Where was John's mother born?", two documents containing "John's father was born in the US" and "Linda's mother was born in the US" can be retrieved, as they exhibit high semantic similarity due to the similar keywords "John", "mother", or "born" despite describing logically distinct relationships. (2) *Semantic search struggles with complex queries involving multiple entities or multi-hop reasoning.* The query "Who is younger, Linda or John?" contains multiple entities, often causing graph traversal to settle in local optima and retrieve documents about only one entity. Similarly, for the multi-hop query "Where was Linda's mother born?", if information about Linda and her mother is distributed across different documents, semantic search only returns

documents about Linda, missing crucial contextual information. To address these limitations, Allan-Poe augments semantic search with logical edges utilizing knowledge graphs.

Knowledge graphs can be constructed from the corpus using deep language models such as BERT [19] or LLMs [27]. For each node in our hybrid index, we store associated entities alongside the fused vector and maintain an entity-to-node mapping. We then extract inter-entity relations from the knowledge graph and represent them as logical edges. Formally, let $V(X)$ denote the entity set of node X , and $G(V, R)$ represent the knowledge graph with entities V and relations R . The logical edges for node X comprise triplets $\{(s, r, t) \mid s \in V(X), r \in R, t \in V \setminus V(X)\}$, where s and t are entities connected by relation r . Thus, any two entities from different nodes that are related in the knowledge graph establish a logical edge between their corresponding document nodes. During search, we dynamically leverage logical edges through a fine-grained entity-document unified strategy that enhances query capability while preserving efficiency (detailed in Section 4.2). Notably, logical edge augmentation is optional, representing a trade-off between potential accuracy improvements and the substantial computational cost of knowledge graph construction [90].

4 Hybrid Index Construction and Query

Section 3 presents the basic structure of the hybrid index in Allan-Poe. The integration of multiple retrieval paths in the index structure introduces significant complexity to both construction and query processing, creating substantial efficiency challenges. To address these issues, this section describes our approach to efficient hybrid index construction and high-performance retrieval leveraging GPU acceleration.

4.1 Efficient Index Construction on GPU

Allan-Poe's hybrid index construction presents additional complexity due to the challenges brought by hybrid distance computation and heterogeneous edge establishment. This substantial overhead motivates our use of GPUs to accelerate the construction process. As illustrated in Figure 7, the Allan-Poe indexing pipeline comprises four key procedures.

Procedure 1: Initial k -NN Graph Construction with Hybrid Distance. We employ the NN-Descent algorithm [20] to construct an approximate k -nearest neighbor (k -NN) graph from the fused vectors. NN-Descent operates on the principle that "a neighbor's neighbors are likely neighbors"—it iteratively refines the graph by evaluating 2-hop neighbors and updating connections based on distance comparisons. This approach has demonstrated superior efficiency to incremental methods on GPUs [50, 75]. Within this framework, hybrid distance calculation between fused vectors represents the most computationally intensive operation, involving two distinct computation patterns. To accelerate this process on GPUs, we assign an entire warp to collaboratively compute each distance. For dense vector computation, each thread fetches four operands using CUDA vectorized instructions [28], maximizing memory bandwidth utilization. Threads then multiply their operands in parallel, storing intermediate results in registers. After processing all dense dimensions, the algorithm computes distances for both learned and statistical sparse vectors using identical processing. Sparse vectors




Figure 7: Index construction process in Allan-Poe.

are stored in CSR format [34], with non-zero values and indices maintained in separate arrays. Distance calculation is transformed into a parallel set intersection: each thread fetches an index from the candidate’s sparse vector and searches for it in the explored node’s sparse vector using parallel binary search. This design enables the frequently used explored nodes’ vectors to be cached in shared memory, minimizing access overhead. When matches occur, the corresponding values are multiplied and accumulated with the dense vector results. Finally, per-thread values are aggregated using warp-level reduction operations in CUDA.

Procedure 2: RNG-IP Joint Pruning. For the inner product metric, vectors with large norms (lengths) frequently appear in maximum inner product search results, making connections to such vectors an effective strategy for improving search efficiency [72]. Recent approaches either apply inner product (IP) pruning directly or augment RNG-pruned edges with high-norm nodes in graph indexes [16, 17]. However, these methods increase index size and often create uneven node degrees, which hinders aligned parallel processing on GPUs. To overcome these limitations, we combine RNG and IP pruning strategies with GPU-specific optimizations, achieving efficient inner product search while preserving both the original index size and aligned degree structure. Our joint pruning approach operates in two phases. The first phase employs RNG pruning from CAGRA to refine neighbors and reduce search complexity. As shown in Figure 7, we calculate the detourable routes for neighbors in the k -NN graph in parallel and sort neighbors by their detourable route counts. The second phase applies the IP pruning strategy² to remove neighbors with small vector norms. To efficiently achieve the pruning, we parallelize the inner product calculations between the candidate neighbor and current neighbors (e.g., between candidate G and neighbors $X \in \{B, C, D, E, F\}$ in Figure 7), where each warp is responsible for the distance calculation between a current neighbor and the candidate neighbor. Finally, $d/2$ neighbors and $d/2$ reverse neighbors are concatenated to form each node’s final edge list, where d represents the target index degree.

Procedure 3: Keyword-aware Neighbor Recycling. This procedure recycles pruned edges from Procedure 2 according to the strategy established in Section 3.3. A brute-force implementation would incur substantial computational overhead, as checking the condition $K(A) \cap K(X) \subseteq K(Y)$ for each node A and its pruned neighbors X against all current neighbors Y requires numerous set intersection operations. To optimize this process, we leverage computations already performed in previous procedures. Specifically, we assign

a Boolean keyword flag to each neighbor of each node. During the second phase of Procedure 2, while computing statistical sparse vector similarities via set intersection between neighbors of node A , we check whether non-intersecting keywords exist in A ’s keyword set $K(A)$. If so, we set the corresponding neighbor’s keyword flag to 1. For efficient $K(A)$ lookups, we maintain it in a hash map within GPU shared memory. For example, in Figure 7, node E is scheduled for pruning during Procedure 2 due to its numerous detourable routes. During the second phase’s distance calculations between E and current neighbors (B, C, D) in Procedure 2, we examine non-intersecting keywords from E (e.g., $K(E) \setminus K(B) = \{k_1, k_4\}$) and verify their presence in $K(A)$ using the hash map. Here, k_4 exists in $K(A)$ but not in any current neighbor, violating the subset condition. Consequently, E ’s keyword flag is set to 1. Finally, we traverse all pruned neighbors and recycle those with activated keyword flags as keyword edges.

Procedure 4: Logical Edge Augmentation. In this final procedure, we establish logical edges by mapping knowledge graph entities to their corresponding document nodes in the graph index and creating connections between nodes whose entities share relationships in the knowledge graph.

Algorithm 1 summarizes the complete index construction pipeline. First, the initial k -NN graph is constructed (lines 1-4). Subsequently, for each node in the graph, the neighbors are sorted by their detourable route counts (lines 6-7), and the IP pruning strategy is applied to filter neighbors (lines 10-13). During IP pruning distance calculations, we simultaneously set keyword flags to enable neighbor recycling (lines 14-15). Finally, we augment the graph index with logical edges extracted from the knowledge graph (line 18).

