

CAPS: A Practical Partition Index for Filtered Similarity Search

Gaurav Gupta
gaurav.gupta@rice.edu
Rice University
Houston, TX, USA

Vihan Lakshman
vihan.lakshman@gmail.com
ThirdAI Corp.
Houston, TX, USA

Jonah Wonkyu Yi
jonah.w.yi@rice.edu
Rice University
Houston, TX, USA

Chen Luo
cheluo@amazon.com
Amazon Search
Palo Alto, CA, USA

Benjamin Coleman
benjamin.ray.coleman@gmail.com
Google Deepmind
Mountain View, CA, USA

Anshumali Shrivastava
anshumali@rice.edu
Rice University, ThirdAI Corp
Houston, TX, USA

ABSTRACT

With the surging popularity of approximate near-neighbor search (ANNS), driven by advances in neural representation learning, the ability to serve queries accompanied by a set of constraints has become an area of intense interest. While the community has recently proposed several algorithms for constrained ANNS, almost all of these methods focus on integration with graph-based indexes, the predominant class of algorithms achieving state-of-the-art performance in latency-recall tradeoffs. In this work, we take a different approach and focus on developing a constrained ANNS algorithm via space partitioning as opposed to graphs. To that end, we introduce Constrained Approximate Partitioned Search (CAPS), an index for ANNS with filters via space partitions that not only retains the benefits of a partition-based algorithm but also outperforms state-of-the-art graph-based constrained search techniques in recall-latency tradeoffs, with only 10% of the index size.

KEYWORDS

Constrained near neighbor, Hybrid Query Processing, Vector Databases

ACM Reference Format:

Gaurav Gupta, Jonah Wonkyu Yi, Benjamin Coleman, Vihan Lakshman, Chen Luo, and Anshumali Shrivastava. 2018. CAPS: A Practical Partition Index for Filtered Similarity Search. In *Proceedings of 3rd International Workshop on Interactive and Scalable Information Retrieval methods for eCommerce (ISIR-eCom) 2024 @ WSDM-2024*, <https://isir-ecom.github.io/> (ISIR-eCom 2024 @ WSDM-2024). ACM, New York, NY, USA, 11 pages. <https://doi.org/XXXXXXX.XXXXXXX>

1 INTRODUCTION

In recent years, the information retrieval community has seen a surge of interest in performant similarity search over dense feature representations. Fueled by the breakthrough success of deep neural networks in mapping raw, unstructured data, such as text and images, into semantically coherent embeddings, vector databases now underpin numerous applications in web search [18], product

recommendations [25, 27, 30], social networks [34], video analytics [39], computational biology [12, 16], and many more.

With the tremendous interest and commercial adoption of similarity search indexes, *Filtered Near Neighbor Search*, which involves retrieving near-neighbor vectors subject to a set of filters, has emerged as an essential feature in many applications. This setting frequently appears in the context of product search, where shoppers often issue queries coupled with constraints such as “free shipping” or “under \$10.” Various graph-based similarity search algorithms have been adapted to provide constraint satisfaction functionality. For example, AIRSHIP [40], NHQ [36], and DiskANN [13, 22] all incorporate the filters into the graph walk subroutine that forms the core of their respective search algorithms. However, these approaches all carry limitations (Table 1) that hinder their practical applicability, such as an inability to support a variable number of constraints per query or a lack of support for conjunctive ANDs (e.g. “free shipping AND under \$10”).

Algorithm	Variable Constraints	Conjunctive ANDs	Sparse Attributes
NHQ [36]	✗	✓	✗
AIRSHIP [40]	✓	✓	✗
DiskANN [13]	✓	✗	✓
CAPS (ours)	✓	✓	✓

Table 1: Summary of capabilities of open-source Filtered Near Neighbor Search algorithms.

While these aforementioned prior works in Filtered Near Neighbor Search focus on integrating with popular graph-traversal algorithms, many industry-scale similarity search databases still use space partition-based approaches due to their unique advantages [14, 17]. For example, graph-based ANNS algorithms can incur frequent cache misses due to random node access [8] and cannot easily be parallelized [28]. Conversely, partition-based algorithms are trivial to parallelize and have a cache and accelerator-friendly data access pattern. To our knowledge, however, there is currently no efficient algorithm or structure for constrained search over partitions short of filtered brute force.

Motivated by these challenges, we propose CAPS (Constrained Approximate Partitioned Search), a new filtered search algorithm based on a hierarchical partitioning that guarantees that the returned set of vectors matches all constraints while demonstrating improved efficiency in the high-recall regime.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ISIR-eCom 2024 @ WSDM-2024, March 8, 2024, Mérida, Mexico

© 2018 Association for Computing Machinery.
ACM ISBN 978-1-4503-XXXX-X/18/06...\$15.00
<https://doi.org/XXXXXXX.XXXXXXX>

1.1 Contributions

We begin with an exploratory experiment on the nature of filtered near-neighbor search. Additionally, we develop methods to address uncorrelated attributes and conjunctive predicates (i.e. AND constraints). We make the following specific contributions.

Constrained Space Partitioning: To our knowledge, we present the first non-trivial approximate algorithm for filtered near neighbor search via space partitions. One of the key problems with filtered search in a partitioned index is that most indexes have large partitions that may only sparsely satisfy the constraint. To address this, we introduce a Huffman tree-inspired hierarchical sub-partitioning algorithm that shards the data based on vector embeddings followed by attributes, which we call an *Attribute Frequency Tree* (AFT). We also experiment with a learned strategy for jointly partitioning the embedding and attributes, which achieves superior results on certain benchmarks.

Practical Utility: We present, to our knowledge, the first efficient and practically useful constrained search algorithm that supports a varying number of attributes per query, conjunctive ANDs and support for rare attributes. Our proposed index also uses a fraction of the space compared to previously proposed graph-based filtered search algorithms.

Real-world Evaluation: Prior work in this area relies on synthetic tasks constructed from standard near-neighbor benchmarks with varying distributional assumptions on the attributes (e.g. clustered or random). We apply our methods to a large-scale industrial e-commerce application and find that real-world constraints are power-law distributed. We use the power law structure to enhance our sub-partition tree algorithm and further improve performance.

Table 2: Notation Used in the Paper

Symbol	Description	Symbol	Description
D	Dataset	x	data point vector
L	number of attributes	q	query
$C(\cdot)$	constraint function	N	Number of points
a	data point attributes	b	query attribute
$f(\cdot)$	assignment function	B	number of partitions
P	Partition	p	sub-partition
h	Huffman tree height	$A(\cdot)$	sub-partition attribute
m	num cluster to probe		

2 FILTERED NEAR NEIGHBOR SEARCH

Modern search applications require simultaneous similarity search of vector embedding data and filtering by categorical attributes. We call this framework *filtered near neighbor search* because it combines ideas from similarity search with more traditional relational database operations.

Attribute Model: We adopt the industry standard attribute model, where an attribute is a tag for an item that can take multiple categorical values. An item can have any number of attributes, each with one value of an attribute. We begin with a formal definition of Query Filter and Filtered Search.

