



Steiner-Hardness: A Query Hardness Measure for Graph-Based ANN Indexes

Zeyu Wang
Fudan University
zeyuwang23@m.fudan.edu.cn

Qitong Wang*
LIPADE, Université Paris Cité
qitong.wang@u-paris.fr

Xiaoxing Cheng
Tongji University
vozeo@tongji.edu.cn

Peng Wang
Fudan University
pengwang5@fudan.edu.cn

Themis Palpanas
LIPADE, Université Paris Cité
themis@mi.parisdescartes.fr

Wei Wang
Fudan University
weiwang1@fudan.edu.cn

ABSTRACT

Graph-based indexes have been widely employed to accelerate approximate similarity search of high-dimensional vectors. However, the performance of graph indexes to answer different queries varies vastly, leading to an unstable quality of service for downstream applications. This necessitates an effective measure to test query hardness on graph indexes. Nonetheless, popular distance-based hardness measures like LID lose their effects due to the ignorance of the graph structure. In this paper, we propose *Steiner-hardness*, a novel connection-based graph-native query hardness measure. Specifically, we first propose a theoretical framework to analyze the minimum query effort on graph indexes and then define *Steiner-hardness* as the minimum effort on a representative graph. Moreover, we prove that our *Steiner-hardness* is highly relevant to the classical Directed *Steiner Tree* (DST) problems. In this case, we design a novel algorithm to reduce our problem to DST problems and then leverage their solvers to help calculate *Steiner-hardness* efficiently. Compared with LID and other similar measures, *Steiner-hardness* shows a significantly better correlation with the actual query effort on various datasets. Additionally, an unbiased evaluation designed based on *Steiner-hardness* reveals new ranking results, indicating a meaningful direction for enhancing the robustness of graph indexes.

PVLDB Reference Format:

Zeyu Wang, Qitong Wang, Xiaoxing Cheng, Peng Wang, Themis Palpanas, and Wei Wang. *Steiner-Hardness: A Query Hardness Measure for Graph-Based ANN Indexes*. PVLDB, 17(13): 4668 - 4682, 2024.
doi:10.14778/3704965.3704974

PVLDB Artifact Availability:

The source code, data, and/or other artifacts have been made available at <https://github.com/DSM-fudan/Steiner-hardness>.


1 INTRODUCTION

Approximate Nearest Neighbor (ANN) search has recently gained high importance. Compared to exact search, ANN can provide high-quality approximate answers at a significantly reduced query time.

*Qitong Wang is the corresponding author.

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.

Proceedings of the VLDB Endowment, Vol. 17, No. 13 ISSN 2150-8097.
doi:10.14778/3704965.3704974



(a) Performance variance on Deep (b) Detriment to a RAG task

Figure 1: Query performance variance on graph indexes. (a) Histograms of NDC to reach 90% recall on Deep [95] dataset. (b) A real example on a RAG task [119], where the low recall of hard queries impairs model accuracy.

To achieve this, ANN indexes are used, which are built before querying. Among the various ANN index families [103], graph indexes¹ like HNSW [76] have become the state-of-the-art for in-memory ng-approximate (i.e., with no guarantees for the accuracy of the results [33]) search [11, 34, 67]. In graph indexes, vectors are represented as vertices in a graph, and connected by edges based on some kind of proximity between vectors. During query answering, the algorithm starts from an entry vertex and travels along the edges to compute the *k*NN answer.

Although graph indexes show superior query performance, it is observed that the performance of different queries varies vastly [10, 12, 65]. That is, graph indexes answer some queries efficiently as expected (i.e., for simple queries), while for other queries, they perform much worse (i.e., for hard queries). This problem widely occurs in various datasets. As shown in Figure 1(a), to reach 90% recall, the Number of Distance Calculations (NDC)² of the queries varies by 3 orders of magnitude for the Deep dataset [95].

Such a large variance in performance will impair the downstream tasks. For example, in recommendation systems, a reduced search accuracy might drift the user away since irrelevant products are recommended to them, if their profiles correspond to hard queries [43, 49, 74]. Another example is shown in Figure 1(b), where graph indexes are used in a Retrieval-Augmented Generation (RAG) framework [119]. Specifically, HNSW index retrieves related items from documents, and these items act as the prompt for language models. In this case, the recall of ANN search influences model accuracy. In a careful parameter setting, the average recall is 0.97; yet, some

¹In this paper, we use *graph indexes* rather than graph-based indexes for simplicity.

²NDC is widely used in the literature [21, 29, 103] as a machine-independent indicator of query time.




Figure 2: Comparison of LID and ME on the same dataset.

hard queries suffer from recall below 0.2, leading to suboptimal answers by the language model. Note that in actual applications the frequency of such hard queries is not known, and a reliable ANN index is required to be able to effectively handle all queries. In this context, an essential problem is required to be addressed: *How can we design a hardness measure for graph indexes to differentiate simple from hard queries?*

The problem of LID. In the past years, Local Intrinsic Dimensionality (LID) [12] has become the most widely used measure to gauge the difficulty of queries and datasets. Formally, given a query q , LID at distance r is defined as $LID_q(r) = \lim_{\epsilon \rightarrow 0^+} \frac{\ln(F((1+\epsilon)r)/F(r))}{\ln(1+\epsilon)}$, where $F: \mathbb{R} \rightarrow [0, 1]$ is the cumulative distribution function of distances to q . When using LID as the hardness, r is set to be the distance between q and the k -th NN. LID describes the increasing rate of mass w.r.t. the radius at some radius r . The more densely data are distributed at k NN neighborhood, the larger LID is which indicates a harder query.

Generally speaking, LID, along with other existing hardness measures [6, 56, 124], describes the hardness in terms of the data distribution in the high-dimensional space. It works well on partition-based indexes [12, 123, 124] since it evaluates how hard it is to distinguish k NN and other points w.r.t. the distance to the query. However, on graph indexes, the effort to answer a query is directly determined by the connections (i.e., edges), rather than the distance. Thus, for the points that are close to the query but not in the k NN answer, if they are not located on the path from the entry point towards the k NN answers, they will never be accessed during querying.

Figure 2 shows an illustrative example of measuring the hardness of a query q with $k=5$, where n_1 to n_5 are the five NNs. We show the principle of LID in Figure 2(a), where the number of points with slightly larger distance than n_5 to q (i.e., orange points) are also considered for hardness estimation. However, Figure 2(b) shows the shortest path from an entry point to all the NNs. Using the given graph connections, the shortest path will not access the orange points of Figure 2(a). In other words, the point density in the k NN neighborhood does not necessarily influence the hardness of a query in a graph index. This raises the need for a graph-native connection-based hardness measure. Such a measure should be designed based on the query effort analysis on the graph (rather than on the data distribution, like LID).

To achieve this goal, there are three major research challenges that need to be tackled in order:




Figure 3: The correlation between Steiner-hardness (b) and NDC to reach 90% recall is much stronger than LID (a).

(1) (*Query-Algorithm-Independent Effort Analysis*) Given graph index G and query q , how can we estimate the (minimum) effort to answer q in G ?

(2) (*Index-Independent Query Hardness Measure*) How can we judge how hard it is to answer a query q with some graph index?

(3) (*Efficiency*) How can we calculate the hardness efficiently?

Query Effort Analysis. To analyze the query effort on graph indexes, many theoretical studies have been proposed recently [37, 38, 63, 80, 82, 120]. However, there are still four major limitations that render them impractical. First, most studies [37, 38, 63, 80, 82] assume a uniform data distribution, which is very far from the distributions observed in real datasets. Second, the graph structures considered in these studies, such as MSNET [37, 38, 80], ρ -Graph [63, 82], and exact Vamana [54], differ from the graph indexes used in practice. Third, the expected- or worst-case results are insufficient to explain the behaviors observed in practice [97]. Last but not least, all these studies focus on an over-simplified single-direction (i.e., $ef=1$) greedy search algorithm for the 1NN problem.

To overcome these limitations, we propose a theoretical framework to analyze the *Minimum Effort* (ME) to answer queries using a graph index. Our framework introduces a lower bound of the real effort to answer a query using a given graph, based only on the connections of that graph. That is, ME represents the effort of an optimal query algorithm on the given graph. Our framework incorporates three major improvements. First, our analysis can be applied to any graph index and any data distribution. Second, our framework is designed to describe the query effort under different *recall targets*, which is important for the ANN problem. Third, our framework focuses on the analysis of the real bottleneck in the query answering process: as we explain in Section 4, most of the query time is spent on identifying the rest of the k NN answers after having identified the first of the k NN answers. However, this part of the query answering process has not been studied in detail in the past. With these improvements, our theoretical framework can effectively estimate the query effort on a given graph index and explain the performance variance.

Query Hardness Measure. To further define an index independent query hardness measure based on the ME, it is important to extract the common feature of the current indexes. We observe that current graph indexes, even though are of various types, have two common structural features. First, most of the edges are short-range connections. They can all be roughly viewed as sub-graphs of a large K Graph. Second, they use edge-pruning rules to

control the out-degree and sparsify the graph, in expectation of a better navigability-sparsity tradeoff [99]. To this effect, we select MRNG (Monotonic Relative Neighborhood Graph) [38] as the representative graph, which prunes edges from short links with direction-based pruning rules. Then we define our hardness measure, *Steiner-hardness*, as the ME on this MRNG. Figure 3 displays the correlation between LID, *Steiner-hardness* and the NDC to reach 90% recall for the queries on Deep dataset with HNSW. We observe that the correlation is significantly improved with our *Steiner-hardness*, which demonstrates the effectiveness of the proposed measure. **Efficient algorithms to calculate *Steiner-hardness*.** Although *Steiner-hardness* is effective for measuring the query difficulty, it can be computationally intensive when implemented in a naive manner. To overcome this problem, we first draw the connections between ME and the classical Directed *Steiner* Tree (DST) problem [122] and its variants [47], and then leverage the DST solvers to solve ME efficiently. Moreover, we design a novel union-find set based algorithm to reduce our ME to the standard DST problems, which provides more than 1000x speedup when compared to the naive method.

Unbiased Workloads Generation. Finally, we observe that the hardness distribution of current query workloads is biased. As shown in Figure 1, most of the queries are simple, and only a small minority are hard. When using the average performance to show the ability of an index, the result will be dominated by the simple queries, and thus, become over-optimistic. Unfortunately, current benchmarks do not pay particular attention to the distribution of the query workloads [11, 29, 67, 99]. To tackle this problem, we propose a method to build an unbiased workload where queries fall within the same distribution as the dataset, and the hardness of queries follows a uniform distribution. Using this workload, we can stress-test the graph indexes, and evaluate their performance when faced with queries of various hardness.

Our contributions can be summarized as follows.

(1) We propose a theoretical framework to analyze the practical effort to answer approximate k NN queries on graph indexes, which effectively estimates the actual query effort (Section 4).

(2) We develop a novel *Steiner-hardness* measure for queries on graph indexes based on our theoretical framework, along with an efficient algorithm to calculate it with the help of DST solvers. *Steiner-hardness* effectively differentiates the difficulty of different queries for graph indexes (Sections 5 and 6).

(3) We propose a method for building unbiased workloads with variable hardness. This method can then be used to evaluate the comprehensive performance of graph indexes (Section 6.1).

(4) We conduct extensive experiments on public datasets to verify the effectiveness of the proposed analytical framework, *Steiner-hardness* and workload generation methods. Moreover, we evaluate current graph indexes with the new unbiased query workloads. The result provides new insights for index selection (Section 7).

2 PRELIMINARIES

Problem Setting In this paper, we assume a high-d vector database db containing N vectors of length d . The exact k NN query can be defined as follows.