Updates of the Hybrid Index. For the insertion of new nodes, the k nearest neighbors of each newly inserted node are determined by merging two candidate sets: (1) the k -NN retrieved from the existing index using base semantic edges, and (2) the k -NN identified by performing NN-Descent among the newly inserted nodes. The k -nearest neighbors of each new node are transmitted to the subsequent procedures (pruning and edge augmentation), which remain identical to the initial construction. For node deletion, Allan-Poe adopts the mark-deletion strategy where removed nodes remain in the index during search but are filtered from final results.

4.2 Flexible Query Processing on GPU

Given the hybrid index constructed in Section 4.1, the query algorithm needs to be carefully designed to enable efficient retrieval across flexible path combinations while maintaining performance. This subsection presents our GPU-accelerated query processing

²For node A with neighbor set R , candidate G is excluded if $\exists X \in R$ such that $IP(G, X) > IP(G, G)$. Similarly, X is filtered if $IP(G, X) > IP(X, X)$ [72].

Algorithm 1: Construction of the Hybrid Index

Input: fused vector data V , knowledge graph KG , numbers of iterations it for NN-Descent, degree d
Output: graph-based hybrid index G

```
1  $G \leftarrow$  Randomly initialize  $k$  neighbors for each node;  
2 for ( $curIt \leftarrow 0$ ;  $curIt < it$ ;  $curIt++$ ) do  
3   foreach  $u \in G$  in parallel do  
4     Explore  $u$ 's 2-hop neighbors and update  $N(u)$ ;  
5 foreach  $u \in G$  in parallel do  
6   Calculate detourable routes for  $v \in N(u)$ ;  
7   Sort  $N(u)$  according to the number of detourable routes;  
8    $SE \leftarrow$  the first node in  $N(u)$ ,  $KE \leftarrow \emptyset$ ;  
9   foreach  $v \in N(u)$  do  
10    Calculate inner product between  $v$  and nodes in  $SE$ ;  
11    Set  $v.keywordFlag$  based on the intersection results;  
12    if  $|SE| < d \wedge \forall w \in SE \text{ s.t. } IP(w, v) < IP(v, v)$  then  
13       $SE \leftarrow SE \cup \{v\}$ ;  
14    else if  $v.keywordFlag = 1$  then  
15       $KE \leftarrow KE \cup \{v\}$ ;  
16     $SE \leftarrow d/2$  nodes in  $SE$  and  $d/2$  reverse neighbors;  
17     $G[u].semanticEdge \leftarrow SE$ ,  $G[u].keywordEdge \leftarrow KE$ ;  
18     $G[u].logicalEdge \leftarrow$  extend based on  $KG$ ;  
19 return  $G$ 
```

algorithm, which efficiently handles keyword and knowledge graph augmentations without compromising query latency.

4.2.1 Query on the Base Semantic Edges. As established in Section 3.2, base semantic edges constructed from fused vectors support arbitrary weight combinations across retrieval paths. Given a weight vector $w = [w_d, w_s, w_f]$ for dense, learned sparse, and statistical sparse vectors respectively, the query vector is formulated as $f_{concat}(q, w_d, w_s, w_f) = [w_d * f_{dense}(q), w_s * f_{sparse}(q), w_f * f_{full}(q)]$. Single-path retrieval is achieved by setting the corresponding weight to 1 (or any non-zero value) and others to 0. For instance, the fused query vector $f_{concat}(q, 1, 0, 0)$ retrieves documents using only dense vector similarity. This approach extends naturally to two-path and three-path configurations. To optimize search efficiency, we select entry points from nodes with the smallest vector norms. The computationally intensive hybrid distance calculations during query processing are accelerated using the same GPU-optimized strategy described in Section 4.1. The only difference is that during the set intersection operation, each thread fetches an index from the document's sparse vector and searches for it in the query's sparse vector. This design reduces time complexity by searching the typically smaller query vector.

4.2.2 Query with Keyword Augmentation. Allan-Poe enables users to specify required keywords in queries, ensuring retrieved documents contain these terms. However, loading keyword edges during every node traversal would incur substantial overhead. To address this, we implement dynamic keyword edge loading: when expanding a node's neighbors into the candidate pool, we check for keyword commonality (already computed during distance calculation) and load keyword edges only for nodes sharing keywords with

Algorithm 2: Query with the Hybrid Index

Input: fused vector data V , graph-based hybrid index G , query q , specified keyword set KW , specified entity set E , multi-path weights $[w_d, w_s, w_f, w_k]$
Output: q 's approximate k nearest neighbors

```
1 Generate the query vector  $q_v = f_{concat}(q, w_d, w_s, w_f)$ ;  
/* Initialize the entry points */  
2 if query with knowledge graph then  
3    $cand \leftarrow$  nodes containing user-specified entities;  
4   for  $v \in cand$  do  
5      $v.hop = 0$ ; // initial entities in the query  
6      $v.ent = e$ ; //  $e \in E$   
7 else  
8    $cand \leftarrow$  nodes with small vector length;  
9 Calculate  $dis(q_v, v)$  where  $v \in cand$  in parallel;  
/* Begin to search */  
10 while  $|cand| > 0$  do  
11    $u \leftarrow$  the nearest unvisited neighbor to  $q_v$  in  $cand$ ;  
12    $z \leftarrow$  the furthest neighbor to  $q_v$  in  $topk$ ;  
13    $N(u) \leftarrow G[u].semanticEdge$ ;  
14   if  $K(u) \cap KW \neq \emptyset$  then  
15      $N(u) \leftarrow N(u) \cup G[u].keywordEdge$ ;  
16   if  $u.ent \neq none$  then  
17      $N(u) \leftarrow N(u) \cup G[u].logicalEdge$ ;  
18   for unvisited  $o \in N(u)$  in parallel do  
19     if  $u.ent \neq none$  then  
20        $o.ent \leftarrow$  entity in  $o$  having relations with  $u.ent$ ;  
21       if  $o.ent \neq none$  then  
22          $o.hop \leftarrow u.hop + 1$ ;  
23          $dis(o, q_v) \leftarrow dis(o, q) - w_k/o.hop$ ;  
24       if  $dis(o, q_v) < dis(z, q_v)$  then  
25          $cand \leftarrow cand \cup \{o\}$ , push  $o$  to  $topk$ ;  
26       if  $|topk| > k$  then  
27         pop the furthest node  $m$  from  $topk$ ;  
28         push  $m$  to  $kwCand$  if  $K(m) \cap KW \neq \emptyset$ ;  
29 while  $|res| < k$  do  
30   push  $s \in topk \cup kwCand$  s.t.  $K(s) \cap KW \neq \emptyset$  to  $res$ ;  
31 return  $res$ 
```

the query. Crucially, we do not restrict traversal exclusively to keyword-matched nodes, as this would impair accuracy by excluding potential pathway nodes. Instead, we employ a twin candidate pool approach [6], maintaining a secondary pool for keyword-satisfying nodes excluded from the primary pool due to larger distances. Upon query completion, we merge both pools and filter for nodes containing the required keywords.