DEFINITION 1. Query Filter
A query filter is a function $C : \mathbb{N}^L \times \mathbb{N}^L \rightarrow \{\text{True}, \text{False}\}$ such that,

given a data attribute vector $a \in \mathbb{N}^L$ and a query attribute vector $b \in \mathbb{N}^L$, $C(a, b) = \text{True}$ if $a[i] = b[i] \forall i$ and False otherwise.

The near-neighbor search problem has many theoretical formulations, leading to slightly different problem statements. For example, the (R, c) -approximate near neighbor problem asks that we retrieve all points within radius R of the query and none other ([9, 19]). The k -nearest neighbor problem [7] is to identify all of the k closest points to a query with high probability. We state the constrained problem in a way that is agnostic to the specific choice of the near-neighbor formulation.

DEFINITION 2. Filtered Near-Neighbor Search
Given a dataset $D = \{(x_1, a_1), \dots, (x_N, a_N)\}$ of embedding vectors $x_i \in \mathbb{R}^d$ and associated attributes $a_i \in \mathbb{N}^L$, a query q and a query filter $C(q, b)$, solve the near neighbor problem of interest on the restricted dataset $D_C = \{(x_i, a_i) \in D : C(b, a_i)\}$.

Definition 1 defines the query filter constraint, and Definition 2 defines the constrained similarity search problem as a *filtered* version of the standard similarity search. This is well-motivated by industrial applications, where practitioners often wish to perform a query with the SQL-style predicate “WHERE $\text{dist}(x, q) < R$ AND $x \in D_C$.”

3 CHALLENGES OF FILTERED SEARCH

In this section, we seek to understand the conditions under which existing constrained search strategies fail to retrieve results in a performing manner.

There is a fundamental tension between satisfying the constraints and taking the optimal path through the dataset for similarity search. While this could, in principle, be solved by constructing an index for each possible attribute (or by incorporating the attribute into the vector embedding), such solutions are insufficient in practice for two reasons.

- Often, the attributes are either unknown or of high cardinality, precluding algorithms that pre-index the data. In other words, D_C must be identified *dynamically* and on a query-dependent basis.
- Elements that do not satisfy the constraint must often be examined en-route to the solution. For example, it is not straightforward to choose between a slightly-further partition with many items in D_C and a slightly-closer partition with few valid items.

Ultimately, all constrained search approaches must find a way to integrate the searching and filtering operations efficiently. Below, we summarize the two possible approaches (as shown in Figure 1).

Search, then filter: A natural solution is to post-filter the items that are identified by the similarity search index. This approach works well when many items in the dataset satisfy the constraint, because then it is likely that only a few extra results will need to be returned from the index. However, it is problematic when D_C is sparsely populated.

Sparsity is a particularly serious issue in high dimensions, because valid candidates may be spread over an exponentially large space (in d). For graph-based algorithms, this manifests as an exploration set that grows rapidly due with the fan-out of each node [29]. For partition-based algorithms, this manifests as low recall unless a large number of clusters are probed [17].

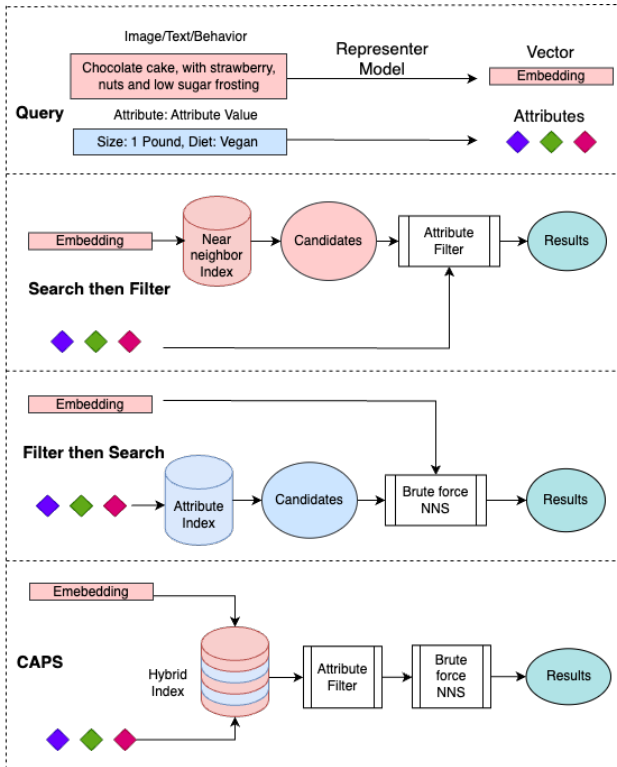


Figure 1: Search engine queries can be typed text, images, and/or user behavior, transformed into vector embeddings. Queries also involve implicitly or explicitly selected filters known as query attributes. Existing methods handle these hybrid queries through either search-then-filter (where the ANN model provides a candidate set, which gets filtered by attribute filter) or filter-then-search (where an attribute index produces a candidate set which is ranked to top-k using brute force similarity). We introduce a novel hybrid search index which provides a single step attribute filter and vector search.

Filter, then search: Another viable approach is to apply an attribute-based filter *followed by* a brute force similarity search over the items that match. This approach has the exact opposite sparsity impact: highly sparse queries are very easy, because then we need only explore a few items, but it will be prohibitively expensive if many items satisfy the constraints.

3.1 The unhappy middle

Extremely sparse constraints can be solved quickly by cascading an inverted index (or randomized approximation) with a brute-force search. Extremely dense constraints can be solved quickly by search-then-filter. Therefore, the only algorithmically interesting regime is when there are too many items to brute force D_C but too few to post-filter a standard similarity search. We call this region *the unhappy middle*. Unfortunately, most real-world queries fall into this regime. To illustrate this problem, we conducted an experiment using a partition-based search over MNIST [26] with synthetic attributes. We randomly distribute the attributes i.i.d. throughout

the whole dataset with different probabilities, then query while constraining on each attribute. We adjust the fraction of elements that satisfy the constraint (the sparsity) to investigate the effects of sparsity on search.

Figure 2 shows the effect of sparsity on the number of items examined (left) and the overall search latency (right). The results confirm our intuition: highly sparse and highly dense constraints can be adequately addressed with separate search and filtering. However, there is a central region (highlighted) where it is expensive to brute-force the entire set of valid results and equally expensive to find enough search results for post-filtering¹




Figure 2: Distance calculations (left) and latency (right) for filtered search on MNIST, subject to recall > 95%. For sparse attributes, it is optimal to filter first, then search. For dense attributes, it is sufficient to filter the search output.

We acknowledge that the differentiation between pre-filtering and post-filtering approaches has been addressed in prior discussions, including the work by Pinecone [5]. However, our contribution lies in the empirical analysis of the unhappy middle problem, which, to the best of our knowledge, has not been previously characterized. This analysis highlights the importance of developing improved algorithms in order to make further advancements in hybrid query processing.

4 RELATED WORK

Approximate solutions to the near-neighbor search problem exhibit tremendous algorithmic diversity. At various times over the last 30 years, locality-sensitive hashing [2, 37], quantization and partitioning methods [10, 12, 14, 15, 23], trees [4, 21, 32], and graph-based methods [20, 28] have each represented state of the art. Out of this, graphs and partitions have emerged as the two dominant data structures.