Algorithm 1 Greedy search (graph G , query q , entry point ep , ef)

```

1:  $pq$  = a priority queue with unlimited capacity, initialized with  $ep$ 
2:  $H$  = a max-heap with capacity  $ef$ 
3: while  $pq$  is not empty do
4:    $d_{v_c}, v_c$  = pop an element from  $pq$ 
5:    $d_{v_{top}}, v_{top}$  = the heap top of  $H$ 
6:   if  $d_{v_c} > d_{v_{top}}$  then
7:     break
8:   for each neighbor  $v$  of  $v_c$  which has not been accessed do
9:     if  $D(v, q) < d_{v_{top}}$  then
10:      Insert  $(D(v, q), v)$  into  $pq$  and  $H$ 
11:      mark  $v$  as accessed
12:   resize  $H$  to be  $ef$ 
13: return  $k$  smallest elements in  $H$ 

```

DEFINITION 1 (k NN QUERY). Given an integer k , a query vector q of dimensionality d , and a distance measure D , a k NN query retrieves from db the set of vectors $N_k = \{n_1, n_2, \dots, n_k\}$ such that for any other vector v in db and any $n_i \in N_k$, $D(n_i, q) \leq D(v, q)$.

For simplicity, we denote the distance between query and the i -th nearest neighbor by $d_i = D(q, n_i)$. Approximate k NN query returns approximate answers, $AN_k = \{an_1, an_2, \dots, an_k\}$, instead of the exact k NN results. Recall is often used to indicate the quality of the approximate answers. Formally, $Recall@k = \frac{|AN_k \cap N_k|}{k}$.

Query on Graph Indexes Graph indexes use a directed graph $G(V, E)$ to index vectors, where each vector is represented as a vertex in V , and the edges in E connect vectors based on some kind of proximity. Graph indexes commonly use greedy search to retrieve k NN. As shown in Algorithm 1 and Figure 2(c), the search starts from an entry point ep , which is often selected randomly, and then computes the distance between the neighbors of ep to the query q . The accessed points are stored in a priority queue pq . In the next step, the algorithm selects the closest point to q from pq as the next stop to visit and repeats the above process. Note that not all accessed points can enter pq : the algorithm maintains a size-bounded heap H and only the points that are closer than some point in H are qualified to enter pq . Finally, the algorithm terminates when all the points in pq are farther than the points in H to q . The algorithm is greedy because only relatively close points to the query can be accessed. For example, a point of N_k who is the neighbor of a distant point to q will not be accessed. A direct way to escape such “local optimum” is to increase the capacity of H , i.e., ef , which is also the knob to tune the efficiency-accuracy trade-off in graph indexes.

3 RELATED WORK

3.1 ANN indexes

Partition-based indexes. Partition-based indexes usually first reduce the dimensions of vectors and then build indexes to partition the lower-dimensional space. When querying, the index only accesses data in the partitions that are near to the query. These indexes often provide a quality guarantee for approximate queries [34]. Several index families are designed based on this rationale. The Locality Sensitive Hashing (LSH) indexes [39, 42, 53, 72, 118] use hash functions to project high-d vectors and build spatial indexes. Some tree indexes [17, 32, 69, 70, 78, 79, 100–102, 104, 105, 110, 112, 113] use summarization techniques such as iSAX [87] to reduce dimensionality and build ad-hoc tree indexes to partition the space. DET-LSH is an LSH solution with quality guarantees that employs iSAX,

leading to improved performance [109]. Others directly partition the space hierarchically with hyper-planes [15, 25, 31]. Pivot-based indexes [9, 18, 19, 21, 41, 59, 77] cluster data according to the distance to a given group of pivots [121].

Graph indexes. In the past decade, the superior query performance of graph indexes has attracted great interest from the research and industrial [45, 98, 108, 111, 116] communities. The simplest graph index is KGraph [30], where each vector connects to its K nearest neighbors. Inspired by the small-world phenomenon [62], NSW [75] adds more long-range links to achieve better navigability, and HNSW [76] uses the RNG rule to sparsify the graph and bounds the out-degree of each point. In this way, HNSW achieves a balance between sparsity and navigability leading to a significant improvement. DPG [67], NSG [38] and NSSG [37] further study the influence of the angle between edges on search performance when sparsifying the graph. τ -MNG [80], DEG [48] and [46, 92, 114] also study the navigability-sparsity tradeoff in theory and practice. Some works focus on the selection of entry points [73, 117] while others optimize the search process [20, 40, 65, 71, 106, 115]. The evolution of the graph index design is discussed and verified in [103]. ELPIS improves on the HNSW performance by combining it with an iSAX-based tree [13]. Wang et al. [99] conducts a comprehensive survey on graph indexes.

3.2 Hardness Measures

A few hardness measures are leveraged to measure the query hardness. The most popular measure is LID [12, 50, 51]. Besides LID, there are several other local intrinsic dimension models [52, 61, 86] following a similar rationale. In this paper, we select the most popular LID as in [12]. Due to the computational difficulty, LID is usually estimated by the following equation, which is a maximum-likelihood estimation [7, 64]. Formally, $LID_q = -(\frac{1}{k} \sum_{i=1}^k \ln \frac{d_i}{d_k})^{-1}$. Relative Contrast (RC) [56] and Query Expansion (QE) [6] are also studied in [12] as hardness measures based on data distribution. RC presents a global view for other dataset points by using d_{mean} as an indicator while QE considers the density in the local area. Formally, $RC_q = \frac{d_{mean}}{d_k}$ where d_{mean} is the mean distance between q and all the vectors in db , and $QE_q = \frac{d_{2k}}{d_k}$. Larger RC and QE mean simpler queries. ϵ -hardness is also proposed for pruning-based tree indexes [123, 124]. Given the pruning ability of an index on a specific query, there are some close points to the query that cannot be pruned. These points are viewed as the ME to answer this query. Formally, ϵ -hardness of a query q is defined as $|\{v|v \in db \wedge D(q, v) \leq (1 + \epsilon)d_k\}|$, where ϵ is a user-determined parameter. A recent study predicts the hardness on-the-fly, while the query is executed halfway, using learned models based on the found best-so-far answers [65]. However, it cannot reveal the intrinsic hardness.

3.3 Theoretical Study of Query Complexity on Graph Indexes

Although graph indexes are proposed based on heuristics [62], there are also some studies trying to analyze the complexity of query

answering on graph indexes. The concept of MSNET [26] is introduced in [38] and the expected length of the search path between two points in the dataset in MSNET is $O(n^{\frac{2}{d}} \log n)$ when using a simplified greedy search algorithm under the uniform distribution. Moreover, it proposes a prototype MRNG which is an instance of MSNET. MRNG modifies the RNG [90] to be a directed graph. On MRNG, a point will link to its nearest neighbors unless the neighbors are located in a very close direction. (More details provided in Section 6.) The case when the query is not in the dataset is discussed in [37]. τ -MG [80] is proposed to optimize MRNG when $d_1 < \tau$ while in [120] MRNG is generalized by bounding the out-degree and the considered candidates in the building stage. [63, 82] study this problem on ρ -Graph of different ranges of dimensionality. The worst-case analyses are first studies in [54] under datasets with bounded doubling dimension on exact Vamana graph [58]. However, this study is still limited in $k=1$ and the simplified greedy search algorithm and the handcrafted hard queries somewhat deviate from the original data distribution. Delaunay graph [27] is sometimes utilized to explain this [38, 80, 99], which however, quickly becomes a complete graph as d grows. Recently, [97] tries to build connections between the size of strongly connected components of k NN-induced subgraph and the recall of a query, which inspires our work.

4 CHARACTERIZE THE EFFORT OF ANSWERING A QUERY

We now analyze the effort to answer approximate k NN queries on any given graph index, which lays the foundation of our hardness measure. For graph indexes, the querying time cost is dominated by the time spent on distance calculations between the database points and the query point. The number of Distance Calculations (NDC) is also the number of the accessed points in the graph when querying. Thus, to analyze the query effort, we focus on calculating the minimum number of accessed points on a given graph (Section 4.1). This minimum effort (i.e., ME) provides a strict lower bound for the cost of any query answering algorithm. However, since the common query algorithm is greedy, with a far higher cost than the theoretical bound, we pose three constraints when analyzing the ME to make it more practical (Sections 4.2 and 4.3). Finally, to efficiently calculate the proposed ME (as well as our hardness measure), we map our definitions to the classical DST problems, and adapt DST solvers to our problem (Sections 4.4 and 4.5).

4.1 Basic ME Definition

We first present a basic definition for ME. Intuitively, it is defined as the minimum number of accessed points required to reach a given recall Acc on a graph G .

DEFINITION 2 (ME@Acc). Consider a directed graph index $G(V, E)$, where V is formed by the vectors in db , a query q and its k NN set N_k . Assume in the query process, one can only travel the graph along the edges E . Then given a least recall requirement $Acc \in [0, 1]$, $ME@Acc$ is defined as the minimum number of accessed vertices to obtain at least $Acc * k$ vertices in N_k , from any possible entry point $ep \in V$.

Figure 4 (a) and (b) show examples for $ME@Acc$ when $Acc=1$ and 0.8 respectively. When $Acc=1$, the optimal entry point is v_1




Figure 4: An illustrative example of our ME definitions with $k=5$. The orange points and edges form Y . (a) $Acc=100\%$, (b) $Acc=80\%$, (c) entry point is limited to be in N_k , (d) $p=0.4$, (e) Limited range of candidates, (f) $ME - exhaustive$.

and it needs to access at least 12 points in G . When $Acc=0.8$, the optimal entry point is changed to be v_2 and it only needs to obtain $0.8 * 5 = 4$ points in N_k to reach 80% recall@5.

Limitations of the basic ME definition. To achieve $ME@Acc$, three tough conditions should be satisfied:

① *Optimum entry point.* Every point $v \in V$ can be the entry point, and has a corresponding shortest routing path p_v^3 . ME is computed by starting from the entry point v^* with the smallest p_{v^*} among all points $v \in V$. v^* is called the optimal entry point.

② *Infinite range of accessible candidates.* All neighbors of a visited point can be accessed without qualification, as long as they can provide shortcuts to the kNN . Note that the greedy search algorithm requires that only vertices close enough to the query are qualified to be accessed, as described in Section 2.

③ *Optimum search path and terminals.* Given the optimal entry point and the unlimited candidates that can be accessed, ME requires finding the optimum search path. Moreover, since we can select any $Acc * k$ points in N_k as terminals, we should meanwhile select the optimum group of terminals that can be accessed with the shortest routing path.

Although these conditions are very strict, $ME@Acc$ is the tightest lower bound of the effort to answer a query on the given graph, for any possible query algorithms, present and future. Nevertheless, for common greedy search algorithms (Algorithm 1), this bound is inaccurate to represent the practical effort since these conditions are not satisfied. Therefore, in the following section, we adjust ME to the greedy search algorithm by adding three constraints.

4.2 Adapt ME for Greedy Search

To render the estimated ME closer to the actual effort of the current greedy search algorithm, we constrain Definition 2 in three ways: (1) the entry point should be selected from a limited range, (2) the range of accessible points in the search path should be limited, and (3) the optimal search path should satisfy the above two constraints. More formally:

³Actually, the accessed vertices and edges form a tree instead of a simple path as we allow backtracking. The shortest routing path means a minimum number of nodes in the tree. For simplicity, we still use the term "path" in the following.




Figure 5: Query time breakdown on 1,000 queries ($k=50$).

DEFINITION 3 ($ME_\delta^p @ Acc$). Consider a directed graph index $G(V, E)$, a query q and its kNN set N_k . Given a least recall requirement $Acc \in [0, 1]$, a probabilistic lower bound $p \in (0, 1]$, and a distance threshold δ , the constrained ME , denoted by $ME_\delta^p @ Acc$, is defined as the number of points in a subgraph $Y(V_Y, E_Y)$, where Y is induced from G by V_Y , satisfying:

① (*Entry point*) There exists a set $N'_k \subset N_k$ with $|N'_k| \geq p * k$. For $\forall n \in N'_k$, there exists a path on subgraph Y from n to at least $Acc * k$ points in N_k .