4.2.3 Query with Knowledge Graph Augmentation. Allan-Poe further enables users to specify key entities in queries to enhance retrieval through logical similarity. As discussed in Section 3.4, logical edges address two key challenges: (1) *complex queries with multiple entities or multi-hop*, and (2) *semantically similar but logically unrelated results*. To mitigate local optima in multi-entity queries, we employ entity-based entry point selection, choosing

nodes containing user-specified entities as initial entry points via the entity-node mapping. For query processing augmented by logical edges, the core principle is that nodes containing entities related to user-specified entities exhibit higher logical similarity, which should reduce their effective hybrid distance. Based on this, during the query process, for each explored node, we first expand the candidate pool using base semantic edges from the current node. If that node is within x hops of user-specified entities in the knowledge graph, we additionally expand the candidate pool using its logical edges. It's worth noting that not all the logical edges of this node are loaded, but only those edges whose source entities are within x hops of the target entities. Each neighbor expanded via logical edges is annotated with its hop distance from the query entities, thereby avoiding the need to recalculate the hop distance from scratch. Furthermore, we verify whether candidates expanded via base semantic edges are knowledge graph neighbors of entities in the explored node. We incorporate logical similarity by rewarding nodes based on their hop distance from query entities: fewer hops yield greater reward (i.e., reduced effective distance). This approach integrates fine-grained entity relations into the document-level graph search, effectively addressing both logically unrelated results and multi-hop query challenges.

Algorithm 2 presents the pseudo-code for Allan-Poe's query processing. The algorithm begins by fusing retrieval vectors (line 1) and initializing the candidate pool with path-appropriate entry points (lines 2-9). During search, the neighbor list is initialized with base semantic edges (line 13), while keyword and logical edges are dynamically loaded based on the current node (lines 14-17). For each unvisited neighbor, we compute its distance from the query vector and incorporate logical similarity (lines 19-23), then expand the candidate pool accordingly (lines 24-28). Finally, results are filtered to ensure they contain the queried keywords (lines 29-30).

5 Experiments

In this section, we conduct comprehensive experiments to evaluate the performance of Allan-Poe and compare it with existing state-of-the-art retrieval methods.

5.1 Experiment Settings

5.1.1 Datasets. For comprehensive evaluations, we use 6 real-world datasets of varying scales, which have been widely used in related works [10, 41, 79, 98]. Among them, **NaturalQuestions** (NQ) [82] and **MS MARCO** (MS) [61] include simple queries, while **2WikiMultiHopQA** (WM) [55] and **HotpotQA** (HP) [92] contain complex, multi-hop queries. Table 1 summarizes the detailed information of each dataset. We employ the BGE-M3 model [14] to generate the dense vectors with a dimension of 1024, the SPLADE model [25] to generate the sparse vectors, and the BM25 algorithm [69] to generate the full-text vectors.

5.1.2 Methods. We evaluate our proposed Allan-Poe against 6 state-of-the-art competitors representing both hybrid and single-path retrieval paradigms:

- **SEISMIC** [11] is the state-of-the-art method supporting only the sparse vector search.

Table 1: Statistics of Datasets. "D. Dim", "S. Dim", and "F. Dim" denote the dimensions of the dense, sparse, and full-text vectors, respectively.

Dataset	#Corpus	#Queries	D. Dim	S. Dim	F. Dim
NQ [82]	1,000,000	1,000	1,024	30,522	852,356
MS [61]	1,000,000	1,000	1,024	30,522	831,592
WM [55]	414,743	1,000	1,024	30,522	529,931
HP [92]	509,176	1,000	1,024	30,522	699,002
NQ-9633 [82]	9,633	100	1,024	30,522	42,834
WM-6119 [55]	6,119	100	1,024	30,522	33,357

- **CAGRA** [62] is the state-of-the-art GPU-based method supporting only the dense vector search.
- **IVF-Fusion** [10] is a hybrid search method adopting the fusion retrieval paradigm. It combines dense vectors and sparse vectors using the Johnson-Lindenstrauss (JL) transformation [42] and utilizes an inverted index to accelerate the search. We implemented a GPU version of IVF-Fusion for fair comparison.
- **Infinity** [2] is a modern database featuring with efficient hybrid search. It adopts the multi-path retrieval paradigm, which supports combinations of dense vector, sparse vector, and full-text search.
- **ThreeRouteGPU** is our implemented GPU-based hybrid search method adopting the multi-path retrieval paradigm. It constructs separate CAGRA indexes for dense vectors, learned sparse vectors, and statistical sparse vectors (reducing dimension via JL transformation). Results retrieved from the three routes are then fused using fixed weights.
- **HippoRAG** [41] is one of the state-of-the-art GraphRAG algorithms integrating knowledge graph and dense vectors to enhance the retrieval of relevant documents. Due to high knowledge graph construction costs, we compare it with All-Poe using logical edges only on the smaller datasets NQ-9633 and WM-6119.
- **Allan-Poe** is our proposed method adopting the fusion retrieval paradigm. We denote different retrieval configurations as: dense for dense vectors only, sparse for sparse vectors only, full for full-text only, TwoPath for dense+sparse combination, and ThreePath for all three paths combined.

5.1.3 Metrics. We evaluate the indexing efficiency by measuring the construction time, the retrieval efficiency by *Queries Per Second* (QPS), and the retrieval accuracy by $nDCG@k$. Without additional explanation, the value of k is set to 10.

5.1.4 Platforms. All experiments are conducted on a server featuring an Intel Xeon Silver 4310 CPU@2.10GHz, 125GB RAM, and a Nvidia GeForce RTX 3090 GPU (24G). We implement Allan-Poe in C++/CUDA under CUDA 12.2.

5.2 Overall Performance for Query Processing

In this section, we compare the QPS and $nDCG@10$ across all methods on four real-world datasets. All the compared hybrid search methods employ equal weighting for all retrieval paths unless specified otherwise. The CPU-based methods, i.e., SEISMIC and Infinity, utilize 48 threads during query processing. The results are depicted




Figure 8: Comparison of all methods on 4 real-world datasets.

in Figure 8. For Infinity and SEISMIC, we report single data points since nDCG values show minimal variation with parameter tuning.

Experimental results shown in Figure 8 demonstrate that the corresponding retrieval configurations of Allan-Poe outperform SEISMIC, CAGRA, IVF-Fusion, Infinity, and ThreeRouteGPU by 4.4-13.6 $\times$, 1.5-2.5 $\times$, 17.6-41.8 $\times$, 27.1-186.4 $\times$, and 1.6-4.8 $\times$, respectively. On datasets *WM* and *HP*, Allan-Poe-ThreePath achieves the highest accuracy (nDCG@10). Despite requiring more distance calculations, Allan-Poe-ThreePath maintains higher nDCG@10 than other methods at equivalent QPS levels on these datasets, demonstrating the benefit of complementary information from multiple retrieval paths. From another perspective, Allan-Poe-dense and CAGRA solely employ dense vectors to retrieve documents, which consistently underperform the hybrid search paradigm across most datasets in nDCG@10 regardless of the QPS. For example, on *MS*, the nDCG@10 of Allan-Poe-dense and CAGRA reach only 0.5, compared to Allan-Poe-ThreePath’s nDCG@10 of 0.56. This confirms that the retrieval path using dense vectors alone is insufficient for end-to-end retrieval due to information loss in embedding models. Consequently, introducing more retrieval paths to complement the lost information of dense vectors is a promising way to enhance the end-to-end query efficiency.