Graphs: Graph-based methods have traded positions with partition-based methods at the top of the leaderboard for several years [3]. Graph algorithms locate near neighbors by walking the edges of a graph where each point is (approximately) connected to its k nearest neighbors. This area has focused on improving graph properties using diversification, pruning, hierarchical structures, and other heuristics [29]. Recent work in this area has focused on systems properties of graph implementations [8], dimension reduction and theoretical analysis [31], and selective, pruned computation [6]. It

¹We note that the latency numbers (Figure 2) are somewhat artificial, as we used a simplified and unoptimized version of our algorithm using Python. Nonetheless, we believe this offers valuable insight into the nature of constrained search.

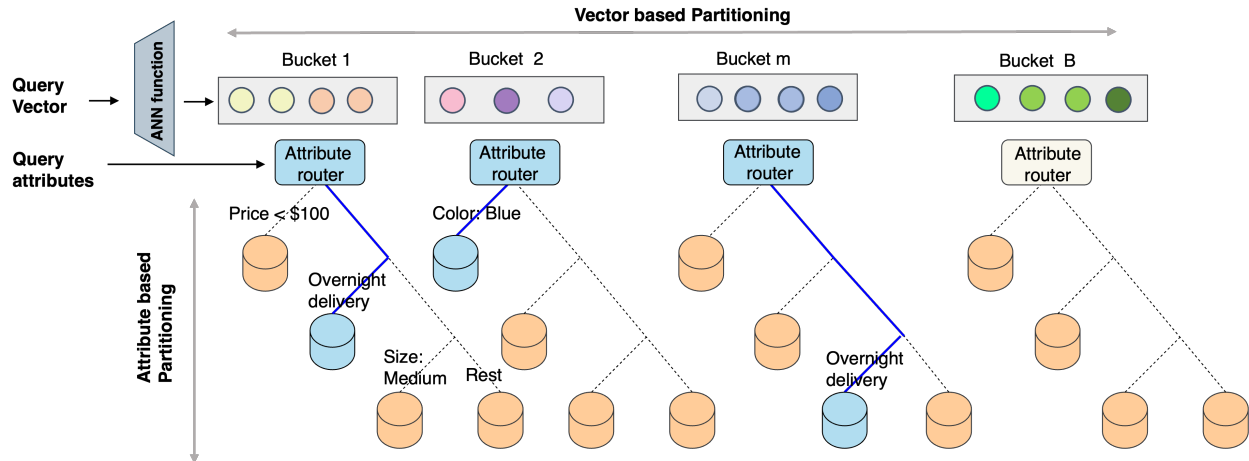


Figure 3: CAPS is adaptable to any near neighbor partitioning approach. Inspired by Huffman coding, each near neighbor partition is further sub-partitioned based on attribute occurrence frequency (refer to Section 5.1). These sub-partitions are disjoint and indexed by common attribute values. For a given query q , we select the closest m partitions to the vector part of the query and the corresponding sub-partition further on the Attribute Frequency Tree (AFT). Once the set of m sub-partitions is retrieved, CAPS performs a lexicographic search based on attributes and brute force nearest neighbor search based on vectors.

is also common for a graph index to serve as the partition identification step of a partitioned index (e.g. Neural LSH [11] and FAISS with HNSW pre-indexing).

Partitions: A large family of successful similarity search algorithms, such as FAISS-IVF [17], splits the dataset into partitions/clusters. This improves the search latency because the query does not need to examine all data points but only a few clusters that are likely to contain neighbors. The search algorithm identifies a small subset of nearby clusters, typically via a near-neighbor search on cluster identifiers. This follows the brute force near neighbor search over the points in clusters.

4.1 Hybrid query search

A variety of *hybrid* approaches have been proposed recently to integrate the constraint satisfaction into the search index.

AIRSHIP: AIRSHIP is a constrained near-neighbor algorithm based on the HNSW near-neighbor graph. Most graph algorithms employ a variation of beam search to walk through the graph and identify neighbors. Beginning from a randomly selected set of initial nodes, beam search maintains a list of “exploration candidates.” The best candidates (those closest to the query) are retained during the search as a tentative set of near-neighbors. Once the algorithm has explored a sufficiently large fraction of the graph and exhausted its supply of candidates, it returns the best candidates as the neighbors.

DiskANN: Originally introduced in 2019, DiskANN [22] is a graph-based search index capable of indexing and searching over billion-scale vector collections with only 64GB of RAM. The authors subsequently built on this work by introducing capabilities for real-time updates [35] and support for filtered queries [13]. The key insight behind the filtered DiskANN family of algorithms is to incorporate knowledge of the attributes into the graph indexing stage so that the greedy search can leverage this information.

However, this algorithm only supports the simplified case of a single constraint per query. While a single filter suffices to support disjunctive ORs (via a union of independent searches), the authors note that supporting conjunctive ANDs would require new ideas.

NHQ: Native Hybrid Query [36] is also a navigable proximity graph-based constrained near-neighbor search algorithm. It builds a graph like HNSW, using a fused distance metric given by $d_f(\cdot) = d_v(\cdot) + wd_a(\cdot)$, where the distance between any two points is given by the weighted summation of Euclidean distance between corresponding vectors and hamming distance between the binary attribute vector.

AnalyticDB-V: AnalyticDB [38] seeks to provide support for SQL-style queries over data that consists of both relational tables and embeddings of unstructured entities. The core insight behind this work is to develop four query plans and an associated cost optimizer to efficiently determine how to order the filtering and search operations. These query plans include brute force search, product quantization bitmap scans, and graph-based searching. We do not evaluate this method in our experiments due to a lack of publicly available code. Still, we note that our contributions complement AnalyticDB in that CAPS can be viewed as a potential query plan within this larger database engine.

5 CAPS: INTERLEAVED FILTERING AND SEARCH

Our proposed method allows for simultaneous filtering and search within a clustered near-neighbor index. We are looking for an algorithm that performs equally well in the dense-attribute regime (the head), the sparse regime (the tail), and the unhappy middle. To enable constrained search, we modify the clustering portion of the algorithm and augment the cluster identification process with constraint checks. The intuition for this direction is as follows.

5.1 Intuition

Consider a standard partitioned index, where we identify partitions likely to contain nearby points. Once a suitable set of partitions have been found, the neighbors which satisfy the constraints can be found via brute-force attribute matching followed by the metric distance computation on the positive attribute matches. If we examine m different partitions $\{P_1, \dots, P_m\}$, then we will require a total of $\sum_{bin=1}^m |P_{bin}|$ brute-force attribute matches and distance computations.

We can apply filtering within the partition to reduce this computational burden, but this still requires passing every point in the partition through the filter. CAPS breaks this limitation by sub-partitioning each cluster (at indexing time) based only on the attributes. By layering a similarity search index with an attribute index, we effectively perform a *filter-then-search* step after doing a preliminary *search-then-filter* over the centroids.