② (*Range of accessible candidates*) $\forall v \in V_Y, D(v, q) \leq (1 + \delta) * d_k$.

③ (*Optimum*) Among all the qualified subgraph, Y has the minimum number of vertices.

For clarity, such a subgraph Y is also denoted by $Y_p^\delta @ Acc$.

In the following, we write ME for short when it does not cause ambiguity. Note that $Y_p^\delta @ Acc$ does not always exist and in this case we assume $ME_\delta^p @ Acc = \infty$. We explain the rationale of this definition in a progressive way as shown in Figure 4(b) to (e).

① **The entry point is constrained to be one of the kNN answers of the query.** In the basic ME definition, we require a globally optimal entry point, from which we can construct the shortest search path. However, it is very hard for the current algorithm to find such an entry point. Nonetheless, we observe that using the current graph indexes, it is easy to obtain one of the kNN . Specifically, we break the query process into two phases: Phase 1 starts from the entry point and continues until we first access one of the kNN points; Phase 2 starts immediately afterwards, and ends at the termination of the algorithm, as in [48, 97]. As shown in Figure 5,

Phase 1 costs less than 7% of the total query time. This indicates that in practice *the bottleneck of the query process is Phase 2, not Phase 1*. To this effect, we focus on modeling the effort in Phase 2, we omit the effort in Phase 1, and we use one of k NN as the entry point. The example is shown in Figure 4(c), where n_2 is selected as the optimum entry point which can reach at least 4 NNs in Y . This change of the entry point only increases ME by 1.

Since which nearest neighbor can be first accessed in Phase 1 is unknown and somewhat random, we further constrain ME by allowing at least $p * k$ points in N_k (i.e., points in N'_k) are able to reach $Acc * k$ points in N_k . That is, as long as Phase 1 accesses one of the points in N'_k , the query cost in Phase 2 is bounded by $ME = |Y_{p@Acc}^\infty|$, where the symbol ∞ means there is no constraint on the range of accessible candidates of Y . In this case, we view the parameter p as a probabilistic lower bound to reach the target recall. As shown in Figure 4(d) where p is set to be 40%, there are $40\% * 5 = 2$ points who can reach at least 4 NNs. Specifically, starting from n_2 we can reach $\{n_2, n_4, n_3, n_1\}$, and from n_5 we can reach $\{n_5, n_4, n_1, n_3\}$. Here $N'_k = \{n_2, n_5\}$.

② **The accessible points in the search path are constrained to be near to the query.** In our basic ME definition, all the points in the graph are available to be accessed, as long as they can contribute to a shorter search path. However, recall that in the current greedy search algorithm, the points that are closer to the query will be visited first, and the points that are farther than all the points in the ef -sized heap H to the query will be skipped. This greedy procedure limits the range of accessible points in the search path, i.e., no farther than the largest distance between the query and the points in H . Consequently, we add a parameter δ to limit the accessible points when searching within $(1 + \delta) * d_k$. Notice that with a limited distance threshold, some shortcuts that contain distant points will be given up. As shown in Figure 4(e), v_3 is replaced by two closer points while v_4 and v_5 are replaced by four closer points, which increase ME by 3.

Determine δ with the critical point δ_0 . In the greedy search algorithm, the distance threshold is determined on-the-fly by the heap H , which makes it impossible to model the exact case when analyzing the effort. Nevertheless, we observe that there is a special value for parameter δ , which is the minimum value such that there exists one qualified subgraph Y . We name this value the *critical point*, denoted by δ_0 . Formally:

DEFINITION 4 (CRITICAL POINT δ_0). *Given a least recall requirement $Acc \in [0, 1]$, a probabilistic lower bound $p \in (0, 1]$, the critical point $\delta_0^p@Acc$ is defined as the minimum distance threshold under which there exists a subgraph $Y_{p@Acc}^{\delta_0}$ satisfying the three constraints in Definition 3.*

Note that δ_0 describes the minimum requirement of the range of accessible candidates. As δ increases from δ_0 , $ME_{\delta}^p@Acc$ is monotonically non-increasing and gradually converges to $ME_{\infty}^p@Acc$. In other words, δ_0 is the parameter that maximizes $ME_{\delta}^p@Acc$ among all possible δ . In practice, the quality of the points in heap H quickly increases to a stable range in Phase 1 on high-dimensional data [68]. This indicates that δ is rather small in Phase 2, on which our ME

focuses. In this case, a smaller δ can better reflect the actual querying situation, which is verified in the experiments in Section 7.6. We elaborate on how to find δ_0 in Section 5.

4.3 Incorporate Decision Cost into ME

Besides the above constraints, there is still a major difference between $ME_{\delta}^p@Acc$ and actual query cost, i.e., the decision cost of each step when querying. Remind that in Algorithm 1, to decide which point as the next stop, we need to compute the distance between each neighbor of the current point and the query. These neighbors accessed in the greedy search algorithm are not included in $Y_{p@Acc}^{\delta}$ since they do not contribute to the navigability of Y , and thus out of the consideration of $ME_{\delta}^p@Acc$. To this effect, we further modify our definition of ME to reflect the decision cost.

DEFINITION 5 (DECISION COST). *Given a graph index $G(V, E)$, and a subgraph Y of G , the decision set of a point $v \in Y$ is defined as*

$$DS(v) = \begin{cases} \{v\} & \text{if } v \text{ has no out-neighbors in } Y \\ \{x | x = v \vee (v, x) \in E\} & \text{otherwise} \end{cases} \quad (1)$$

The decision cost of v is the cardinality of the decision set, i.e. $Cost(v) = |DS(v)|$. Furthermore, the decision cost of the subgraph Y is defined as $Cost(Y) = |\bigcup_{v \in Y} DS(v)|$.

The decision cost of a point acts as the weight of the point, and then the definition of ME is further modified as the minimum weight of the qualified subgraph.

DEFINITION 6 ($ME_{\delta}^p@Acc$ -exhaustive). *$ME_{\delta}^p@Acc$ -exhaustive is defined by modifying the third constraint of $ME_{\delta}^p@Acc$, i.e.,*

③ *(Optimum with decision cost) Among all the qualified subgraph, Y has the minimum cost $Cost(Y)$.*

We use ME -exhaustive for short when no ambiguity. As shown in Figure 4(f), the neighbors in the original graph G is also counted in ME -exhaustive, e.g., v_3 and v_4 , despite that they do not contribute to the reachability. This poses a punishment on the vertices with a large out-degree which introduce higher decision costs during querying. For example, although v_5 can directly build a bridge from n_5 to n_4 , the large out-degree incurs a high decision cost and thus it is replaced by three other points for the optimum.

By now, we have built the analytical framework for the ME of query answering in graph indexes. All our definitions of ME can be utilized to describe different query algorithms on graph indexes. It's clear that on the same graph and settings, $ME@Acc \leq ME_{\delta}^p@Acc \leq ME_{\delta}^p@Acc$ -exhaustive. In the following, we utilize $ME_{\delta_0}^p@Acc$ -exhaustive to describe the effort of the greedy search algorithm. We also note that our analytical framework is flexible by combining different constraints on the basic format of ME (Definition 2) such that it can be adapted for bounding the effort of more advanced search algorithms in the future.

4.4 Map ME Definitions to DST Problems

The core problem of solving ME is to find a qualified subgraph Y , which is however, an NP-hard problem [122]. To calculate it efficiently, in this subsection, we prove Y can be mapped to DST [122] and its variants. Specifically, we prove the equivalence or relevance

between the DST problems and three definitions of ME, i.e., $ME@Acc$, $ME_{\delta}^p@Acc$ and $ME_{\delta}^p@Acc - exhaustive$ using Theorem 1, 2 and 3, respectively. Based on these analyses, we can calculate our ME with the help of solutions for the classical DST problems.

Map $ME@Acc$ to DST. We start from $ME@Acc$ (Definition 2) where $Acc=100\%$.

DEFINITION 7 (DIRECTED STEINER TREE (DST)). Given a directed graph $G(V, E)$, a root vertex $r \in V$, a group of terminals $T \subset V$, DST is the subgraph of G with minimum edge cost satisfying that for $\forall v \in T$, there is a path from r to v on DST, denoted by DST_T^r and the cost is denoted by $|DST_T^r|$.

Note that the original DST problem is defined on edge-weighted graphs [122], and in this paper, we only study the unweighted version, which is a special case of the original problem.

THEOREM 1. Given a graph $G(V, E)$, a query q with its kNN N_k ,

$$ME@100\% = \min_{r \in V} |DST_{N_k}^r| \quad (2)$$

Map $ME_{\delta}^p@Acc$ to vDSN. The constraint on the limited range of accessible candidates (2) is equivalent to the minimum DST on a subgraph, while the limitation on the entry point in N_k is equivalent to the minimum DST when the root $r \in N_k$. We then use directed Steiner network (DSN) problem [23] to connect ME with accuracy Acc and probabilistic lower bound p constraints.

DEFINITION 8 (VERTEX-FOCUSED DIRECTED STEINER NETWORK (vDSN)). Given a directed graph $G(V, E)$, a collection vertex pairs $P \subset V \times V$, vDSN is the subgraph of G with minimum number of vertices satisfying that for $\forall (v, w) \in P$, there is a path from v to w on vDSN, denoted by $vDSN_P$ and the cost is denoted by $|vDSN_P|$.

THEOREM 2. Given a directed graph $G(V, E)$, a query q with its kNN N_k , a least recall Acc and a probabilistic lower bound p , then

$$ME_{\infty}^p@Acc = \min |vDSN_P| \quad (3)$$

where $P \subset N_k \times N_k$, and the number of starting points in P are at least $p * k$ and for each starting point, the number of different terminals are at least $Acc * k$.

Map $ME_{\delta}^p@Acc - exhaustive$ to node-weighted vDSN. For $ME - exhaustive$, there is no specific Steiner-related problem matching the definition. However, we can upper bound it by the node-weighted vDSN problem, that is, the target is modified to find the minimum sum weights of vertices in the subgraph. For simplicity, we still denote the cost by $|vDSN_P|$.

THEOREM 3. Given a graph $G(V, E)$, a vertex weighting function $W(v) = Cost(v)$ for $v \in V$, we have

$$ME_{\infty}^p@Acc - exhaustive \leq \min |vDSN_P| \quad (4)$$

where P has the same requirements as Theorem 2.

4.5 Calculate ME using Steiner Tree Solvers

Based on these theorems, the three definitions of ME is mapped to the DST, vDSN, and node-weighted vDSN problems.

Finding the exact DST is an NP-hard and APX-hard problem [16, 44], which means that there is even no polynomial-time constant-approximation algorithm. To the best of our knowledge, the exact [85, 91] and approximate [16, 44, 107] algorithms are very time-consuming. Meanwhile, efficient heuristic algorithms can provide

high-precision results in practice [1, 55, 91]. Therefore, we use the shortest path-based algorithm to find the shortest paths from root vertex r to each of the terminals $t \in T$ and unite them to build DST.

The DSN problem is harder than DST problem. None of exact [35] algorithms or approximate [16, 24, 28, 36] algorithms are efficient. To this effect, we design a heuristic algorithm. We first group the set P of vertex pairs by the starting vertex and reduce the DSN problem to multiple DST problems. Then we find DSTs one by one as described above and unite them to form the final result.

Lastly, to efficiently estimate $ME - exhaustive$, we adapt the shortest path-based algorithm with the node-weighted Dijkstra algorithm [2]. The weight of a node is set to the cost defined in Definition 5. We omit the algorithm details due to the lack of space.