However, as noted in Section 2.2, additional retrieval paths do not guarantee improved accuracy in all scenarios. For instance, on datasets *NQ* and *MS*, Allan-Poe-TwoPath and Allan-Poe-sparse achieve optimal performance, respectively. This occurs because all participating paths influence final accuracy—on *NQ* and *MS*, while dense and sparse retrieval perform well, full-text retrieval underperforms and reduces overall accuracy. Moreover, for dataset *MS*, Allan-Poe-sparse using a single retrieval path outperforms the hybrid retrieval methods because all the documents in *MS* have a short length with simple semantics, which can be efficiently handled by sparse vectors. These results emphasize the importance of supporting flexible path combinations without index reconstruction, corresponding to the flexibility dimension in Figure 5.

For different retrieval paradigms adopted by existing methods (fusion retrieval and separate multi-path retrieval), the fusion retrieval paradigm (represented by Allan-Poe-ThreePath) outperforms the separate multi-path retrieval paradigm (represented by ThreeRouteGPU and Infinity) on *MS* and *HP* in nDCG@10 regardless of the QPS, while achieving comparable nDCG@10 on *NQ* and *WM*. As discussed in Section 2.2, separate multi-path retrieval can miss relevant documents, reducing accuracy, which underscores the advantage of fusion retrieval.

Table 2: Comparison of index build time and index size.

Methods	Build Time (s)				Index Size (MB)			
	<i>NQ</i>	<i>MS</i>	<i>WM</i>	<i>HP</i>	<i>NQ</i>	<i>MS</i>	<i>WM</i>	<i>HP</i>
SEISMIC	98.09	114.34	50.83	87.38	2921	1993	1526	1904
CAGRA	16.29	17.62	7.45	8.96	126	126	52	64
IVF-Fusion	2.63	2.42	1.77	1.41	136	136	131	134
Infinity	487.04	440.93	186.54	263.19	5738	4541	1962	2701
ThreeRouteGPU	49.20	48.89	22.56	27.18	378	378	156	192
Allan-Poe	40.08	36.02	19.50	24.25	186	186	78	95

In terms of efficiency, the different configurations of Allan-Poe achieve the highest QPS among all methods, even with multiple retrieval paths. This performance stems from our GPU optimizations, particularly the hybrid distance calculation that addresses the primary retrieval overhead.

5.3 Evaluations of Indexing Efficiency

In Table 2, we report the index construction time as well as the index size across all the compared methods. Among three-path retrieval methods, Allan-Poe achieves the fastest build time while maintaining a compact index size. Specifically, compared to the GPU-based method ThreeRouteGPU, Allan-Poe demonstrates 1.2 $\times$ faster construction and 2.0 $\times$ smaller index size. These advantages are even more pronounced against Infinity, with 11.2 $\times$ faster build time and 21.0 $\times$ reduction in index size. These results validate the effectiveness of our GPU-accelerated construction optimizations in Section 4.1, which includes the hybrid distance acceleration and the parallel pruning implementation for heterogeneous edges. The compact index size further demonstrates the effectiveness of our unified design, which significantly reduces storage overhead and system complexity compared to separate index paradigms. Among all methods, IVF-Fusion achieves the fastest construction due to its simple inverted index structure, but suffers from consistently poor accuracy. CAGRA maintains a small index by supporting only single-path retrieval for dense vectors, but its retrieval performance lags behind hybrid approaches.

5.4 Evaluations of Keyword Retrieval

To evaluate the effectiveness of keyword specification in queries, we employ the Qwen3 LLM [88] to simulate users specifying required keywords for the first 100 queries from the representative datasets *NQ* and *WM*. Figure 9 compares results with and without




Figure 9: Comparison of retrieval methods w/o keywords.

Table 3: Comparison of retrieval methods on small datasets. ‘kg’ denotes the methods augmented by the knowledge graph.

Methods	NQ-9633		WM-6119	
	nDCG@10	QPS	nDCG@10	QPS
SEISIMC	0.732	838.10	0.765	659.40
CAGRA	0.732	1254.82	0.761	1070.75
IVF-Fusion	0.740	4030.35	0.780	3316.88
Infinity	0.749	211.43	0.741	232.37
ThreeRouteGPU	0.736	1167.36	0.765	1085.5
HippoRAG	0.747	1.28	0.805	1.04
Allan-Poe-full	0.605	14716.81	0.699	6306.72
Allan-Poe-full (kg)	0.666	9212.35	0.817	4572.75
Allan-Poe-dense	0.728	13581.5	0.759	9487.67
Allan-Poe-dense (kg)	0.738	9428.36	0.818	5941.49
Allan-Poe-sparse	0.730	13303.75	0.759	6553.34
Allan-Poe-sparse (kg)	0.744	8318.01	0.832	5240.93
Allan-Poe-ThreePath	0.751	12098.23	0.770	5508.97
Allan-Poe-ThreePath (kg)	0.751	8520.28	0.834	4126.26

keyword supplementation, reporting the highest nDCG@10 and corresponding QPS for each method. Keyword supplementation improves nDCG@10 by 1.2% on average across retrieval paths, with only a 3.2% QPS reduction, demonstrating its effectiveness. However, some methods show minimal accuracy gains (e.g., 0.2% for three-path retrieval on NQ), as many relevant documents already contain the required keywords, limiting the filtering impact. Interestingly, certain retrieval paths with keyword supplementation achieve higher QPS than their non-supplemented counterparts (e.g., full-text search on NQ and dense vector search on WM). This suggests that keyword constraints enable earlier convergence to optimal accuracy by focusing the search on more relevant nodes within a small candidate pool in query process.

5.5 Evaluations of Logical Augmentation

To evaluate knowledge graph augmentation while managing the knowledge graph construction costs, we use Qwen3 to construct knowledge graphs for the two smaller datasets NQ-9633 and WM-6119, and simulate user queries by specifying required entities for




Figure 10: Performance w/o RNG-IP joint pruning.

the first 100 queries. Table 3 presents the evaluation results. Overall, Allan-Poe-ThreePath (kg) exhibits the highest nDCG@10 while maintaining competitive QPS. On NQ-9633, which contains simple queries without multi-hop reasoning, knowledge graph augmentation provides modest accuracy improvements—particularly for single-path Allan-Poe variants, with slight gains for the three-path configuration. In contrast, the accuracy improvements on WM-6119 are significant due to the complex multi-hop queries in this dataset, demonstrating the effectiveness of the knowledge graph integration in Allan-Poe. While HippoRAG (a GraphRAG method) achieves higher nDCG@10 than approaches without knowledge graphs, it still underperforms compared to Allan-Poe. This gap occurs because HippoRAG does not effectively integrate knowledge graph information with document-level semantic similarity, and its community search can introduce redundant documents that impair query efficiency despite being useful for global queries. Consequently, integrating knowledge graphs with document-level vector search represents a promising direction for future research.

5.6 Effectiveness of Heterogeneous Edges

5.6.1 Effectiveness of RNG-IP Joint Pruning. Figure 10 compares the performance of Allan-Poe-ThreePath with and without the RNG-IP joint pruning strategy on two representative datasets. The joint RNG-IP pruning strategy improves both retrieval efficiency and accuracy compared to using RNG pruning alone, demonstrating its effectiveness in enhancing index quality. Notably, while the distance metric is Inner Product, using IP pruning alone (without RNG) yields lower performance than RNG pruning alone because IP pruning eliminates fewer candidate neighbors, providing limited reduction in search computation cost.