Our sub-partition construction algorithm is motivated by two observations. First, the sub-partitions should be non-overlapping. Second, we should refuse to search a sub-partition if none of its contents satisfy the constraint.² The first observation manifests as a constraint in the sub-partitioning algorithm: for each cluster P_{bin} , $P_{bin} = \bigcup_{j=1}^{h+1} P_{bin,j}$ and $P_{bin,j} \cap P_{bin,i} = \emptyset$ for $j \neq i$. In other words, each cluster has $h + 1$ disjoint sub-partitions.

Our second observation means one can check all sub-partitions for query filter match. Instead, we can quickly identify relevant sub-partitions in $O(1)$ time. This can be done if we cluster points with one common attribute value into the same sub-partitions so that only a few sub-partitions are likely to be valid. This is achieved by organizing the points using a frequency-based partitioning over the attributes. We explain this partitioning scheme in the following section and in Figure 3. This reduces the query complexity to $\sum_{bin=1}^m \sum_{j=1}^{h+1} C(b, p, bin, j) |P_{bin,j}|$, where $C(b, p, bin, j) = \mathbb{I}(b = A(P_{bin,j}))$, which checks whether the sub-partition contains the query attribute b .

However, it is also possible to integrate the attributes into the partitioning. This can improve performance if we have reason to suspect that the attributes correlate or cluster with the embeddings. While certain algorithms, such as NHQ, explicitly fuse the attribute and vector similarity measures into a single distance metric, this is highly dataset-dependent and is difficult to tune. This is especially problematic in production settings, where some – but not all – attributes are correlated with embeddings.

5.2 Index creation

At index time, we perform two-level partitioning of the data D . The first level is done by applying any existing partitioning method on the embedding vectors. Balanced Kmeans clustering from FAISS-IVF [17] is a common choice in partitioned indices. This partitioning does not require access to the attribute information, but instead clusters the dense vectors using a similarity measure such as the Euclidean distance, Cosine distance, or Inner-product.

Recent work on learned partitioning provides a nice avenue to integrate attributes into the first-level clustering. For example, the BLISS method [15] uses an iterative learning scheme to learn the

²Equivalently, if any point in a sub-partition is valid, we *should* search the sub-partition.

Algorithm 1 CAPS Index construction

Input: Partitioning Model $f(\cdot)$, dataset $D = (x_i, a_i)$

Learn partitioning Model $f(\cdot) = \mathcal{L}(D)$

for $i \in 1..N$ **do**

$bin = f(x_i, a_i)$, $P_{bin} \leftarrow (i)$

end for

for $bin \in 1..B$ **do**

$S = \emptyset$

for $j \in 1..h$ **do**

$t = \text{top-attribute}(P_{bin} - S)$

$p_j \leftarrow (i), t \in a_i, i \in (P_{bin} - S)$

$A(p_j) = t \text{ \#attribute tag}$

$S = S \cup p_j$

end for

$P_{h+1} \leftarrow (P_b - S)$

end for

Output: Hierarchical partitions P, p

partitions for any binary *relevant-vs-non-relevant* similarity match oracle. Therefore, we can use the true filtered near neighbors of a set of points in the dataset to learn the BLISS partitions. We consider two different setups. In CAPS-BLISS1, we do not use the attribute or filtering information, but instead train and index partitions based only on the vector near neighbors' labels. In CAPS-BLISS2, we use the filtered near neighbors labels against the data vectors.

We will use the function $f(\cdot)$ to denote the first-level partition function that assigns a partition to an input point. By applying $f(\cdot)$ to the elements of the dataset, we get the first-level partitions. The next step is to construct second-level sub-partitions that shard by attribute. We create these partitions by exploiting the power law distribution of the attributes, using a hierarchy called an , which we call an *Attribute Frequency Tree* (AFT). A given partition P is split with the points satisfying the most frequent attribute value, and the remaining set is split further with the same strategy. This creates a tree where each level has two disjoint sub-partitions, with one having the selected attribute value. In the interest of lower index memory, we truncate the tree after a fixed number of levels denoted by the height h . This truncation scheme is particularly well-suited to attributes that follow a power law distribution. We refer to Algorithm 1 and Section 6.1 for a detailed analysis of the tree depth h .

5.3 Query Algorithm

The query algorithm uses the learned partitioning function $f(\cdot)$ and selects the top m similar partitions, where m is the number of partitions to probe. For the search to be efficient, we choose $m \ll B$. The query proceeds by selecting sub-partitions within each of the m partitions, using the truncated AFT. Each level of the tree splits into a leaf and a sub-tree, where the split is performed by an attribute that is common to all items in the leaf (an *identifier*); thus we have $O(h)$ identifiers. We hash the positions of the identifiers so that the process of identifying sub-partitions in each first-level partition has $O(1)$ complexity. To get the final candidate set for brute force, we join all of the valid sub-partitions into a small candidate set that (mostly) contains valid points. In the setting where we constrain on

Algorithm 2 CAPS Query

Input: Partitioning Model $f(\cdot)$
dataset $D = (x_i, a_i)$
Input: Query point $(q, b) \in Q$
 $\mathcal{R}(q) = \phi$
Select top m partitions
 $\{bin_1, \dots, bin_m\} = f(q)$
for $bin \in \{bin_1, \dots, bin_m\}$ **do**
 IF $A(p_j) = b$ for any $j \in 1..h$
 $\mathcal{R}(q) = \mathcal{R}(q) \cup p_j$
 ELSE
 $\mathcal{R}(q) = \mathcal{R}(q) \cup p_{h+1}$
 candidate set for $q = \mathcal{R}(q)$
end for
 $R(q) = \text{Attribute Match}(\mathcal{R}(q), q)$
Result = topk Vector Match ($R(q)$)
Output: Result

multiple attributes (e.g., an OR constraint), multiple sub-partitions from each partition are included in the join.

6 EVALUATION

Baselines: We use compare CAPS with the existing SOTA algorithms for hybrid search. This includes NHQ [36], AIRSHIP [40], Filtered DiskANN’s filtered index (FI) and Stched Index (SI) [13]. DiskANN is limited to working only with one attribute per query; hence we evaluate DiskANN with only one attribute per query (even though multiple attributes are specified). Where possible, we use the original authors’ implementations.

Experiment setup: We measure the recall of the top 100 filtered neighbors (following definition 1) against the filtered search index output. The $\text{Recall}_{K@K}$ is the intersection of the top K items retrieved from the index with the true top K items. Indexing performed using 64 threads and we report single-thread query results (as is standard for kNN experiments [3]).

Attribute structure: We adopted the attribute structure as same as used in NHQ [36]. We see the similar structure of attributes on retail websites. For example, an item/query can take one value for each attribute of Color(Black, Red, Khaki, Yellow, White), Price(under \$50, \$70, \$100) and Size (S, M, L, XL).

Datasets: We use six publicly available real-world near-neighbor datasets SIFT, Glove-100, GIST, Crawl, Audio, and Msong (more details in Appendix Table 4). Because the community currently lacks any open constrained search benchmark datasets, we randomly generate attributes for each dataset with $L = 3$ following the same procedure as [36]. Each attribute can have multiple categorical values and is distributed with exponential distribution motivated by a real-world scenario in Figure 6.