5 FIND THE CRITICAL POINT δ_0

Given the solutions for ME, we still need algorithms to calculate the critical point δ_0 , and the optimum set of vertex pairs P . In this subsection, we propose an efficient exact algorithm to achieve these targets at the same time.

The basic structure of the algorithm. To obtain δ_0 , the basic algorithm is as follows. First, we initialize Y as empty and then add vertices to Y in the ascending order of distance to the query one by one. After each insertion, we check whether the current subgraph Y is a qualified $Y_p^{\delta}@Acc$. In this process, the list of the sorted vertices can be supported by many exact kNN indexes in CPU [60, 66, 89] and algorithms in GPU [60, 88, 93, 94], and the major time cost comes from the graph qualification. To check the graph, a naive method is to iterate all the points in N_k and for each point, we use BFS to obtain their reachable points in N_k . Once the number of points that can reach at least $Acc * k$ points in N_k on Y is no less than $p * k$, the current graph is qualified and $\delta_0 = D(n_i, q)$, where n_i is the latest inserted point. Obviously, the repeated reachability checking is redundant and incurs unnecessary costs.

Avoid redundant computation by maintaining reachability information on the fly. we design a novel algorithm to maintain the reachability information of all inserted points on the fly. The basic idea is to group together the points that are reachable to each other in Y as a union-find set, (i.e., the reachable group), and maintain the reachability between different reachable groups through a graph (i.e., union-find set graph (USG)). When we check the reachability of a point, we can directly apply BFS on the small USG instead of the complete subgraph Y to reduce the complexity.

Update USG by loop detection. To minimize the size of USG and thus the complexity of the reachability check, we merge all the reachable groups belonging to a loop in USG as a single one. That is, after updating, USG will be a directed acyclic graph (DAG). See Figure 6 as an example, where v_{10} is newly inserted. We first add v_{10} as an individual reachable group into USG, along with corresponding edges. After that, we detect a new loop in USG that contains three reachable groups (i.e., RG_1 , RG_2 , and RG_3). Then we merge these three reachable groups into a single one in USG (i.e., RG_5), and also unite their corresponding union-find sets. Note that the edges inside the loop will be removed while the edges that connect reachable groups outside the loop will be reserved for the new reachable group (e.g., the edge connecting RG_1). In this way, the subgraph with 16 vertices and 19 edges is reduced to the USG




Figure 6: The updates of USG when adding a new vertex v_{10} into the subgraph. In (b), the vertices in the same circle form a reachable group.

with 6 vertices (i.e., reachable groups) and 4 edges. The following theorems prove that the loop merge can keep the property of a reachable group. The proof is omitted due to lack of space.

THEOREM 4. *Given a subgraph $G_i(N_i, E_i)$, two reachable groups $RG_1, RG_2 \subset N_i$, $RG_1 \cap RG_2 = \emptyset$, then if $\exists v_1 \in RG_1, v_2 \in RG_2$, and v_1 can reach v_2 in G_i , then $\forall v \in RG_1, w \in RG_2, v$ can reach w in G_i .*

COROLLARY 1. *In a USG of a subgraph G_i , if there exists a loop, $RG_i \rightarrow RG_{i+1} \rightarrow \dots \rightarrow RG_j \rightarrow RG_i$, then for $\forall v, w \in RG_t, i \leq t \leq j, v$ can reach w in G_i .*

In the implementation, we use the classical Depth-First-Search (DFS) algorithm to detect whether there is a loop on USG. This loop detection is executed every time a new point is inserted into USG. After detecting the loop, we collect all the incoming and outgoing edges that point to and start from the reachable groups in the loop, respectively. Then we remove all these reachable groups and related edges from USG, and insert a new vertex into USG, along with the collected edges beforehand for this new vertex.

The details of the algorithm are shown in Algorithm 2. After preparing the NN list and initialization (lines 1-3), the main body of the algorithm is a loop that takes one NN to increment the subgraph each time (lines 5-7). The loop can be divided into three stages. In the first stage (lines 8-12), we locate the in- and out-neighbors of the new point and add edges on USG. In the second stage (lines 13-15), we detect and merge the loops in USG. The union-find sets are also updated accordingly. In the third stage (lines 16-23), we check the reachability on USG. Then for each root, we use BFS to find which points in N_k are reachable from the root and count qualified points (lines 19-25). Finally, $D(n_i, q)$ is returned as δ_0 along with the found vertex pairs set P (line 27).

Complexity Analysis. Assuming the average number of iterations is \hat{it} , the complexity of the first stage is $O(\hat{it})$ since \hat{it} will be the expected size of Ins , and also the number of vertices in the complete subgraph. For the second stage, the loop detection costs $O(|USG|)$ where $|USG|$ denotes the sum number of vertices and edges in USG, usually ranging from tens to a few hundred, while loop merge needs only constant time. The third stage needs at most k times of BFS search, so the complexity is $O(k|USG|)$. Overall, the time complexity is $O(\hat{it}(\hat{it} + k|USG|))$. In contrast, the complexity of the naive method is $O(k\hat{M}\hat{it}^2)$, where \hat{M} is the expected average out-degree of the subgraph, usually ranging from tens to hundreds. Given that \hat{it} usually ranges from hundreds to tens of thousands,

Algorithm 2 Find δ_0 and P (graph G , query q , reversed graph $revG$, k , recall target Acc , probabilistic lower bound p)

```

1: Get sufficient number of NN of  $q$ .
2: Initialize an empty union-find set  $UF$ , an empty union-find set graph  $USG$ , and an empty hash set  $Ins$ .
3: Initialize the result vertex pair set  $P$ .
4: while True do
5:    $n_i = \text{pop}$  the first element from NN.
6:   Add  $n_i$  as an individual union-find set to  $UF$ .
7:   Add  $n_i$  as an individual vertex in  $USG$ .
    $\triangleright$  (1) Subgraph construction.
8:    $N_{out} = G[n_i] \cap Ins, N_{in} = revG[n_i] \cap Ins$ .
9:   for  $n_{out}$  in  $N_{out}$  do
10:     Add an directed edge from  $n_i$  to  $UF[n_{out}]$  in  $USG$ , if not exists.
11:   for  $n_{in}$  in  $N_{in}$  do
12:     Add an directed edge from  $UF[n_{in}]$  to  $n_i$  in  $USG$ , if not exists.
    $\triangleright$  (2) Reduce USG to a DAG by loop detection.
13:   while there exists a loop in  $USG$  do
14:     Unite all the RG in this loop as a single RG.
15:     Unite all the union-find sets in this loop in  $UF$ .
    $\triangleright$  (3) Qualification check: the reachability of  $kNN$ .
16:   Initialize  $R_k$  as an empty hash map, where key is the root of a union-find set, and the value is a set containing the elements in this union-find set.
17:   for  $1 \leq j \leq \min(k, i)$  do
18:      $R_k[UF[n_j]].\text{add}(n_j)$ 
19:   for  $(s, Set) \in R_k$  do
20:     Use BFS to find which keys in  $R_k$  are reachable in  $USG$  from  $s$ , and store them into a set  $reach$ .
21:      $reach = reach \cup Set$ 
22:     if  $|reach| \geq Acc * k$  then
23:        $P.\text{add}(s, t)$ , for  $\forall t \in reach$ 
24:   if  $|P.\text{keys}().\text{distinct}()| \geq p * k$  then
25:     break
26:   Add  $n_i$  into  $Ins$  and clear  $P$ .
27: return  $D(n_i, q)$  as  $\delta_0$ , and  $P$ 

```

our algorithm can provide $k * \hat{M}$ times speedup according to this complexity analysis.

6 Steiner-HARDNESS

Now we describe our *index-independent Steiner-hardness* for queries on all current graph indexes, based on our definitions of ME. Specifically, we use $ME_{\delta_0}^p @Acc - \text{exhaustive}$ on a representative graph structure, (approximate) MRNG to define *Steiner-hardness*.

DEFINITION 9 (Steiner-HARDNESS). *Given a dataset db we build approximate MRNG on db with a parameter efC . Then, Steiner-hardness for a query q is defined as*

$$\text{Steiner}_q = ME_{\delta_0}^p @Acc - \text{exhaustive on MRNG} \quad (5)$$

where Acc is the given recall requirement, and p is the probabilistic lower bound, as in Definition 6.

To build MRNG, for any point v , all other points in db are sorted according to the distance to v in ascent. Then we insert edges from v to these points one by one, as long as the angle between the new edge and all existing edges are no smaller than 60° . Since all the current graph indexes only select neighbors from a pool of NNs of limited size efC ⁴, we approximate MRNG by also considering only efC NNs rather than all points. In this case, the approximate MRNG can be viewed as a pruned version of K Graph where $K = efC$. The reasons for selecting approximate MRNG as the representative graph are as follows.

(1) *Approximate MRNG and current graph indexes are all pruned versions of a large K Graph.* As shown in Figure 7, we count the percentage of the edges in graph indexes that occur in K Graph. The results show that over 80% of the edges of current graph indexes

⁴It is denoted by efC in HNSW, C in NSG and τ -MNG, and k_{ext} in DEG.




Figure 7: Overlap ratio of edges in current graph indexes and K Graph with varying K .

are in a K Graph with $K \geq 500$. Therefore, approximate MRNG represents a common structure for advanced graph indexes.

(2) Both approximate MRNG and other graph indexes limit the number of edges. Approximate MRNG selects edges according to the spatial distribution with bounded out-degree. Furthermore, it is proved to own a nice navigability [120]. Thus, approximate MRNG reaches a balance of sparsity-navigability trade-off, which matches the rationale of current graph indexes.

(3) The structure of MRNG only relies on the data distribution in high- d space and is not influenced by randomness or any particular index-building strategy.

6.1 Unbiased Workload Generation

Based on our *Steiner-hardness*, we can build unbiased query workloads for evaluating graph indexes. To comprehensively and fairly evaluate the graph index, the produced queries must follow the following two principles:

- (1) **Validity.** The queries should follow the distribution of the dataset.
- (2) **Uniformness.** The hardness of the queries should follow a uniform distribution across the hardness spectrum.

Validity avoids the queries which are out of the distribution of the dataset. For example, in a dataset where the values in the vectors are between 0 and 1, a query vector with values beyond 100 is meaningless for evaluation. Uniformness guarantees that the workload will not exhibit bias neither for simple, nor for hard queries. Based on these two principles, the core idea of our approach is a generation-and-check method. That is, we first generate sufficient vectors by over-sampling from the dataset⁵ and then check the hardness of these vectors. Finally we preserve the vectors that satisfy the uniformness principle.

Specifically, we use Gaussian Mixture Model (GMM) [3] to learn the data distribution of the dataset. According to our pre-experiments, the distribution of the Mahalanobis distance [5] – which is adopted to test whether two groups of vectors are in the same distribution [57] – on generated data is very similar to the dataset. It verifies the effectiveness of GMM. We use the GMM model to produce sufficient candidate vectors, and compute *Steiner-hardness* of them. Secondly, we build histograms on the new data w.r.t. their *Steiner-hardness*. Specifically, we remove the extreme values, i.e., too simple and hard queries, and then uniformly split the range of hardness by h equi-length segments. Given the cardinality of the

⁵In the query set of the original workload, the hard query is usually insufficient for evaluation. So we need to generate more queries as candidates.

Table 1: Dataset statistics.

	Base size	Query size	Dim.	Source	Type
Deep	1,000,000	10,000	96	[22]	Image
GIST	1,000,000	1,000	960	[8]	Image
Glove	1,183,514	10,000	100	[81]	Text
Rand	1,000,000	10,000	100	$U(-1, 1)$	Syn.
Gauss	1,000,000	10,000	100	$N(0, 1)$	Syn.
Tiny10M	10,000,000	1,000	150	[73]	Image
Deep1B	1,000,000,000	10,000	96	[14]	Image

new workload as Q , we randomly sample $\lceil Q/h \rceil$ queries from each segment to build the workload. In this way, the new workload encompasses the entire spectrum of *Steiner-hardness*.