5.6.2 Effectiveness of Keyword Edges. Figure 11 compares Allan-Poe with three retrieval paths and its full-text search configuration with and without keyword edges. As discussed in Section 3.3, the introduction of keyword edges is to restore keyword-based retrieval capability lost in full-text search during vector fusion. As shown in Figure 11, keyword edges improve nDCG@10 for full-text search by 1% and 4% on NQ and WM, respectively. These improvements extend to three-path search on WM, but are less pronounced on NQ, where baseline full-text search accuracy is substantially lower than other retrieval paths.

5.6.3 Effectiveness of Logical Edges. Without logical edges, Allan-Poe achieves nDCG@10 of 0.655 (full-text), 0.737 (dense), 0.734 (sparse), and 0.751 (three-path) on NQ-9633, and 0.727, 0.746, 0.739,




Figure 11: Performance w/o keyword edges.




Figure 12: Performance of various weights for retrieval paths.

and 0.760 on WM-6119 according to our experiments. Compared to the results in Table 3, these values show significant degradation, indicating that logical edges effectively compensate for semantic edge limitations by enhancing query-document relevance.

5.7 Weights of Retrieval Paths

To investigate the impact of retrieval path weighting, we evaluate Allan-Poe-TwoPath and Allan-Poe-ThreePath under various weight configurations. For the two-path configuration (dense + sparse vectors), the fused distance is computed as $\alpha \cdot \text{sim}_d(q, d) + (1 - \alpha) \cdot \text{sim}_s(q, d)$, where $\alpha \in [0, 1]$, and q, d denote query and document respectively. For the three-path configuration, the distance function is $\alpha \cdot [\text{sim}_d(q, d) + w_{\text{opt}} \cdot \text{sim}_s(q, d)] + (1 - \alpha) \cdot \text{sim}_f(q, d)$, where w_{opt} represents the optimal dense-sparse weight derived from Allan-Poe-TwoPath evaluations. Results are presented in Figure 12.

Optimal weights correlate strongly with individual path performance: higher-accuracy paths warrant greater weighting to achieve overall high accuracy. For instance, on the NQ dataset, sole dense vector retrieval ($\alpha = 1$ for the line of dense+sparse) achieves higher nDCG@10 than sparse retrieval ($\alpha = 0$ for the line of dense+sparse), resulting in an optimal $\alpha = 0.7$ that favors dense vectors. Similarly, since dense+sparse retrieval substantially outperforms full-text search on NQ, the optimal three-path configuration allocates 0.9 weight to dense+sparse and 0.1 to full-text. The evaluation results in Figure 12 also demonstrate that three-path retrieval can surpass or at least have comparable accuracy with two-path retrieval under




Figure 13: Comparison of inserting various data volumes.

Table 4: Comparison of indexing overhead.

Datasets	Rebuild	Insert 20%	Insert 10%	Insert 5%
NQ	40.08s	5.82s	2.86s	1.40s
WM	19.50s	2.20s	1.06s	0.51s

appropriate weight selection. Based on these findings, we derive an empirical weighting criterion based on the nDCG gap of two paths:

- **nDCG gap < 5%:** Equal weighting ($\alpha \in [0.4, 0.6]$);
- **nDCG gap 5-10%:** Favor higher-accuracy path ($\alpha \in [0.7, 0.8]$);
- **nDCG gap > 10%:** Strongly favor higher-accuracy path ($\alpha \in [0.9, 1]$).

5.8 Evaluations of Data Insertion

As established in Section 4.1, Allan-Poe supports efficient data insertion and mark-and-delete operations to accommodate data updates. We evaluate both insertion efficiency and its impact on retrieval quality by inserting varying data volumes into pre-built hybrid indexes and measuring subsequent search performance. As shown in Figure 13, the performance decrease of the updated index is marginal compared to the rebuilt index, which is up to 1% of nDCG@10 with the same QPS. Figure 13 shows that the updated index experiences only marginal performance degradation compared to a complete rebuild, with at most a 1% reduction in nDCG@10 while maintaining equivalent QPS. Furthermore, as shown in Table 2, our insertion strategy incorporates 20% new data with only 14.5% of the computational overhead required for a full index rebuild. These results demonstrate Allan-Poe’s capability to efficiently handle data updates in dynamic environments.

6 Conclusion

This paper presents Allan-Poe, a unified, GPU-accelerated hybrid index that integrates dense vector, sparse vector, full-text, and knowledge graph retrieval. We first analyze the limitations of existing retrieval paradigms and derive design principles for effective hybrid indexing. Guided by these principles, we design an all-in-one graph-based index featuring an isolated heterogeneous edge storage that integrates multiple retrieval paths within a unified structure while minimizing maintenance overhead. Furthermore, we optimize index construction through hybrid distance acceleration, RNG-IP joint pruning, and keyword-aware neighbor recycling, leveraging massive GPU parallelism to accelerate the entire construction pipeline. Finally, we introduce a unified query processing strategy that dynamically fuses information from all retrieval paths to achieve high-accuracy results. Our approach also innovatively augments document-level vector search with knowledge graph reasoning to

handle complex queries. Comprehensive experiments on real-world datasets demonstrate that Allan-Poe consistently outperforms state-of-the-art methods in both efficiency and accuracy.