Groundtruth: For the ground truth top K constrained near neighbors, we perform a (slow) exact search over the set of valid points for each query, i.e., the ground truth search results match the query attributes exactly and are the near neighbors in D_C (see Definition 1).

Metric: The metric in use is defined by

$$\text{Recall}_{K@K} = \frac{|\text{Index-top}K \cap \text{True-top}K|}{K} \quad (1)$$

Here the K is the number of generated outputs and the ground-truth true constrained near neighbors.

Hardware: We use a Ubuntu 20.04 machine equipped with 2 AMD EPYC 7742 64-core processors and 1.48TB of RAM. All experiments are performed in RAM. CAPS is written in C++ and compiled using the GNU Compiler. Refer to the anonymous GitHub link³ for the code.

Discussion: Figure 4 compares CAPS with FAISSkmeans, BLISS1 and BLISS2 with baselines on Recall100@100 against QPS. AIRSHIP suffers from the sparsity of the constraints highly for all datasets apart from SIFT. DiskANN achieves high recall but it has slower latency on a single thread. NHQ is competitive, but CAPS excels mostly in the high recall region. We also note that NHQ contains an attribute-embedding fusion hyperparameter that is challenging to tune well. Table 3 presents the overhead added by the indexing structure (e.g., graph or hierarchical partition). We achieve a very low index size with 10x smaller overhead than the next-cheapest baseline (NHQ). Please refer to the supplementary section for the dataset and code details.

6.1 Ablation Study

Attribute frequency tree height (h): We evaluate CAPS-FAISSkmeans with a varying number of subpartitions ($h + 1$), which are formed at each level of the truncated attribute frequency tree. These subpartitions are non-overlapping groups that can be easily identified using a specific attribute. Consequently, retrieving a single subpartition from each partition can be done in constant time during a query. As the size of the subpartitions becomes smaller, the computational burden associated with brute force query filtering reduces, in effect reducing the query latency (Figure 5).

Varying number of attributes (L): We also vary the number of attributes L involved with the query. In other words, any missing attribute in the query might lead to probing more than one subpartition from a partition. Figure 5 (1-2) shows an absence fraction representing the probability of an attribute value missing from the query. Our base set of experiments (Figure 4) represents absence 0. On the other side, absence 1 will denote only vector-based near-neighbor search.

Retail system applicability: To evaluate CAPS on search in retail systems, we conducted a case study with the Amazon.com search auto-completion service, which provides customers with search suggestions given a partially typed query. This service requires a low-latency near-neighbor search algorithm to find the corresponding candidate queries. In addition, the suggested query needs to meet certain search constraints, such as preserving the same product type and brand as the original query.

In this case study, we utilize the production model to generate the keyword’s embedding [24] and extract possible attributes⁴. The frequency distribution of attribute values is presented in figure 6, and the dataset dimensions are $N = 8M$ and $d = 768$. It is

³<https://anonymous.4open.science/r/constrainedANN-1DDE>

⁴We use a sample of 11 possible binary attributes for this case study

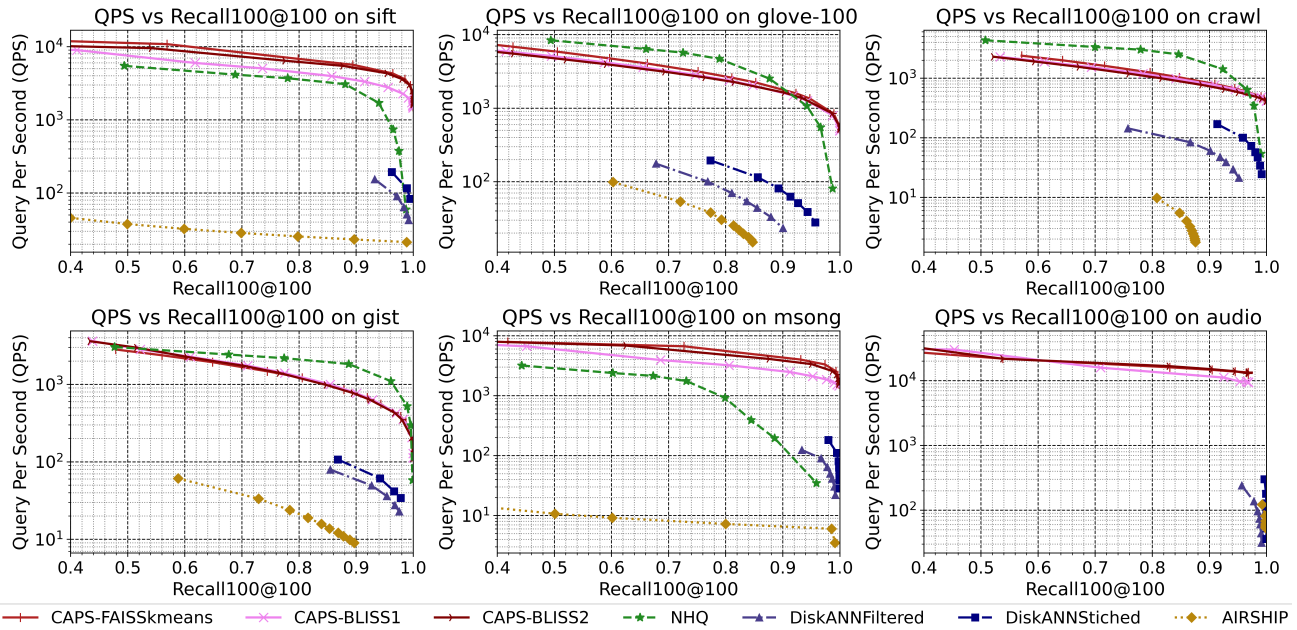


Figure 4: Recall100@100 vs Query per second tradeoff on public datasets (top-right is best).

Table 3: Indexing overhead (MB) and Construction time (Sec.) of CAPS and against the baselines.

Data	NHQ		AIRSHIP		DiskANN-FI		DiskANN-SI		CAPS-FAISSkm		CAPS-Bliss	
	Size	Time	Size	Time	Size	Time	Size	Time	Size	Time	Size	Time
SIFT	75	26	60	218	357	346.6	320.4	292.2	8.2	69.4	8.8	129
Glove	34	131.8	72	255	254	371.11	142.8	205.5	9.5	86.8	11	177
GIST	43	47.1	90.4	899	295.2	565.5	192.8	848.8	12	150.6	9.7	171
Crawl	67	168	79.4	850	591.4	738.9	284.2	611.1	18	377.7	18	415
Audio	3.3	2.83	3.38	7.1	17.38	32.2	10.38	37.8	0.62	12.6	0.86	8.6
Msong	73	78	90.4	551	397.6	373.3	192.8	360	9.3	118.3	9.1	147