7 EXPERIMENTS

7.1 Experimental Setup

Experiments were conducted on an machine with two 40-core Intel(R) Xeon(R) Platinum 8383C 2.70GHz CPUs and 1TB DDR4 main memory. All the codes are implemented in C/C++, compiled by g++ 10.1.0 with $-O3$ optimization. We use five public datasets and two synthetic dataset with varying sizes and dimensionalities, as shown in Table 1. The results of Gauss dataset are omitted since they are similar to Rand dataset. Other omitted datasets are shown in our code repository. Since only HNSW and K Graph can be built on Deep1B in our machine due to memory limitations, in this paper, we only use Deep1B to test the scalability of our approaches.

Algorithms and parameters. We select five representative graph indexes for evaluation, including four advanced indexes and the K Graph as a baseline. HNSW [76] and NSG [38] are the state-of-the-arts according to previous benchmarks [96]. We only test the base layer of HNSW for evaluation, since the upper layers of HNSW only provide benefits in very low dimensionality [68, 103]. We select τ -MNG [80] and DEG [48] as the latest indexes. We disable the dynamic edge optimization of DEG to control the construction time as in the original paper. The parameters for building these indexes are tuned for the best query performance. When evaluating our analytical framework, Acc and k are set to 98% and 50, respectively by default. efC is set to 2048. The default dataset and index are Deep and HNSW. We tune p in the range 0.85~1 to achieve the optimal correlation.

7.2 Effectiveness of ME

We first evaluate whether $ME_{\delta_0}^p @Acc - exhaustive$ can describe the actual query effort on the graph indexes. Given a query and a recall target, we use the least NDC to reach the recall as the query effort. As shown in Figure 8, the correlation coefficients of ϵ -effort are on average **0.66**, **0.17**, **0.70**, and **0.74** in GIST, Rand, Glove and Tiny10M datasets for all evaluated indexes⁶, whereas our $ME - exhaustive$ is **0.95**, **0.98**, **0.92** and **0.98** respectively. Clearly, $ME - exhaustive$ can describe the query effort more precisely than ϵ -hardness due to the focus on the graph connections.

⁶ τ -MNG and DEG indexes and Deep dataset are omitted due to lack of space, whose results are similar to other settings.

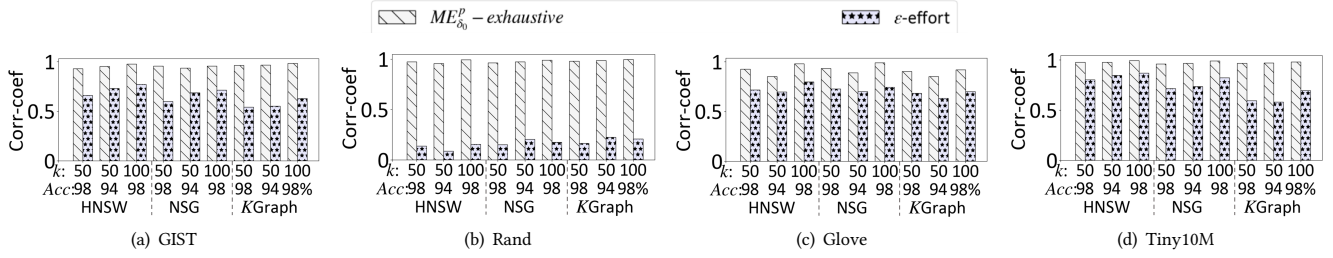


Figure 8: Pearson correlation coefficients between $ME_{\delta_0}^p - exhaustive$ and actual query effort (NDC).

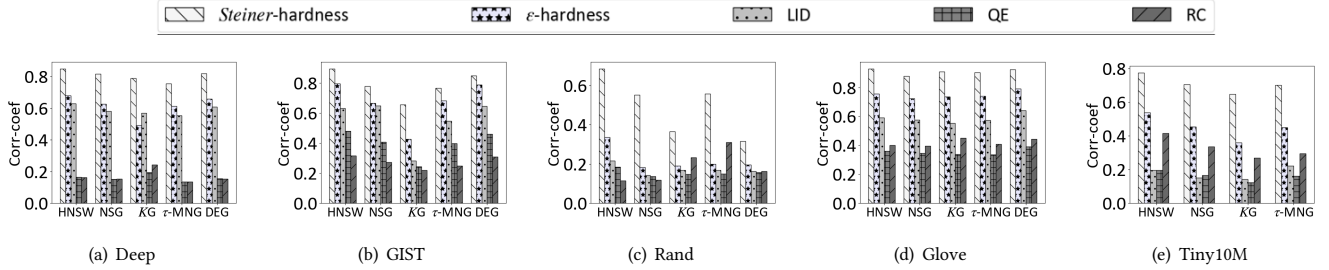


Figure 9: Correlation coefficients between different hardness measures and the average query effort (NDC)

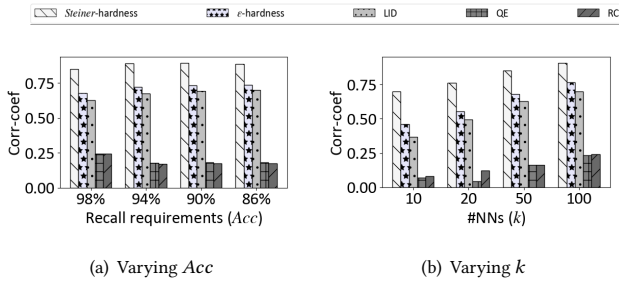


Figure 10: Correlation coefficients of hardness measures under different parameters.

7.3 Effectiveness of *Steiner-hardness*

We then study the effectiveness of our index-independent *Steiner-hardness*. That is, whether a query with larger *Steiner-hardness* requires more effort to be answered for graph indexes. As studied in [71, 103], the insertion order when building graph indexes significantly impacts the effort to answer a given query. Therefore, to mitigate randomness, we shuffle the dataset and build multiple indexes with different insertion orders on the same dataset, and take the average query effort on these indexes as the actual effort.

Figure 9 shows the correlation coefficients between the query effort and different hardness measures on different datasets and indexes⁷. Our *Steiner-hardness* shows the strongest correlation among all the hardness measures, with on average **0.75** correlation coefficient, whereas ϵ -hardness, LID, QE and RC are **0.50**, **0.42**, **0.26**

⁷DEG fails to be built on Tiny10M due to logic errors.

and **0.31**, respectively, far smaller than *Steiner-hardness*. These results demonstrate that *Steiner-hardness* can effectively represent the intrinsic hardness of the query on graph indexes.

It is observed that the correlation on real datasets is stronger than synthetic datasets. It might be because that real datasets are often distributed in a clustered manner [83, 84] and the k NN of the query are often highly interconnected in the graph [103], which indicates that they will be visited together when querying. In this way, the number of close points correlates more to the actual query effort than in the synthetic datasets.

7.4 Comparison of Old and New Workloads

We first check the hardness distribution of the queries in our newly generated workloads which are produced by adopting the approach described in Section 6.1. The new workload contains 1,000 queries, constructed by randomly selecting 50 queries from 20 hardness segments, i.e., $h=20$. The result is shown in Figure 11, where we compare our new unbiased workload against the original workload on three representative datasets⁸. Assume simple queries are the ones with hardness in the 20% lowest range of the whole hardness spectrum. Then the percentage of simple queries on the original workload is **78%**, **80%**, and **82%** on the three datasets respectively, while **20%** on the new unbiased workloads. This verifies our workloads do not show bias for simple or hard queries.

In Figure 12, we compare the query accuracy under old and new workloads using HNSW. To fairly display the difference, we set NDC of each query as a fixed number and calculate the recall. In the original workloads, the simple queries which cover the most of the workload, can quickly be answered with a high recall. While

⁸The box plot is shown in the most popular way with outliers [4].




Figure 11: Distribution of query hardness in the workloads.




Figure 12: Distribution of recall in the workloads.




Figure 13: Evaluate indexes on original query workloads.




Figure 14: Evaluate indexes on unbiased query workloads.

the recall of hard queries is gradually improved as NDC increases. On the contrary, in our unbiased workload, when NDC is small, the recall values of the queries are uniformly distributed.

7.5 Index Evaluation on New Workloads

To comprehensively show the query performance, we design a new metric that shows the distribution of NDC when all the queries in the workload reach a certain recall in Figure 13 and 14, for the old and the new unbiased workloads, respectively. The best method w.r.t. the median value under each setting is marked with a star. We summarize the major insights as follows.

- (1) *The performance significantly downgrades when using new unbiased workloads compared to the old one on real datasets.* The NDC range (in y-axis) increases >10 and >3 times on Deep and Glove datasets respectively, from Figure 13 to 14.
- (2) *As the recall target increases, the performance variance of graph indexes goes larger.* For example, on Deep dataset and HNSW index, after removing outliers, when recall=86%, the range of NDC spans from 491 to 14,765, and 718 to 38,288 when recall=86%. This indicates that high recall targets can lead to a significant effort for the graph index, as has been observed before [71, 97].
- (3) *Latest indexes show improvement on old workloads, but can be beaten by HNSW or NSG on unbiased workloads.* In Figure 13, τ -MNG or DEG wins the first or second place on most settings on real datasets even with our new metric. While in Figure 14, they become inferior to HNSW and NSG w.r.t. the median value or the variance. This indicates that latest indexes might “overfit” to the old biased workloads, while the classical HNSW and NSG remain in the pool of graph index choices.

- (4) *Advanced indexes do not always outperform KGraph.* On some settings of Rand and Tiny10M datasets, KGraph is a strong competitor. Although the number of edges in KGraph is 2.5x more than the other graphs to reach the best performance, its better navigability leads to a superior overall query performance.

Performance comparison between simple and hard queries.

We select 1,000 simplest and hardest queries from the new workload and benchmark the indexes. As shown in Figure 15, even though all the indexes can answer hard queries, the time cost is higher than simple queries from **10** to **50** times (see the differences of the x-axis). For the ranking result of simple queries, all indexes perform well on three datasets, except for KGraph on Deep and GIST. While for hard queries, τ -MNG and DEG show unstable performance on some cases. HNSW is the most robust followed by NSG.

Discussion. According to our evaluation, the major problem of current graph indexes is the stability on hard queries. The huge performance variance hinders their effects in practice. KGraph surprisingly works well on hard datasets, demonstrating that current edge pruning rules might not always work well. Besides, since the same hard query could be answered with variable efficiency on different index instances, combining multiple index instances is also a promising idea to improve the reachability of the graph.

7.6 Ablation Study of ME

In Figure 16, we present ablation study results to understand the effect of the three constraints of ME in Section 4.2. Figure 16 (b)-(d) show the deterioration of the correlation between the estimated effort and the actual query effort. We first remove the constraint on entry points (Constraint ①) by setting $p = 1/k$, that is, using the optimal entry point in N_k to get the ME. As shown in Figure 16(b),

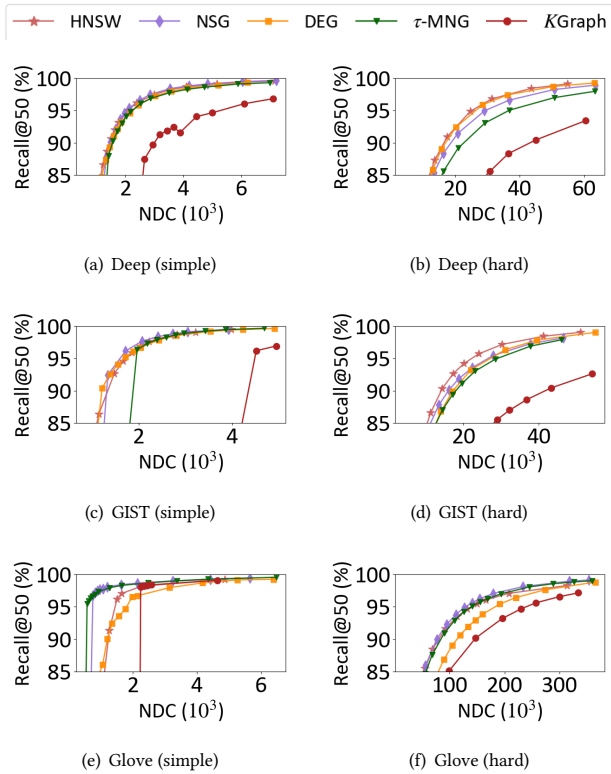


Figure 15: Benchmark results on 1,000 simple and hard queries from our unbiased workloads, respectively.