References

- [1] 2025. Elastic Hybrid Search. <https://www.elastic.co/what-is/hybrid-search>.
- [2] 2025. Infinity Hybrid Search. <https://infiniflow.org/blog/best-hybrid-search-solution>.
- [3] 2025. Milvus Hybrid Search. https://milvus.io/docs/hybrid_search_with_milvus.md.
- [4] 2025. Milvus Rerank Methods. <https://milvus.io/docs/reranking.md>.
- [5] 2025. Weaviate Hybrid Search. <https://docs.weaviate.io/weaviate/search/hybrid>.
- [6] Anas Ait Aomar, Karima Echihabi, Marco Arnaboldi, Ioannis Alagiannis, Damien Hilloulin, and Manal Cherkaoui. 2025. RWalks: Random Walks as Attribute Diffusers for Filtered Vector Search. *SIGMOD* 3, 3 (2025), 1–26.
- [7] Yihao Ang, Yifan Bao, Qiang Huang, Anthony KH Tung, and Zhiyong Huang. 2024. Tsgassistant: An interactive assistant harnessing llms and rag for time series generation recommendations and benchmarking. *PVLDB* 17, 12 (2024), 4309–4312.
- [8] Roi Blanco and Paolo Boldi. 2012. Extending BM25 with multiple query operators. In *SIGIR*. 921–930.
- [9] Sebastian Bruch, Siyu Gai, and Amir Ingber. 2023. An analysis of fusion functions for hybrid retrieval. *TOIS* 42, 1 (2023), 1–35.
- [10] Sebastian Bruch, Franco Maria Nardini, Amir Ingber, and Edo Liberty. 2024. Bridging dense and sparse maximum inner product search. *TOIS* 42, 6 (2024), 1–38.
- [11] Sebastian Bruch, Franco Maria Nardini, Cosimo Rulli, and Rossano Venturini. 2024. Efficient inverted indexes for approximate retrieval over learned sparse representations. In *SIGIR*. 152–162.
- [12] Yuzheng Cai, Jiayang Shi, Yizhuo Chen, and Weiguo Zheng. 2024. Navigating labels and vectors: A unified approach to filtered approximate nearest neighbor search. *SIGMOD* 2, 6 (2024), 1–27.
- [13] Yukun Cao, Zengyi Gao, Zhiyang Li, Xike Xie, S Kevin Zhou, and Jianliang Xu. 2025. Lego-graphrag: Modularizing graph-based retrieval-augmented generation for design space exploration. *PVLDB* 18, 10 (2025), 3269–3283.
- [14] Jianlv Chen, Shitao Xiao, Peitian Zhang, Kun Luo, Defu Lian, and Zheng Liu. 2024. Bge m3-embedding: Multi-lingual, multi-functionality, multi-granularity text embeddings through self-knowledge distillation. *arXiv* (2024).
- [15] Sibe Chen, Ju Fan, Bin Wu, Nan Tang, Chao Deng, Pengyi Wang, Ye Li, Jian Tan, Feifei Li, Jingren Zhou, et al. 2025. Automatic database configuration debugging using retrieval-augmented language models. *SIGMOD* 3, 1 (2025), 1–27.
- [16] Tingyang Chen, Cong Fu, Xiangyu Ke, Yunjun Gao, Yabo Ni, and Anxiang Zeng. 2025. Stitching Inner Product and Euclidean Metrics for Topology-aware Maximum Inner Product Search. In *SIGIR*. 2341–2350.
- [17] Tingyang Chen, Cong Fu, Kun Wang, Xiangyu Ke, Yunjun Gao, Wenchao Zhou, Yabo Ni, and Anxiang Zeng. 2025. Maximum Inner Product is Query-Scaled Nearest Neighbor. *PVLDB* 18, 6 (2025), 1770–1783.
- [18] Gordon V Cormack, Charles LA Clarke, and Stefan Buettcher. 2009. Reciprocal rank fusion outperforms condorcet and individual rank learning methods. In *SIGIR*. 758–759.
- [19] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*. 4171–4186.
- [20] Wei Dong, Charikar Moses, and Kai Li. 2011. Efficient k-nearest neighbor graph construction for generic similarity measures. In *WWW*. 577–586.
- [21] Karima Echihabi, Panagioti Fatourou, Kostas Zoumpatianos, Themis Palpanas, and Houda Benbrahim. 2022. Hercules against data series similarity search. *PVLDB* 15, 10 (2022), 2005–2018.
- [22] Darren Edge, Ha Trinh, Newman Cheng, Joshua Bradley, Alex Chao, Apurva Mody, Steven Truitt, Dasha Metropolitansky, Robert Osazuwa Ness, and Jonathan Larson. 2024. From local to global: A graph rag approach to query-focused summarization. *arXiv* (2024).
- [23] Thibault Formal, Stéphane Clinchant, Hervé Déjean, and Carlos Lassance. 2024. Splate: Sparse late interaction retrieval. In *SIGIR*. 2635–2640.
- [24] Thibault Formal, Carlos Lassance, Benjamin Piwowarski, and Stéphane Clinchant. 2024. Towards effective and efficient sparse neural information retrieval. *TOIS* 42, 5 (2024), 1–46.
- [25] Thibault Formal, Benjamin Piwowarski, and Stéphane Clinchant. 2021. SPLADE: Sparse lexical and expansion model for first stage ranking. In *SIGIR*. 2288–2292.
- [26] Cong Fu, Chao Xiang, Changxu Wang, and Deng Cai. 2019. Fast approximate nearest neighbor search with the navigating spreading-out graph. *PVLDB* 12, 5 (2019), 461–474.
- [27] Jiajie Fu, Haitong Tang, Arijit Khan, Sharad Mehrotra, Xiangyu Ke, and Yunjun Gao. 2025. In-context Clustering-based Entity Resolution with Large Language Models: A Design Space Exploration. *SIGMOD* 3, 4 (2025).
- [28] Trevor Gale, Matei Zaharia, Cliff Young, and Erich Elsen. 2020. Sparse gpu kernels for deep learning. In *SC*. 1–14.
- [29] Jianyang Gao, Yutong Gou, Yuexuan Xu, Yongyi Yang, Cheng Long, and Raymond Chi-Wing Wong. 2025. Practical and asymptotically optimal quantization of high-dimensional vectors in euclidean space for approximate nearest neighbor search. *SIGMOD* 3, 3 (2025), 1–26.