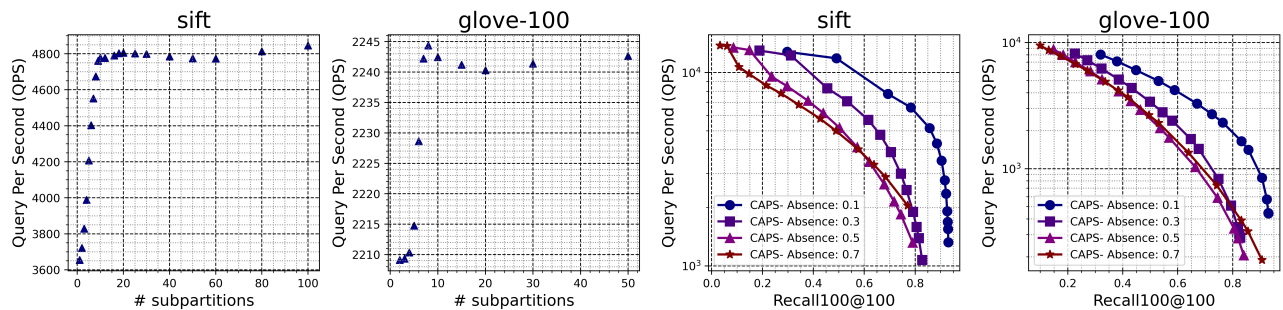



Figure 5: (1-2) The query per second (QPS) increases with an increase in the number of sub-partitions ($h+1$). For each experiment, Recall values remain unchanged. (3-4) QPS vs Recall for a variable number of query attributes. A higher absence fraction implies a fewer number of attributes per query.

evident that the Amazon attribute distribution follows a power-law distribution pattern with a number of rare constraints. We compared our method with the existing approximate near-neighbor

search algorithm in production. As shown in table of figure 6, CAPS outperforms the production algorithm in terms of query time and recall.



Methods	Query Per Sec (QPS) (↑)	R@100 (↑)
Production	1.0	1.0
CAPS-FAISSkmeans	5.56	1.2

Figure 6: Attribute values distribution for Amazon search queries.

6.2 Experiment on Larger attribute lengths

Plot 7 shows the QPS vs Recall experiment on the larger attribute length. We vary attribute length L from 1 to 10 and notice that:

- For 1 attribute, NHQ is about 13% better than CAPS
- For 2 attributes, NHQ is about the same as CAPS
- For 3 attributes, CAPS is about 40% better than NHQ
- For 10 attributes, CAPS is significantly better than NHQ (About 300%)

Additionally, we can understand the sensitivity of CAPS with L using theoretical arguments. The query time (as proved in Appendix, section 9) is function of L such that $QT \leq O(aN^{\log(L+\gamma d)} + b(L + \gamma d)N^{-\log(L+\gamma d)})$, where a and b are terms independent of L . This expression has two terms - the first term monotonically increases with L , but the second term decreases to zero. This suggests that as L increases, the query time will at first decrease (due to having a smaller set of NN search candidates) but later increase (due to a more expensive filtering process). In our set of experiments with limited values of $L=1,2,3$ and 10, we observe the same pattern. This is attributed to the fact that higher L decreases the candidate set for brute force vector distance computations but it also increases the latency of attribute distance computation.

We also see that graph based constrained NN search started performing badly with a higher number of attribute types (with increasing L we see mid-high attribute sparsity).

7 LIMITATIONS AND FUTURE WORK

We realize that, search-then-filter and filter-then-search will be optimal (compare to CAPS) with very high or very low attribute sparsity (Figure 2). Additionally, when there are only one or two attributes, then it is easier to directly incorporate the attributes into the distance function. In this case, a graph based index (NHQ) will be better than CAPS. (Figure 7)

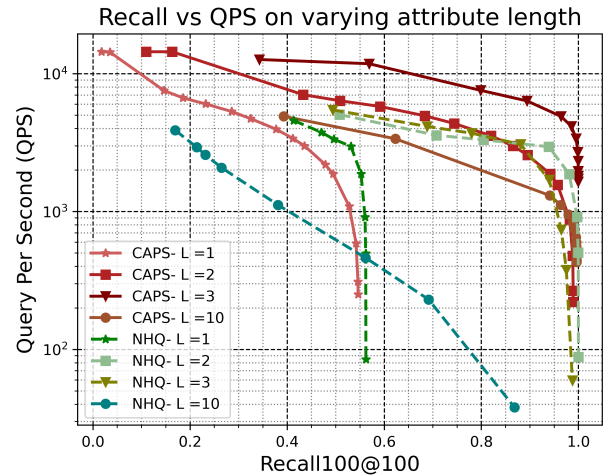


Figure 7: Attribute length L is 1, 2, 3 and 10 for SIFT 1M dataset. The Attribute length is the number of attributes present for each item at indexing time. During the query, the user selects random attributes with probability 1, 0.3, and 0.03 for $L=3,10,100$, respectively. This represents the real search scenario where the user selects only a limited number of attribute constraints.

In addition to address these limitations, we believe that the following areas are promising as a follow up research in future:

Range constraints: At present, CAPS exclusively manages categorical attributes. Range attributes, such as product prices, can be converted into categories through the process of quantizing the ranges (e.g., grouping prices into categories like \$1-\$10). When dealing with multiple ranges, the final outputs of the index can be combined. This works very well in practical use-case, but achieving more detailed handling of range constraints may require a specialized approach.

Dynamic constraints: Future works can provide an improved support for updating items with changing attributes over time.

Negation filters: Future works can provide an efficient method of finding items constrained on absence of an attribute.

8 CONCLUSION

In this work, we present CAPS, a practical algorithm for filtered approximate near-neighbor search that achieves state-of-the-art performance in the high-recall regime. While the majority of previously proposed filtered search algorithms focus on integrating with graph-based traversal, CAPS provides a novel solution via jointly partitioning the space of embeddings and attributes. This partitioning achieves a smaller space footprint than graph indexes while also supporting critical practical features such as a variable number of query attributes and conjunctive constraints. We also find that the hierarchical partitioning scheme at the core of CAPS is particularly well-suited for power law-distributed attributes, which we validate in a case study with the Amazon.com product search engine. These results demonstrate the effectiveness of space partitioning methods in the filtered search setting and provide practitioners with an immediately useful tool for vector search with constraints.