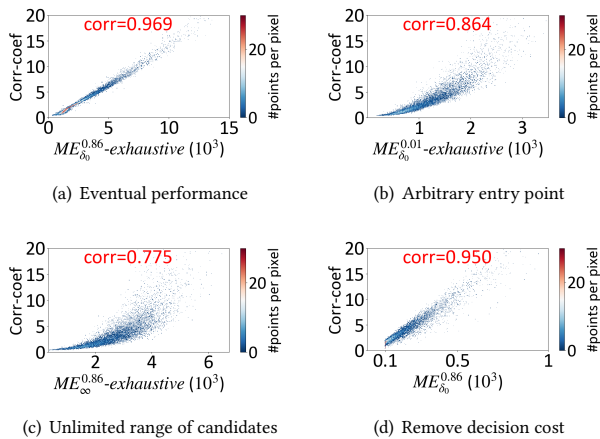


Figure 16: Ablation studies on the definitions of ME.

the correlation becomes weak for hard queries. Figure 16(c) shows a similar result where the limitation on the accessible candidates (Constraint ②) is removed. When removing the decision cost (i.e., Definition 3), as shown in Figure 16(d), simple queries cannot be distinguished. In brevity, the three proposed constraints are all necessary to describe the query effort.

Table 2: One-off pre-processing time of Steiner-hardness

	Deep	GIST	Rand	Glove	Tiny10M	Deep1B
Time (min)	0.9	3.4	9.6	66.8	34.5	16.9 hrs




Figure 17: The time of calculating the hardness 1,000 queries.

7.7 Efficiency and Scalability

We finally test the efficiency and the scalability of our approach.

One-off pre-processing. To calculate *Steiner*-hardness, for each dataset, we need to build the approximate MRNG and its reversed graph as pre-processing. This one-off pre-processing time is reported in Table 2. This time is dominated by building a *KGraph* (>90%) while building MRNG from the *KGraph* is very fast.

Hardness calculation. We report the overall time of calculating the *Steiner*-hardness of 1,000 queries in Figure 17(a), and compare it against the baselines on all the datasets. It is observed that on five million-sized datasets, the time to compute different hardness measures is very close except for RC. Figure 17(b) shows all the hardness measures can be computed with a linear scalability.

Workload generation. Generating the new unbiased workload includes three steps: (1) GMM training and inference, (2) hardness calculation, and (3) query selection. The first step can be done on a million-sized sample or smaller, which costs several to tens of minutes. The time cost of the second step is shown in Figure 17(a) while the third step only costs a constant time.

8 CONCLUSIONS AND FUTURE WORK

In this paper, we propose a practical query effort analyses framework for graph-based ANN indexes. This framework effectively describes the minimum effort under different recall targets and k on a given graph. We further design a novel connection-based query hardness measure, *Steiner*-hardness, based on the framework for graph indexes with efficient algorithms. Moreover, we build unbiased workloads encompassing the entire spectrum of *Steiner*-hardness to fairly stress-test current graph-based indexes.

In our future work, we plan to conduct a more comprehensive evaluation with unbiased query workloads to benchmark both graph- and partition-based indexes, and recommend indexes based on different user requirements.

ACKNOWLEDGMENTS

This work was partially funded by EU project AI4Europe (101070000).

REFERENCES

- [1] 2018. PACE (Parameterized Algorithms and Computational Experiments) 2018-Steiner Tree . <https://pacechallenge.org/2018/steiner-tree/>. Accessed on January 30, 2024.
- [2] 2023. Dijkstra's algorithm. https://en.wikipedia.org/w/index.php?title=Dijkstra%27s_algorithm&oldid=1189493285 [Online; accessed 11-January-2024].
- [3] 2023. Mixture model. https://en.wikipedia.org/w/index.php?title=Mixture_model&oldid=1166525726. [Online; accessed 7-January-2024].
- [4] 2024. Box plot – Wikipedia, The Free Encyclopedia. https://en.wikipedia.org/w/index.php?title=Box_plot&oldid=1225026016. [Online; accessed 1-June-2024].
- [5] 2024. Mahalanobis distance. https://en.wikipedia.org/w/index.php?title=Mahalanobis_distance&oldid=1193836393 [Online; accessed 12-January-2024].
- [6] Thomas D. Ahle, Martin Aumüller, and Rasmus Pagh. [n.d.]. *Parameter-free Locality Sensitive Hashing for Spherical Range Reporting*. 239–256. <https://doi.org/10.1137/1.9781611974782.16>
- [7] Laurent Amsaleg, Oussama Chelly, Teddy Furon, Stéphane Girard, Michael E. Houle, Ken-ichi Kawarabayashi, and Michael Nett. 2015. Estimating Local Intrinsic Dimensionality. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (Sydney, NSW, Australia). 29–38.
- [8] Laurent Amsaleg and Jégou Hervé. [n.d.]. Evaluation of Approximate nearest neighbors: large datasets. <http://corpus-textmex.irisa.fr/> [Online; accessed 15-May-2024].
- [9] 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. *PVLDB* 11, 8 (2018).
- [10] Martin Aumüller and Matteo Ceccarello. 2023. Recent Approaches and Trends in Approximate Nearest Neighbor Search, with Remarks on Benchmarking. *IEEE Data Eng. Bull.* 46, 3 (2023), 89–105.
- [11] Martin Aumüller, Erik Bernhardsson, and Alexander Faithfull. 2020. ANN-Benchmarks: A benchmarking tool for approximate nearest neighbor algorithms. *Information Systems* 87 (2020), 101374. <https://doi.org/10.1016/j.is.2019.02.006>
- [12] Martin Aumüller and Matteo Ceccarello. 2021. The role of local dimensionality measures in benchmarking nearest neighbor search. *Information Systems* 101 (2021), 101807. <https://doi.org/10.1016/j.is.2021.101807>
- [13] Ilias Azizi, Karima Echihabi, and Themis Palpanas. 2023. Elpis: Graph-Based Similarity Search for Scalable Data Science. *Proc. VLDB Endow.* 16, 6 (2023), 1548–1559. <https://doi.org/10.14778/3583140.3583166>
- [14] Dmitry Baranchuk and Artem Babenko. [n.d.]. Benchmarks for Billion-Scale Similarity Search. <https://research.yandex.com/blog/benchmarks-for-billion-scale-similarity-search> [Online; accessed 15-May-2024].
- [15] Eric Bernhardsson. [n.d.]. Annoy. <https://github.com/spotify/>. [Online; accessed 7-January-2024].
- [16] Moses Charikar, Chandra Chekuri, To yat Cheung, Zuo Dai, Ashish Goel, Sudipto Guha, and Ming Li. 1999. Approximation Algorithms for Directed Steiner Problems. *Journal of Algorithms* 33, 1 (1999), 73–91. <https://doi.org/10.1006/jagm.1999.1042>
- [17] Manos Chatzakis, Panagiota Fatourou, Eleftherios Kosmas, Themis Palpanas, and Botao Peng. 2023. Odyssey: A Journey in the Land of Distributed Data Series Similarity Search. *Proc. VLDB Endow.* 16, 5 (2023), 1140–1153. <https://doi.org/10.14778/3579075.3579087>
- [18] Lu Chen, Yunjun Gao, Xinhan Li, Christian S Jensen, and Gang Chen. 2015. Efficient metric indexing for similarity search. In *ICDE*. 591–602.
- [19] Lu Chen, Yunjun Gao, Baihua Zheng, Christian S. Jensen, Hanyu Yang, and Keyu Yang. 2017. Pivot-Based Metric Indexing. *Proc. VLDB Endow.* 10, 10 (jun 2017), 1058–1069. <https://doi.org/10.14778/3115404.3115411>
- [20] Patrick Chen, Wei-Cheng Chang, Jyun-Yu Jiang, 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*). ACM, New York, NY, USA, 3225–3235. <https://doi.org/10.1145/3543507.3583318>
- [21] 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 Neighborhood Search. In *Advances in Neural Information Processing Systems*, Vol. 34. Curran Associates, Inc., 5199–5212.
- [22] James Cheng, Xinyan Dai, Bao Ergute, Ti-chung Cheng, Haopeng Sun, Bin Hu, Xiao Yan, Jinfei Li, and Jie Liu. [n.d.]. GQR: A General and Efficient Querying Method for Learning to Hash. <https://www.cse.cuhk.edu.hk/systems/hash/gqr/datasets.html> [Online; accessed 15-May-2024].
- [23] Eden Chlamtáč, Michael Dinitz, Guy Kortsarz, and Bundit Laekhanukit. 2020. Approximating Spanners and Directed Steiner Forest: Upper and Lower Bounds. *ACM Trans. Algorithms* 16, 3, Article 33 (jun 2020), 31 pages. <https://doi.org/10.1145/3381451>
- [24] Eden Chlamtáč, Michael Dinitz, Guy Kortsarz, and Bundit Laekhanukit. 2020. Approximating Spanners and Directed Steiner Forest: Upper and Lower Bounds. *ACM Trans. Algorithms* 16, 3, Article 33 (jun 2020), 31 pages. <https://doi.org/10.1145/3381451>
- [25] Sanjoy Dasgupta and Yoav Freund. 2008. Random Projection Trees and Low Dimensional Manifolds. In *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing* (Victoria, British Columbia, Canada) (*STOC '08*). ACM, New York, NY, USA, 537–546. <https://doi.org/10.1145/1374376.1374452>
- [26] D.W. Dearholt, N. Gonzales, and G. Kurup. 1988. Monotonic Search Networks For Computer Vision Databases. In *Twenty-Second Asilomar Conference on Signals, Systems and Computers*, Vol. 2. 548–553. <https://doi.org/10.1109/ACSSC.1988.754602>
- [27] D.W. Dearholt, N. Gonzales, and G. Kurup. 1988. Monotonic Search Networks For Computer Vision Databases. In *Twenty-Second Asilomar Conference on Signals, Systems and Computers*, Vol. 2. 548–553. <https://doi.org/10.1109/ACSSC.1988.754602>
- [28] Michael Dinitz, Guy Kortsarz, and Zeev Nutov. 2017. Improved Approximation Algorithm for Steiner K-Forest with Nearly Uniform Weights. *ACM Trans. Algorithms* 13, 3, Article 40 (jul 2017), 16 pages. <https://doi.org/10.1145/3077581>
- [29] Magdalen Dobson, Zheqi Shen, Guy E. Blelloch, Laxman Dhulipala, Yan Gu, Harsha Vardhan Simhadri, and Yihan Sun. 2023. Scaling Graph-Based ANNS Algorithms to Billion-Size Datasets: A Comparative Analysis. [arXiv:2305.04359](https://arxiv.org/abs/2305.04359) [cs.IR]
- [30] Wei Dong, Charikar Moses, and Kai Li. 2011. Efficient K-Nearest Neighbor Graph Construction for Generic Similarity Measures. In *Proceedings of the 20th International Conference on World Wide Web* (Hyderabad, India). 577–586. <https://doi.org/10.1145/1963405.1963487>
- [31] Ishita Doshi, Dhritiman Das, Ashish Bhutani, Rajeev Kumar, Rushi Bhatt, and Niranjan Balasubramanian. 2021. LANNs: A Web-Scale Approximate Nearest Neighbor Lookup System. *Proc. VLDB Endow.* 15, 4 (dec 2021), 850–858. <https://doi.org/10.14778/3503585.3503594>
- [32] Karima Echihabi, Panagiota Fatourou, Kostas Zoumpatianos, Themis Palpanas, and Houda Benbrahim. 2022. Hercules Against Data Series Similarity Search. *PVLDB* 15, 10 (2022).
- [33] Karima Echihabi, Kostas Zoumpatianos, Themis Palpanas, and Houda Benbrahim. 2018. The Lernaean Hydra of Data Series Similarity Search: An Experimental Evaluation of the State of the Art. *Proc. VLDB Endow.* 12, 2 (2018), 112–127. <https://doi.org/10.14778/3282495.3282498>
- [34] Karima Echihabi, Kostas Zoumpatianos, Themis Palpanas, and Houda Benbrahim. 2019. Return of the Lernaean Hydra: Experimental Evaluation of Data Series Approximate Similarity Search. *PVLDB* 13, 3 (2019), 403–420.
- [35] Jon Feldman and Matthias Ruhl. 2006. The Directed Steiner Network Problem is Tractable for a Constant Number of Terminals. *SIAM J. Comput.* 36, 2 (2006), 543–561. <https://doi.org/10.1137/S0097539704441241>
- [36] Moran Feldman, Guy Kortsarz, and Zeev Nutov. 2012. Improved approximation algorithms for Directed Steiner Forest. *J. Comput. System Sci.* 78, 1 (2012), 279–292. <https://doi.org/10.1016/j.jcss.2011.05.009> JCSS Knowledge Representation and Reasoning.
- [37] 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>
- [38] 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.
- [39] Junhao Gan, Jianlin Feng, Qiong Fang, and Wilfred Ng. 2012. Locality-sensitive hashing scheme based on dynamic collision counting. In *SIGMOD*. 541–552.
- [40] 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 (jun 2023), 27 pages. <https://doi.org/10.1145/3589282>
- [41] Jianyang Gao and Cheng Long. 2024. RaBitQ: Quantizing High-Dimensional Vectors with a Theoretical Error Bound for Approximate Nearest Neighbor Search. *Proc. ACM Manag. Data* 2, 3, Article 167 (may 2024), 27 pages. <https://doi.org/10.1145/3654970>
- [42] Aristides Gionis, Piotr Indyk, and Rajeev Motwani. 1999. Similarity Search in High Dimensions via Hashing. In *PVLDB*. 518–529.
- [43] Zhen Gong, Xin Wu, Lei Chen, Zhenzhe Zheng, Shengjie Wang, Anran Xu, Chong Wang, and Fan Wu. 2023. Full Index Deep Retrieval: End-to-End User and Item Structures for Cold-start and Long-tail Item Recommendation. In *Proceedings of the 17th ACM Conference on Recommender Systems* (Singapore, Singapore) (*RecSys '23*). Association for Computing Machinery, New York, NY, USA, 47–57.
- [44] Fabrizio Grandoni, Bundit Laekhanukit, and Shi Li. 2019. O(Log 2 k / Log Log k)-Approximation Algorithm for Directed Steiner Tree: A Tight Quasi-Polynomial-Time Algorithm. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing* (Phoenix, AZ, USA) (*STOC 2019*). 253–264. <https://doi.org/10.1145/3313276.3316349>