- [30] Jianyang Gao and Cheng Long. 2024. Rabbitq: Quantizing high-dimensional vectors with a theoretical error bound for approximate nearest neighbor search. *SIGMOD* 2, 3 (2024), 1–27.
- [31] Siddharth Gollapudi, Neel Karia, Varun Sivashankar, Ravishankar Krishnaswamy, Nikit Begwani, Swapnil Raz, Yiyong Lin, Yin Zhang, Neelam Mahapatro, Premkumar Srinivasan, et al. 2023. Filtered-diskann: Graph algorithms for approximate nearest neighbor search with filters. In *WWW*. 3406–3416.
- [32] Qianwen Gou, Yunwei Dong, Yujiao Wu, and Qiao Ke. 2024. Semantic similarity-based program retrieval: a multi-relational graph perspective. *FCS* 18, 3 (2024), 183209.
- [33] Yutong Gou, Jianyang Gao, Yuexuan Xu, and Cheng Long. 2025. SymphonyQG: Towards Symphonious Integration of Quantization and Graph for Approximate Nearest Neighbor Search. *SIGMOD* 3, 1 (2025), 1–26.
- [34] Joseph L Greathouse and Mayank Daga. 2014. Efficient sparse matrix-vector multiplication on GPUs using the CSR storage format. In *SC*. 769–780.
- [35] Haoyu Han, Harry Shomer, Yu Wang, Yongjia Lei, Kai Guo, Zhigang Hua, Bo Long, Hui Liu, and Jiliang Tang. 2025. Rag vs. graphrag: A systematic evaluation and key insights. *arXiv* (2025).
- [36] Xanh Ho, Anh-Khoa Duong Nguyen, Saku Sugawara, and Akiko Aizawa. 2020. Constructing A Multi-hop QA Dataset for Comprehensive Evaluation of Reasoning Steps. In *COLING*, Donia Scott, Nuria Bel, and Chengqing Zong (Eds.). 6609–6625.
- [37] Sen Hu, Lei Zou, Jeffrey Xu Yu, Haixun Wang, and Dongyan Zhao. 2017. Answering natural language questions by subgraph matching over knowledge graphs. *TKDE* 30, 5 (2017), 824–837.
- [38] Masajiro Iwasaki and Daisuke Miyazaki. 2018. Optimization of indexing based on k-nearest neighbor graph for proximity search in high-dimensional data. *arXiv* (2018).
- [39] Mengxu Jiang, Zhi Yang, Fangyuan Zhang, Guanhuo Hou, Jieming Shi, Wenchao Zhou, Feifei Li, and Sibao Wang. 2025. DIGRA: A Dynamic Graph Indexing for Approximate Nearest Neighbor Search with Range Filter. *SIGMOD* 3, 3 (2025), 1–26.
- [40] Wenqi Jiang, Marco Zeller, Roger Waleffe, Torsten Hoefler, and Gustavo Alonso. 2024. Chameleon: a heterogeneous and disaggregated accelerator system for retrieval-augmented language models. *PVLDB* 18, 1 (2024), 42–52.
- [41] Bernal Jimenez Gutierrez, Yiheng Shu, Yu Gu, Michihiro Yasunaga, and Yu Su. 2024. Hipporag: Neurobiologically inspired long-term memory for large language models. *NeurIPS* 37 (2024), 59532–59569.
- [42] William B Johnson, Joram Lindenstrauss, et al. 1984. Extensions of Lipschitz mappings into a Hilbert space. *Contemporary mathematics* 26, 189–206 (1984), 1.
- [43] Omar Khattab, Mohammad Hammoud, and Tamer Elsayed. 2020. Finding the best of both worlds: Faster and more robust top-k document retrieval. In *SIGIR*. 1031–1040.
- [44] Omar Khattab and Matei Zaharia. 2020. Colbert: Efficient and effective passage search via contextualized late interaction over bert. In *SIGIR*. 39–48.
- [45] Weize Kong, Jeffrey M Dudek, Cheng Li, Mingyang Zhang, and Michael Bender-sky. 2023. Sparseembed: Learning sparse lexical representations with contextual embeddings for retrieval. In *SIGIR*. 2399–2403.
- [46] Guoliang Li, Jianhua Feng, Jianyong Wang, and Lizhu Zhou. 2008. An effective and versatile keyword search engine on heterogenous data sources. *PVLDB* 1, 2 (2008), 1452–1455.
- [47] Jie Li, Haifeng Liu, Chuanghua Gui, Jianyu Chen, Zhenyuan Ni, Ning Wang, and Yuan Chen. 2018. The design and implementation of a real time visual search system on JD E-commerce platform. In *Proceedings of the 19th International Middleware Conference Industry*. 9–16.
- [48] Peizheng Li, Chaoyi Chen, Hao Yuan, Zhenbo Fu, Hang Shen, Xinbo Yang, Qiang Wang, Xin Ai, Yanfeng Zhang, Yingyou Wen, et al. 2025. Neutron-RAG: Towards Understanding the Effectiveness of RAG from a Data Retrieval Perspective. In *SIGMOD*. 163–166.
- [49] Yunyao Li, Ziyang Liu, and Huaiyu Zhu. 2014. Enterprise search in the big data era: Recent developments and open challenges. *PVLDB* 7, 13 (2014), 1717–1718.
- [50] Zhonggen Li, Xiangyu Ke, Yifan Zhu, Bocheng Yu, Baihua Zheng, and Yunjun Gao. 2025. Scalable Graph Indexing using GPUs for Approximate Nearest Neighbor Search. *SIGMOD* 3, 6 (2025).
- [51] Anqi Liang, Pengcheng Zhang, Bin Yao, Zhongpu Chen, Yitong Song, and Guangxu Cheng. 2025. UNIFY: Unified Index for Range Filtered Approximate Nearest Neighbors Search. *PVLDB* 18, 4 (2025), 1118–1130.
- [52] Fang Liu, Clement Yu, Weiyei Meng, and Abdur Chowdhury. 2006. Effective keyword search in relational databases. In *SIGMOD*. 563–574.
- [53] Hao Liu, Zhengren Wang, Xi Chen, Zhiyu Li, Feiyu Xiong, Qinhan Yu, and Wentao Zhang. 2025. HopRAG: Multi-Hop Reasoning for Logic-Aware Retrieval-Augmented Generation. In *ACL*.