REFERENCES

- [1] stackexchange, <https://stackoverflow.com/>.
- [2] Alexandr Andoni and Piotr Indyk. Near-optimal hashing algorithms for approximate nearest neighbor in high dimensions. *Communications of the ACM*, 51(1):117–122, 2008.
- [3] Martin Aumüller, Erik Bernhardsson, and Alexander Faithfull. Ann-benchmarks: A benchmarking tool for approximate nearest neighbor algorithms. *Information Systems*, 87:101374, 2020.
- [4] Alina Beygelzimer, Sham Kakade, and John Langford. Cover trees for nearest neighbor. In *Proceedings of the 23rd international conference on Machine learning*, pages 97–104, 2006.
- [5] James Briggs. The missing where clause in vector search. <https://www.pinecone.io/learn/vector-search-filtering>, 2022.
- [6] Patrick Chen, Wei-Cheng Chang, Jyun-Yu Jiang, Hsiang-Fu Yu, Inderjit Dhillon, and Cho-Jui Hsieh. Finger: Fast inference for graph-based approximate nearest neighbor search. In *Proceedings of the ACM Web Conference 2023*, pages 3225–3235, 2023.
- [7] Benjamin Coleman, Richard Baraniuk, and Anshumali Shrivastava. Sub-linear memory sketches for near neighbor search on streaming data. In *International Conference on Machine Learning*, pages 2089–2099. PMLR, 2020.
- [8] Benjamin Coleman, Santiago Segarra, Alexander J Smola, and Anshumali Shrivastava. Graph reordering for cache-efficient near neighbor search. *Advances in Neural Information Processing Systems*, 35:38488–38500, 2022.
- [9] Mayur Datar, Nicole Immorlica, Piotr Indyk, and Vahab S Mirrokni. Locality-sensitive hashing scheme based on p-stable distributions. In *Proceedings of the twentieth annual symposium on Computational geometry*, pages 253–262, 2004.
- [10] Yihe Dong, Piotr Indyk, Ilya Razenshteyn, and Tal Wagner. Learning space partitions for nearest neighbor search. *arXiv preprint arXiv:1901.08544*, 2019.
- [11] Yihe Dong, Piotr Indyk, Ilya Razenshteyn, and Tal Wagner. Learning space partitions for nearest neighbor search. *arXiv preprint arXiv:1901.08544*, 2019.
- [12] Joshua Engels, Benjamin Coleman, and Anshumali Shrivastava. Practical near neighbor search via group testing. *Advances in Neural Information Processing Systems*, 34:9950–9962, 2021.
- [13] Siddharth Gollapudi, Neel Karia, Varun Sivashankar, Nikit Begwani, Swapnil Raz, Yiyong Lin, Yin Zhang, Neelam Mahapatro, Amit Singh, and Harsha Vardhan. Filtered-diskann: Graph algorithms for approximate nearest neighbor search with filters. In *Proceedings of the ACM Web Conference 2023*, pages 3406–3416, 2023.
- [14] Ruiqi Guo, Philip Sun, Erik Lindgren, Quan Geng, David Simcha, Felix Chern, and Sanjiv Kumar. Accelerating large-scale inference with anisotropic vector quantization. In *International Conference on Machine Learning*, pages 3887–3896. PMLR, 2020.
- [15] Gaurav Gupta, Tharun Medini, Anshumali Shrivastava, and Alexander J Smola. Bliss: A billion scale index using iterative re-partitioning. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 486–495, 2022.
- [16] Gaurav Gupta, Minghao Yan, Benjamin Coleman, Bryce Kille, RA Leo Elworth, Tharun Medini, Todd Treangen, and Anshumali Shrivastava. Fast processing and querying of 170tb of genomics data via a repeated and merged bloom filter (rambo). In *Proceedings of the 2021 International Conference on Management of Data*, pages 2226–2234, 2021.
- [17] Jui-Ting Huang, Ashish Sharma, Shuying Sun, Li Xia, David Zhang, Philip Pronin, Janani Padmanabhan, Giuseppe Ottaviano, and Linjun Yang. Embedding-based retrieval in facebook search. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2553–2561, 2020.
- [18] Po-Sen Huang, Xiaodong He, Jianfeng Gao, Li Deng, Alex Acero, and Larry Heck. Learning deep structured semantic models for web search using clickthrough data. In *Proceedings of the 22nd ACM international conference on Information & Knowledge Management*, pages 2333–2338, 2013.
- [19] Piotr Indyk and Rajeev Motwani. Approximate nearest neighbors: towards removing the curse of dimensionality. In *Proceedings of the thirtieth annual ACM symposium on Theory of computing*, pages 604–613, 1998.
- [20] Masajiro Iwasaki and Daisuke Miyazaki. Optimization of indexing based on k-nearest neighbor graph for proximity search in high-dimensional data. *arXiv preprint arXiv:1810.07355*, 2018.
- [21] Mike Izbicki and Christian Shelton. Faster cover trees. In *International Conference on Machine Learning*, pages 1162–1170. PMLR, 2015.
- [22] Suhas Jayaram Subramanya, Fnu Devvrit, Harsha Vardhan Simhadri, Ravishankar Krishnawamy, and Rohan Kadekodi. Diskann: Fast accurate billion-point nearest neighbor search on a single node. *Advances in Neural Information Processing Systems*, 32, 2019.
- [23] Herve Jegou, Matthijs Douze, and Cordelia Schmid. Product quantization for nearest neighbor search. *IEEE transactions on pattern analysis and machine intelligence*, 33(1):117–128, 2010.
- [24] Haoming Jiang, Tianyu Cao, Zheng Li, Chen Luo, Xianfeng Tang, Qingyu Yin, Danqing Zhang, Rahul Goutam, and Bing Yin. Short text pre-training with extended token classification for e-commerce query understanding. *arXiv preprint arXiv:2210.03915*, 2022.
- [25] Vihan Lakshman, Choon Hui Teo, Xiaowen Chu, Priyanka Nigam, Abhinandan Patni, Pooja Maknikar, and SVN Vishwanathan. Embracing structure in data for billion-scale semantic product search. *arXiv preprint arXiv:2110.06125*, 2021.
- [26] Yann LeCun. The mnist database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>, 1998.
- [27] Sen Li, Fuyu Lv, Taiwei Jin, Guli Lin, Keping Yang, Xiaoyi Zeng, Xiao-Ming Wu, and Qianli Ma. Embedding-based product retrieval in taobao search. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pages 3181–3189, 2021.
- [28] Yu A Malkov and Dmitry A Yashunin. Efficient and robust approximate nearest neighbor search using hierarchical navigable small world graphs. *IEEE transactions on pattern analysis and machine intelligence*, 42(4):824–836, 2018.
- [29] Yu A Malkov and Dmitry A Yashunin. Efficient and robust approximate nearest neighbor search using hierarchical navigable small world graphs. *IEEE transactions on pattern analysis and machine intelligence*, 42(4):824–836, 2018.
- [30] Priyanka Nigam, Yiwei Song, Vijai Mohan, Vihan Lakshman, Weitian Ding, Ankit Shingavi, Choon Hui Teo, Hao Gu, and Bing Yin. Semantic product search. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2876–2885, 2019.
- [31] Liudmila Prokhorenkova and Aleksandr Shekhovtsov. Graph-based nearest neighbor search: From practice to theory. In *International Conference on Machine Learning*, pages 7803–7813. PMLR, 2020.
- [32] Parikshit Ram and Kaushik Sinha. Revisiting kd-tree for nearest neighbor search. In *Proceedings of the 25th acm sigkdd international conference on knowledge discovery & data mining*, pages 1378–1388, 2019.
- [33] Rice. sentence-transformers. Rice News, <https://huggingface.co/sentence-transformers/multi-qa-MiniLM-L6-cos-v1>, 2020.
- [34] Aneesh Sharma, C Seshadhri, and Ashish Goel. When hashes met wedges: A distributed algorithm for finding high similarity vectors. In *Proceedings of the 26th International Conference on World Wide Web*, pages 431–440, 2017.
- [35] Aditi Singh, Suhas Jayaram Subramanya, Ravishankar Krishnawamy, and Harsha Vardhan Simhadri. Freshdiskann: A fast and accurate graph-based ann index for streaming similarity search. *arXiv preprint arXiv:2105.09613*, 2021.
- [36] Mengzhao Wang, Lingwei Lv, Xiaoliang Xu, Yuxiang Wang, Qiang Yue, and Jiongkang Ni. Navigable proximity graph-driven native hybrid queries with structured and unstructured constraints. *arXiv preprint arXiv:2203.13601*, 2022.
- [37] Yiqiu Wang, Anshumali Shrivastava, Jonathan Wang, and Junghee Ryu. Flash: Randomized algorithms accelerated over cpu-gpu for ultra-high dimensional similarity search. *arXiv preprint arXiv:1709.01190*, 2017.
- [38] Chuangxian Wei, Bin Wu, Sheng Wang, Renjie Lou, Chaoqun Zhan, Feifei Li, and Yuanzhe Cai. Analyticdb-v: a hybrid analytical engine towards query fusion for structured and unstructured data. *Proceedings of the VLDB Endowment*, 13(12):3152–3165, 2020.
- [39] Michael Wray, Hazel Doughty, and Dima Damen. On semantic similarity in video retrieval. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3650–3660, 2021.
- [40] Weijie Zhao, Shulong Tan, and Ping Li. Constrained approximate similarity search on proximity graph. *arXiv preprint arXiv:2210.14958*, 2022.