- [45] 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>
- [46] Ben Harwood and Tom Drummond. 2016. FannG: Fast approximate nearest neighbour graphs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 5713–5722.
- [47] M Hauptmann and M Karpinski. 2013. A Compendium on Steiner Tree Problems. (2013).
- [48] Nico Hezel, Kai Uwe Barthel, Konstantin Schall, and Klaus Jung. 2023. Fast Approximate Nearest Neighbor Search with a Dynamic Exploration Graph using Continuous Refinement. arXiv:2307.10479 [cs.IR]
- [49] Kohei Hirata, Daichi Amagata, Sumio Fujita, and Takahiro Hara. 2022. Solving Diversity-Aware Maximum Inner Product Search Efficiently and Effectively. In *Proceedings of the 16th ACM Conference on Recommender Systems (RecSys '22)*. Association for Computing Machinery, New York, NY, USA, 198–207. <https://doi.org/10.1145/3523227.3546779>
- [50] Michael E. Houle. 2013. Dimensionality, Discriminability, Density and Distance Distributions. In *2013 IEEE 13th International Conference on Data Mining Workshops*. 468–473. <https://doi.org/10.1109/ICDMW.2013.139>
- [51] Michael E. Houle. 2017. Local Intrinsic Dimensionality I: An Extreme-Value-Theoretic Foundation for Similarity Applications. In *Similarity Search and Applications*. Springer International Publishing, Cham, 64–79.
- [52] Michael E. Houle, Hisashi Kashima, and Michael Nett. 2012. Generalized Expansion Dimension. In *2012 IEEE 12th International Conference on Data Mining Workshops*. 587–594. <https://doi.org/10.1109/ICDMW.2012.94>
- [53] Qiang Huang, Jianlin Feng, Yikai Zhang, Qiong Fang, and Wilfred Ng. 2015. Query-aware locality-sensitive hashing for approximate nearest neighbor search. *PVLDB* 9, 1 (2015), 1–12.
- [54] Piotr Indyk and Haiké Xu. 2023. Worst-case Performance of Popular Approximate Nearest Neighbor Search Implementations: Guarantees and Limitations. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- [55] Ishindaniil. [n.d.]. directed-steiner-tree. <https://github.com/ishindaniil/directed-steiner-tree>. Accessed: 2024-01-30.
- [56] Masajiro Iwasaki and Daisuke Miyazaki. 2018. Optimization of Indexing Based on k-Nearest Neighbor Graph for Proximity Search in High-dimensional Data. arXiv:1810.07355 [cs.DB]
- [57] 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]
- [58] Suhas Jayaram Subramanya, Fnu Devvrit, Harsha Vardhan Simhadri, Ravishankar Krishnaswamy, and Rohan Kadekodi. 2019. DiskANN: Fast Accurate Billion-point Nearest Neighbor Search on a Single Node. In *Advances in Neural Information Processing Systems*, Vol. 32. Curran Associates, Inc.
- [59] Herve Jegou, Matthijs Douze, and Cordelia Schmid. 2010. Product quantization for nearest neighbor search. *TPAMI* 33, 1 (2010).
- [60] Jeff Johnson, Matthijs Douze, and Hervé Jégou. 2021. Billion-Scale Similarity Search with GPUs. *IEEE Transactions on Big Data* 7, 3 (2021), 535–547. <https://doi.org/10.1109/TBDATA.2019.2921572>
- [61] David R. Karger and Matthias Ruhl. 2002. Finding Nearest Neighbors in Growth-Restricted Metrics. In *Proceedings of the Thirty-Fourth Annual ACM Symposium on Theory of Computing (Montreal, Quebec, Canada) (STOC '02)*. ACM, New York, NY, USA, 741–750. <https://doi.org/10.1145/509907.510013>
- [62] Jon Kleinberg. 2000. The small-world phenomenon: An algorithmic perspective. In *Proceedings of the thirty-second annual ACM symposium on Theory of computing*. 163–170.
- [63] Thijs Laarhoven. 2017. Graph-based time-space trade-offs for approximate near neighbors. CoRR abs/1712.03158 (2017).
- [64] Elizaveta Levina and Peter Bickel. 2004. Maximum Likelihood Estimation of Intrinsic Dimension. In *Advances in Neural Information Processing Systems*, Vol. 17. MIT Press.
- [65] 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 (Portland, OR, USA) (SIGMOD '20)*. ACM, New York, NY, USA, 2539–2554. <https://doi.org/10.1145/3318464.3380600>
- [66] Ke Li and Jitendra Malik. 2017. Fast k-Nearest Neighbour Search via Prioritized DCl. In *Proceedings of the 34th International Conference on Machine Learning (Proceedings of Machine Learning Research)*, Vol. 70. PMLR, 2081–2090.
- [67] Wen Li, Ying Zhang, Yifang Sun, Wei Wang, Mingjie Li, Wenjie Zhang, and Xuemin Lin. 2020. Approximate Nearest Neighbor Search on High Dimensional Data – Experiments, Analyses, and Improvement. *IEEE Transactions on Knowledge and Data Engineering* 32, 8 (2020), 1475–1488. <https://doi.org/10.1109/TKDE.2019.2909204>
- [68] Peng-Cheng Lin and Wan-Lei Zhao. 2019. Graph based Nearest Neighbor Search: Promises and Failures. arXiv:1904.02077 [cs.IR]
- [69] Michele Linardi and Themis Palpanas. 2018. Scalable, Variable-Length Similarity Search in Data Series: The ULISSE Approach. *Proc. VLDB Endow.* 11, 13 (2018), 2236–2248.
- [70] Michele Linardi and Themis Palpanas. 2020. Scalable data series subsequence matching with ULISSE. *VLDB J.* 29, 6 (2020), 1449–1474.
- [71] Jun Liu, Zhenhua Zhu, Jingbo Hu, Hanbo Sun, Li Liu, Lingzhi Liu, Guohao Dai, Huazhong Yang, and Yu Wang. 2022. Optimizing Graph-based Approximate Nearest Neighbor Search: Stronger and Smarter. In *2022 23rd IEEE International Conference on Mobile Data Management (MDM)*. 179–184. <https://doi.org/10.1109/MDM55031.2022.00045>
- [72] Kejing Lu and Mineichi Kudo. 2020. R2LSH: A Nearest Neighbor Search Scheme Based on Two-dimensional Projected Spaces. In *ICDE*. 1045–1056.
- [73] Kejing Lu, Mineichi Kudo, Chuan Xiao, and Yoshiharu Ishikawa. 2021. HVS: Hierarchical Graph Structure Based on Voronoi Diagrams for Solving Approximate Nearest Neighbor Search. *Proc. VLDB Endow.* 15, 2 (oct 2021), 246–258. <https://doi.org/10.14778/3489496.3489506>
- [74] Malte Ludevig, Iman Kamehkhosh, Nick Landia, and Dietmar Jannach. 2018. Effective Nearest-Neighbor Music Recommendations. In *Proceedings of the ACM Recommender Systems Challenge 2018 (Vancouver, BC, Canada) (RecSys Challenge '18)*. Article 3, 6 pages. <https://doi.org/10.1145/3267471.3267474>
- [75] 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.
- [76] Yu A. Malkov and D. A. Yashunin. 2020. Efficient and Robust Approximate Nearest Neighbor Search Using Hierarchical Navigable Small World Graphs. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 42, 4 (2020), 824–836. <https://doi.org/10.1109/TPAMI.2018.2889473>
- [77] Marius Muja and David G. Lowe. 2014. Scalable Nearest Neighbor Algorithms for High Dimensional Data. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 36, 11 (2014), 2227–2240. <https://doi.org/10.1109/TPAMI.2014.2321376>
- [78] Botao Peng, Panagiota Fatourou, and Themis Palpanas. 2021. Fast data series indexing for in-memory data. *VLDB J.* 30, 6 (2021), 1041–1067. <https://doi.org/10.1007/S00778-021-00677-2>
- [79] Botao Peng, Panagiota Fatourou, and Themis Palpanas. 2021. SING: Sequence Indexing Using GPUs. In *37th IEEE International Conference on Data Engineering, ICDE 2021, Chania, Greece, April 19-22, 2021*. IEEE, 1883–1888. <https://doi.org/10.1109/ICDE51399.2021.00171>
- [80] 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>
- [81] Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. GloVe: Global Vectors for Word Representation. In *Empirical Methods in Natural Language Processing (EMNLP)*. 1532–1543. <http://www.aclweb.org/anthology/D14-1162>
- [82] Liudmila Prokhorenkova and Aleksandr Shekhovtsov. 2020. Graph-based Nearest Neighbor Search: From Practice to Theory. In *Proceedings of the 37th International Conference on Machine Learning (Proceedings of Machine Learning Research)*, Vol. 119. 7803–7813.
- [83] Miloš Radovanović, Alexandros Nanopoulos, and Mirjana Ivanović. 2009. Nearest neighbors in high-dimensional data: the emergence and influence of hubs. In *Proceedings of the 26th Annual International Conference on Machine Learning (Montreal, Quebec, Canada) (ICML '09)*. Association for Computing Machinery, New York, NY, USA, 865–872. <https://doi.org/10.1145/1553374.1553485>
- [84] Miloš Radovanović, Alexandros Nanopoulos, and Mirjana Ivanović. 2010. Hubs in Space: Popular Nearest Neighbors in High-Dimensional Data. *J. Mach. Learn. Res.* 11 (dec 2010), 2487–2531.
- [85] Thomas Rothvoß. 2012. Directed Steiner Tree and the Lasserre Hierarchy. arXiv:1111.5473 [cs.DS]
- [86] A. Rozza, G. Lombardi, C. Ceruti, E. Casiraghi, and P. Campadelli. 2012. Novel High Intrinsic Dimensionality Estimators. *Mach. Learn.* 89, 1–2 (oct 2012), 37–65. <https://doi.org/10.1007/s10994-012-5294-7>
- [87] Jin Shieh and Eamonn Keogh. 2008. iSAX: indexing and mining terabyte sized time series. In *SIGKDD*. 623–631.
- [88] Xiaoxin Tang, Zhiyi Huang, David Ebers, Steven Mills, and Minyi Guo. 2015. Efficient Selection Algorithm for Fast k-NN Search on GPUs. In *2015 IEEE International Parallel and Distributed Processing Symposium*. 397–406. <https://doi.org/10.1109/IPDPS.2015.115>
- [89] Yao Tian, Tingyun Yan, Xi Zhao, Kai Huang, and Xiaofang Zhou. 2023. A Learned Index for Exact Similarity Search in Metric Spaces. *IEEE Transactions on Knowledge and Data Engineering* 35, 8 (2023), 7624–7638. <https://doi.org/10.1109/TKDE.2022.3206441>
- [90] Godfried T. Toussaint. 1980. The relative neighbourhood graph of a finite planar set. *Pattern Recognition* 12, 4 (1980), 261–268. [https://doi.org/10.1016/0031-3203\(80\)90066-7](https://doi.org/10.1016/0031-3203(80)90066-7)
- [91] Saskia van der Hoeven. 2023. Efficient solution methods for the directed Steiner tree problem.