- [54] Kejing Lu, Hongya Wang, Wei Wang, and Mineichi Kudo. 2020. VHP: approximate nearest neighbor search via virtual hypersphere partitioning. *PVLDB* 13, 9 (2020), 1443–1455.
- [55] Yuanhua Lv and ChengXiang Zhai. 2011. When documents are very long, bm25 fails!. In *SIGIR*. 1103–1104.
- [56] Chandana Sree Mala, Gizem Gezici, and Fosca Giannotti. 2025. Hybrid Retrieval for Hallucination Mitigation in Large Language Models: A Comparative Analysis. *arXiv* (2025).
- [57] Yu A Malkov and Dmitry A Yashunin. 2018. Efficient and robust approximate nearest neighbor search using hierarchical navigable small world graphs. *TPAMI* 42, 4 (2018), 824–836.
- [58] Antonio Mallia, Giuseppe Ottaviano, Elia Porciani, Nicola Tonellotto, and Rossano Venturini. 2017. Faster BlockMax WAND with variable-sized blocks. In *SIGIR*. 625–634.
- [59] Antonio Mallia, Torsten Suel, and Nicola Tonellotto. 2024. Faster learned sparse retrieval with block-max pruning. In *SIGIR*. 2411–2415.
- [60] Xupeng Miao, Yining Shi, Hailin Zhang, Xin Zhang, Xiaonan Nie, Zhi Yang, and Bin Cui. 2022. HET-GMP: A graph-based system approach to scaling large embedding model training. In *SIGMOD*. 470–480.
- [61] Tri Nguyen, Mir Rosenberg, Xia Song, Jianfeng Gao, Saurabh Tiwary, Rangan Majumder, and Li Deng. 2016. Ms marco: A human-generated machine reading comprehension dataset. (2016).
- [62] Hiroyuki Ootomo, Akira Naruse, Corey Nolet, Ray Wang, Tamas Feher, and Yong Wang. 2024. Cagra: Highly parallel graph construction and approximate nearest neighbor search for gpus. In *ICDE*. 4236–4247.
- [63] Tao Ouyang, Guihang Hong, Kongyange Zhao, Zhi Zhou, Weigang Wu, Zhao-biao Lv, and Xu Chen. 2025. AdaRAG: Adaptive Optimization for Retrieval Augmented Generation with Multilevel Retrievers at the Edge. In *INFOCOM*. 1–10.
- [64] Jiaul H Paik. 2013. A novel TF-IDF weighting scheme for effective ranking. In *SIGIR*. 343–352.
- [65] James Jie Pan, Jianguo Wang, and Guoliang Li. 2024. Survey of vector database management systems. *VLDBJ* 33, 5 (2024), 1591–1615.
- [66] Panos Parchas, Yonatan Naamad, Peter Van Bouwel, Christos Faloutsos, and Michalis Petropoulos. 2020. Fast and effective distribution-key recommendation for amazon redshift. *PVLDB* 13, 12 (2020), 2411–2423.
- [67] Liana Patel, Peter Kraft, Carlos Guestrin, and Matei Zaharia. 2024. Acorn: Performant and predicate-agnostic search over vector embeddings and structured data. *SIGMOD* 2, 3 (2024), 1–27.
- [68] Yun Peng, Byron Choi, Tsz Nam Chan, Jianye Yang, and Jianliang Xu. 2023. Efficient approximate nearest neighbor search in multi-dimensional databases. *SIGMOD* 1, 1 (2023), 1–27.
- [69] Stephen Robertson. 2025. BM25 and all that—a look back. In *SIGIR*. 5–8.
- [70] Kunal Sawarkar, Abhilasha Mangal, and Shivam Raj Solanki. 2024. Blended rag: Improving rag (retriever-augmented generation) accuracy with semantic search and hybrid query-based retrievers. In *MIPR*. 155–161.
- [71] Xibo Sun, Shixuan Sun, Qiong Luo, and Bingsheng He. 2022. An in-depth study of continuous subgraph matching. *PVLDB* 15, 7 (2022), 1403–1416.
- [72] Shulong Tan, Zhixin Zhou, Zhaozhao Xu, and Ping Li. 2019. On efficient retrieval of top similarity vectors. In *EMNLP*. 5236–5246.
- [73] Godfried T Toussaint. 1980. The relative neighbourhood graph of a finite planar set. *Pattern recognition* 12, 4 (1980), 261–268.
- [74] Sairaj Voruganti and M Tamer Özsu. 2025. MIRAGE-ANNS: Mixed Approach Graph-based Indexing for Approximate Nearest Neighbor Search. *SIGMOD* 3, 3 (2025), 1–27.
- [75] Hui Wang, Wan-Lei Zhao, Xiangxiang Zeng, and Jianye Yang. 2021. Fast k-nn graph construction by gpu based nn-descent. In *CIKM*. 1929–1938.
- [76] Jianguo Wang, Chunbin Lin, Yannis Papakonstantinou, and Steven Swanson. 2017. An experimental study of bitmap compression vs. inverted list compression. In *SIGMOD*. 993–1008.
- [77] Jianguo Wang, Xiaomeng Yi, Rentong Guo, Hai Jin, Peng Xu, Shengjun Li, Xiangyu Wang, Xiangzhou Guo, Chengming Li, Xiaohai Xu, et al. 2021. Milvus: A purpose-built vector data management system. In *SIGMOD*. 2614–2627.
- [78] Mengzhao Wang, Xiangyu Ke, Xiaoliang Xu, Lu Chen, Yunjun Gao, Pinpin Huang, and Runkai Zhu. 2024. Must: An effective and scalable framework for multimodal search of target modality. In *ICDE*. 4747–4759.
- [79] Mengzhao Wang, Boyu Tan, Yunjun Gao, Hai Jin, Yingfeng Zhang, Xiangyu Ke, Xiaoliang Xu, and Yifan Zhu. 2025. Balancing the Blend: An Experimental Analysis of Trade-offs in Hybrid Search. *arXiv* (2025).
- [80] Mengzhao Wang, Xiaoliang Xu, Qiang Yue, and Yuxiang Wang. 2021. A comprehensive survey and experimental comparison of graph-based approximate nearest neighbor search. *PVLDB* 14, 11 (2021), 1964–1978.
- [81] Quan Wang, Zhendong Mao, Bin Wang, and Li Guo. 2017. Knowledge graph embedding: A survey of approaches and applications. *TKDE* 29, 12 (2017), 2724–2743.
- [82] Yuhao Wang, Ruiyang Ren, Junyi Li, Xin Zhao, Jing Liu, and Ji-Rong Wen. 2024. REAR: A Relevance-Aware Retrieval-Augmented Framework for Open-Domain Question Answering. In *EMNLP*. 5613–5626.
- [83] Yining Wang, Liwei Wang, Yuanzhi Li, Di He, and Tie-Yan Liu. 2013. A theoretical analysis of NDCG type ranking measures. In *PMLR*. 25–54.
- [84] Orion Weller, Michael Boratko, Iftekhar Naim, and Jinhyuk Lee. 2025. On the Theoretical Limitations of Embedding-Based Retrieval. *arXiv preprint arXiv:2508.21038* (2025).
- [85] Ho Chung Wu, Robert Wing Pong Luk, Kam Fai Wong, and Kui Lam Kwok. 2008. Interpreting TF-IDF term weights as making relevance decisions. *TOIS* 26, 3 (2008), 1–37.
- [86] Yuexuan Xu, Jianyang Gao, Yutong Gou, Cheng Long, and Christian S Jensen. 2024. irangraph: Improvising range-dedicated graphs for range-filtering nearest neighbor search. *SIGMOD* 2, 6 (2024), 1–26.
- [87] Inbal Yahav, Onn Shehory, and David Schwartz. 2018. Comments mining with TF-IDF: the inherent bias and its removal. *TKDE* 31, 3 (2018), 437–450.
- [88] An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. 2025. Qwen3 technical report. *arXiv* (2025).
- [89] Chen Yang, Sunhao Dai, Yupeng Hou, Wayne Xin Zhao, Jun Xu, Yang Song, and Hengshu Zhu. 2024. Revisiting Reciprocal Recommender Systems: Metrics, Formulation, and Method. In *SIGKDD*. 3714–3723.
- [90] Rui Yang, Boming Yang, Xinjie Zhao, Fan Gao, Aosong Feng, Sixun Ouyang, Moritz Blum, Tianwei She, Yuang Jiang, Freddy Lecue, et al. 2025. Graphusion: A RAG Framework for Scientific Knowledge Graph Construction with a Global Perspective. In *WWW*. 2579–2588.
- [91] Shuo Yang, Jiadong Xie, Yingfan Liu, Jeffery Xu Yu, Xiyue Gao, Qianru Wang, Yanguo Peng, and Jiangtao Cui. 2025. Revisiting the index construction of proximity graph-based approximate nearest neighbor search. *PVLDB* 18, 6 (2025), 1825–1838.
- [92] Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William W Cohen, Ruslan Salakhutdinov, and Christopher D Manning. 2018. HotpotQA: A dataset for diverse, explainable multi-hop question answering. In *EMNLP*. 2369–2380.
- [93] Ziqi Yin, Jianyang Gao, Pasquale Balsebre, Gao Cong, and Cheng Long. 2025. DEG: Efficient Hybrid Vector Search Using the Dynamic Edge Navigation Graph. *SIGMOD* 3, 1 (2025), 1–28.
- [94] Runjie Yu, Weizhou Huang, Shuhan Bai, Jian Zhou, and Fei Wu. 2025. AquaPipe: A Quality-Aware Pipeline for Knowledge Retrieval and Large Language Models. *SIGMOD* 3, 1 (2025), 1–26.
- [95] Ye Yuan, Delong Ma, Zhenyu Wen, Zhiwei Zhang, and Guoren Wang. 2021. Subgraph matching over graph federation. *PVLDB* 15, 3 (2021), 437–450.
- [96] Yuxiang Zeng, Yongxin Tong, and Lei Chen. 2023. Litehst: A tree embedding based method for similarity search. *SIGMOD* 1, 1 (2023), 1–26.
- [97] Huayi Zhang, Lei Cao, Yizhou Yan, Samuel Madden, and Elke A Rundensteiner. 2020. Continuously adaptive similarity search. In *SIGMOD*. 2601–2616.
- [98] Haoyu Zhang, Jun Liu, Zhenhua Zhu, Shulin Zeng, Maojia Sheng, Tao Yang, Guohao Dai, and Yu Wang. 2024. Efficient and effective retrieval of dense-sparse hybrid vectors using graph-based approximate nearest neighbor search. *arXiv* (2024).
- [99] Da Zheng, Xiang Song, Chao Ma, Zeyuan Tan, Zihao Ye, Jin Dong, Hao Xiong, Zheng Zhang, and George Karypis. 2020. Dgl-ke: Training knowledge graph embeddings at scale. In *SIGIR*. 739–748.
- [100] Yingli Zhou, Yaodong Su, Youran Sun, Shu Wang, Taotao Wang, Runyuan He, Yongwei Zhang, Sicong Liang, Xilin Liu, Yuchi Ma, et al. 2025. In-depth Analysis of Graph-based RAG in a Unified Framework. *arXiv* (2025).
- [101] Yifan Zhu, Lu Chen, Yunjun Gao, Ruiyao Ma, Baihua Zheng, and Jingwen Zhao. 2024. HJG: An Effective Hierarchical Joint Graph for ANNS in Multi-Metric Spaces. In *ICDE*. 4275–4287.