9 APPENDIX

9.1 Query time complexity

CAPS creates the index with a partition identification function $f(\cdot)$, sub-partition lookup, and index files. The query time is the summation of top m partition identification time + sub-partition identification time + brute force filter from the given sub-partition + brute force distance computation on the constrained satisfied items from the sub-partition. For the k-means this is

$$QT = Bd + B \log m + mO(1) \quad (2)$$

$$+ \sum_{i=1}^m \sum_{j=1}^{h+1} (\mathbb{I}(A(p) = b) |p_{i,j}|) (L + d \mathbb{I}(\text{Attributematch})) \quad (3)$$

The Attribute match is the constrained match on query attribute b and the item attributes of $p_{i,j}$. If the attributes are correlated with the embedding feature, we will observe the $\mathbb{I}(\text{Attributematch})$ be 1 for most of the items in the top partitions match. However, in the worst case scenario, we can get $\mathbb{I}(\text{Attributematch})$ to be uniformly distributed over all the item embeddings. Assuming this uniform distribution probability a.k.a sparsity to be γ , we can write

$$QT \leq Bd + B \log m + mO(1) + \sum_{i=1}^m \sum_{j=1}^{h+1} (\mathbb{I}(A(p) = b) |p_{i,j}|) (L + d\gamma)$$

If the query specifies all the attributes, only one sub-partition is selected from each partition. Hence assuming a balanced sub-partitioning, we have

$$QT \leq Bd + B \log m + mO(1) + \sum_{i=1}^m \frac{|P_i|}{h+1} (L + d\gamma)$$

Under an assumption that the number of partitions is polynomial in N , $|B| = N^t$, where $t < 1$, the sublinear query time is achieved in

$$QT \leq N^t d + N^t \log m + m + m \frac{N^{1-t}}{h+1} (L + d\gamma) = QT_{UB} \quad (4)$$

The minimum of the query time upper bound is achieved when $d(QT_{UB})/dt = 0$

$$t = \frac{\log N \theta}{2 \log N}, \quad \text{Where } \theta = \frac{(L + d\gamma)m}{(h+1)(d + \log m)}$$

With the given expression of t , we have

$$QT_{UB} = N^t (d + \log m) + m + \frac{mN}{h+1} N^{-t} (L + d\gamma)$$

Focusing only on its dependence on L , we can write

$$QT_{UB} \approx O(aN^{\log(L+\gamma d)} + b(L + \gamma d)N^{-\log(L+\gamma d)}) \quad (5)$$

where a and b are terms independent of L

9.2 Index size

The number of bytes used in RAM from index and data is embedding model size ($4Bd$) + sub-partition lookup ($4N + 4B(h+1)$) + sub-partition key ($2B(h+1)r$) + data (including attributes + vectors) ($4Nd + NLr$). We use *float32* precision for all apart from storing attributes. The total size is:

$$\text{Size}(\text{Index} + \text{data}) = 4Bd + 4N + 4B(h+1) + 2B(h+1)r + 4Nd + NLr$$




Figure 8: Comparison on Recall100@100 against QPS for the Serverfault questions search with tags as attributes. The CAPS is used with FaissKmeans clustering based partitioning of the data. The possible number of attributes are 100.

where $r = \log L \bmod 8$, is the attribute precision. For example, for a total of 180 attribute values, we only need $r = 1$, i.e. 1 byte for representation. After simplification, the total size of index and data needed during the query operation is:

$$\text{Size}(\text{Index} + \text{data}) = B(4d + 2(h+1)(2+r)) + N(4d + 1 + rL) \quad (6)$$

9.3 Experiment with real world attribute data

Besides examining near-neighbor data and real retail system data, we also analyze CAPS on real-world attribute and vector dataset. To accomplish this, we sample Serverfault questions from Stack Exchange [1]. Additionally, we record the corresponding tags. These tags serve as attributes for each question text. The dataset comprises 832,692 questions, each accompanied by its respective attributes. As an illustration, a sampled question and its corresponding tag looks as follows:

Our nightly full (and periodic differential) backups are becoming quite large. due mostly to the amount of indexes on our tables; roughly half the backup size is comprised of indexes. We're using the Simple recovery model for our backups. Is there any way, through using FileGroups or some other file-partitioning method, to exclude indexes from the backups? It would be nice if this could be extended to full-text catalogs, as well.

Tags: sql-server, backup, sql-server-2008, indexes

For consistency with NHQ's data format, we restrict the total unique tags to 100 and allocate an attribute membership vector of size 100 to each text row. Additionally, we represent the question text using the sentence-transformers' "multi-qa-mpnet-base-dot-v1" [33] embedding of size 768. The CAPS index was constructed with 256 k-means clusters and an attribute tree height of $h = 5$.

Table 4: Datasets

Name	Embedding dimension	Corpus size	Query size	Source
SIFT	128	1000000	10000	http://corpus-texmex.irisa.fr/
Glove-100	100	1183514	10000	https://nlp.stanford.edu/projects/glove/
GIST	960	1000000	1000	https://www.cs.cmu.edu/enron/
Crawl	300	1989995	10000	https://commoncrawl.org/
Audio	192	53387	200	https://www.cs.princeton.edu/cass/demos.htm
Msong	420	992272	200	http://www.ifs.tuwien.ac.at/mir/msd/
Serverfault	768	832692	10000	https://serverfault.com/questions/

From Figure 8, we infer that CAPS shows a better Recall100@100 vs QPS trade-off than the strongest competitor NHQ.

9.4 Dataset details

We use six publicly available near-neighbor datasets SIFT, Glove-100, GIST, Crawl, Audio, and Msong (Table 4) and Amazon internal dataset. Additionally, we sample 832,692 questions from stack exchange for real attribute dataset comparison.