- [92] Javier Vargas Muñoz, Marcos A. Gonçalves, Zanoni Dias, and Ricardo da S. Torres. 2019. Hierarchical Clustering-Based Graphs for Large Scale Approximate Nearest Neighbor Search. *Pattern Recognition* 96 (2019), 106970. <https://doi.org/10.1016/j.patcog.2019.106970>
- [93] Polychronis Velentzas, Panagiotis Moutafis, and George Mavrommatis. 2021. An Improved GPU-Based Algorithm for Processing the k Nearest Neighbor Query. In *Proceedings of the 24th Pan-Hellenic Conference on Informatics* (Athens, Greece) (*PCI '20*). ACM, New York, NY, USA, 372–375. <https://doi.org/10.1145/3437120.3437343>
- [94] Polychronis Velentzas, Michael Vassilakopoulos, and Antonio Corral. 2021. GPU-aided edge computing for processing the k nearest-neighbor query on SSD-resident data. *Internet of Things* 15 (2021), 100428. <https://doi.org/10.1016/j.iot.2021.100428>
- [95] Skoltech Computer Vision. [n.d.]. Deep billion-scale indexing. <http://sites.skoltech.ru/compvision/noimi> Accessed March 14, 2022.
- [96] Hui Wang, Yong Wang, and Wan-Lei Zhao. 2022. Graph-based Approximate NN Search: A Revisit. arXiv:2204.00824 [cs.IR]
- [97] Hongya Wang, Zhizheng Wang, Wei Wang, Yingyuan Xiao, Zeng Zhao, and Kaixiang Yang. 2020. A Note on Graph-Based Nearest Neighbor Search. arXiv:2012.11083 [cs.LG]
- [98] 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. In *Proceedings of the 2021 International Conference on Management of Data* (Virtual Event, China) (*SIGMOD '21*). ACM, New York, NY, USA, 2614–2627. <https://doi.org/10.1145/3448016.3457550>
- [99] 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, 11 (jul 2021), 1964–1978. <https://doi.org/10.14778/3476249.3476255>
- [100] Qitong Wang and Themis Palpanas. 2021. Deep Learning Embeddings for Data Series Similarity Search. In *KDD '21: The 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Virtual Event, Singapore, August 14-18, 2021*, Feida Zhu, Beng Chin Ooi, and Chunyan Miao (Eds.). ACM, 1708–1716. <https://doi.org/10.1145/3447548.3467317>
- [101] Qitong Wang and Themis Palpanas. 2023. SEAnet: A Deep Learning Architecture for Data Series Similarity Search. *IEEE Trans. Knowl. Data Eng.* 35, 12 (2023), 12972–12986. <https://doi.org/10.1109/TKDE.2023.3270264>
- [102] Yang Wang, Peng Wang, Jian Pei, Wei Wang, and Sheng Huang. 2013. A data-adaptive and dynamic segmentation index for whole matching on time series. *PVLDB* 6, 10 (2013), 793–804.
- [103] Zeyu Wang, Peng Wang, Themis Palpanas, and Wei Wang. 2023. Graph-and Tree-based Indexes for High-dimensional Vector Similarity Search: Analyses, Comparisons, and Future Directions. *IEEE Data Eng. Bull.* 46, 3 (2023), 3–21.
- [104] Zeyu Wang, Qitong Wang, Peng Wang, Themis Palpanas, and Wei Wang. 2023. Dumpy: A Compact and Adaptive Index for Large Data Series Collections. *Proc. ACM Manag. Data* 1, 1, Article 111 (may 2023), 27 pages. <https://doi.org/10.1145/3588965>
- [105] Zeyu Wang, Qitong Wang, Peng Wang, Themis Palpanas, and Wei Wang. 2024. DumpyOS: A data-adaptive multi-ary index for scalable data series similarity search. *The VLDB Journal* (2024), 1–25.
- [106] Zeyu Wang, Haoran Xiong, Zhenying He, Peng Wang, and Wei Wang. 2024. Distance Comparison Operators for Approximate Nearest Neighbor Search: Exploration and Benchmark. arXiv:2403.13491 [cs.DB] <https://arxiv.org/abs/2403.13491>
- [107] Dimitri Watel and Marc-Antoine Weisser. 2016. A Practical Greedy Approximation for the Directed Steiner Tree Problem. *J. Comb. Optim.* 32, 4 (nov 2016), 1327–1370. <https://doi.org/10.1007/s10878-016-0074-0>
- [108] 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>
- [109] Jiuqi Wei, Botao Peng, Xiaodong Lee, and Themis Palpanas. 2024. DET-LSH: A Locality-Sensitive Hashing Scheme with Dynamic Encoding Tree for Approximate Nearest Neighbor Search. *Proceedings of the VLDB Endowment* 17, 9 (2024), 2241–2254.
- [110] Haoran Xiong, Hang Zhang, Zeyu Wang, Zhenying He, Peng Wang, and X. Sean Wang. 2024. CIVET: Exploring Compact Index for Variable-Length Subsequence Matching on Time Series. *Proc. VLDB Endow.* 17, 9 (aug 2024), 2123–2135.
- [111] Yuming Xu, Hengyu Liang, Jin Li, Shuotao Xu, Qi Chen, Qianxi Zhang, Cheng Li, Ziyue Yang, Fan Yang, Yuqing Yang, Peng Cheng, and Mao Yang. 2023. SPFresh: Incremental In-Place Update for Billion-Scale Vector Search. In *Proceedings of the 29th Symposium on Operating Systems Principles* (Koblenz, Germany) (*SOSP '23*). ACM, New York, NY, USA, 545–561. <https://doi.org/10.1145/3600006.3613166>
- [112] Djamel Edine Yagoubi, Reza Akbarinia, Florent Masseglia, and Themis Palpanas. 2017. DPiSAX: Massively Distributed Partitioned iSAX. In *ICDM*.
- [113] Djamel Edine Yagoubi, Reza Akbarinia, Florent Masseglia, and Themis Palpanas. 2020. Massively Distributed Time Series Indexing and Querying. *IEEE Trans. Knowl. Data Eng. (TKDE)* 32, 1 (2020), 108–120.
- [114] Minjia Zhang, Wenhan Wang, and Yuxiong He. 2022. GraSP: Optimizing Graph-Based Nearest Neighbor Search with Subgraph Sampling and Pruning. In *Proceedings of the Fifteenth ACM International Conference on Web Search and Data Mining* (Virtual Event, AZ, USA) (*WSDM '22*). ACM, New York, NY, USA, 1395–1405. <https://doi.org/10.1145/3488560.3498425>
- [115] Pengcheng Zhang, Bin Yao, Chao Gao, Bin Wu, Xiao He, Feifei Li, Yuanfei Lu, Chaoqun Zhan, and Feilong Tang. 2023. Learning-based query optimization for multi-probe approximate nearest neighbor search. *The VLDB Journal* 32, 3 (2023), 623–645.
- [116] 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.
- [117] Xi Zhao, Yao Tian, Kai Huang, Bolong Zheng, and Xiaofang Zhou. 2023. Towards Efficient Index Construction and Approximate Nearest Neighbor Search in High-Dimensional Spaces. *Proc. VLDB Endow.* 16, 8 (apr 2023), 1979–1991. <https://doi.org/10.14778/3594512.3594527>
- [118] Bolong Zheng, Xi Zhao, Lianggui Weng, Nguyen Quoc Viet Hung, Hang Liu, and Christian S. Jensen. 2020. PM-LSH: A Fast and Accurate LSH Framework for High-Dimensional Approximate NN Search. *Proc. VLDB Endow.* 13, 5 (jan 2020), 643–655. <https://doi.org/10.14778/3377369.3377374>
- [119] Shuyan Zhou, Uri Alon, Frank F. Xu, Zhengbao Jiang, and Graham Neubig. 2023. DocPrompting: Generating Code by Retrieving the Docs. In *The Eleventh International Conference on Learning Representations*.
- [120] Dantong Zhu and Minjia Zhang. 2021. Understanding and Generalizing Monotonic Proximity Graphs for Approximate Nearest Neighbor Search. arXiv:2107.13052 [cs.IR]
- [121] Yifan Zhu, Lu Chen, Yunjun Gao, and Christian S Jensen. 2022. Pivot selection algorithms in metric spaces: a survey and experimental study. *The VLDB Journal* (2022), 1–25.
- [122] Leonid Zosin and Samir Khuller. 2002. On Directed Steiner Trees. In *Proceedings of the Thirteenth Annual ACM-SIAM Symposium on Discrete Algorithms* (San Francisco, California) (*SODA '02*). 59–63.
- [123] Kostas Zoumpatianos, Yin Lou, Ioana Ileana, Themis Palpanas, and Johannes Gehrke. 2018. Generating data series query workloads. *The VLDB Journal* 27 (2018), 823–846.
- [124] Kostas Zoumpatianos, Yin Lou, Themis Palpanas, and Johannes Gehrke. 2015. Query Workloads for Data Series Indexes. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (Sydney, NSW, Australia). 1603–1612.