



mmLSH: A Practical and Efficient Technique for Processing Approximate Nearest Neighbor Queries on Multimedia Data

Omid Jafari^(✉), Parth Nagarkar, and Jonathan Montaña

New Mexico State University, Las Cruces, USA
{ojafari,nagarkar,jmon}@nmsu.edu

Abstract. Many large multimedia applications require efficient processing of nearest neighbor queries. Often, multimedia data are represented as a collection of important high-dimensional feature vectors. Existing Locality Sensitive Hashing (LSH) techniques require users to find top-k similar feature vectors for each of the feature vectors that represent the query object. This leads to wasted and redundant work due to two main reasons: 1) not all feature vectors may contribute equally in finding the top-k similar multimedia objects, and 2) feature vectors are treated independently during query processing. Additionally, there is no theoretical guarantee on the returned multimedia results. In this work, we propose a practical and efficient indexing approach for finding top-k approximate nearest neighbors for multimedia data using LSH called *mmLSH*, which can provide theoretical guarantees on the returned multimedia results. Additionally, we present a buffer-conscious strategy to speed up the query processing. Experimental evaluation shows significant gains in performance time and accuracy for different real multimedia datasets when compared against state-of-the-art LSH techniques.

Keywords: Approximate nearest neighbor search · High-dimensional spaces · Locality Sensitive Hashing · Multimedia indexing

1 Introduction

Finding nearest neighbors in high-dimensional spaces is an important problem in several multimedia applications. In multimedia applications, content-based data objects, such as images, audio, videos, etc., are represented using high-dimensional feature vectors. Locality Sensitive Hashing (LSH) [8] is one of the most popular solutions for the approximate nearest neighbor (ANN) problem in high-dimensional spaces. Since it was first introduced in [8], many variants of LSH have been proposed [4, 7, 9, 13] that mainly focused on improving the search accuracy and/or the search performance of the given queries. LSH is known for two main advantages: its sub-linear query performance (in terms of

the data size) and theoretical guarantees on the query accuracy. While the original LSH index structure suffered from large index sizes (in order to obtain a high query accuracy), state-of-the-art LSH techniques [7, 9] have alleviated this issue by using advanced methods such as *Collision Counting* and *Virtual Rehashing*. Thus, owing to their small index sizes, fast index maintenance, fast query performance, and theoretical guarantees on the query accuracy, we propose to build *mmLSH* upon existing state-of-the-art LSH techniques.

Motivation of Our Work: Drawbacks of LSH on Multimedia Data

Popular feature extraction algorithms, such as SIFT, SURF (for images), Marsyas (for audio), etc., extract multiple features that collectively represent the object of interest for improved accuracy during retrieval. Hence, if a user wants to find similar objects to a given query object, nearest-neighbor queries have to be performed for every individual feature vector representing the query object (and then these intermediate results are aggregated to find the final object results (Sect. 4)). Existing techniques treat these individual feature vectors as independent of each other, and hence cannot leverage common elements between these feature vector queries for improved query performance. Most importantly, existing techniques can only give theoretical guarantees on the accuracy of the individual feature vector queries, but not on the final object results, unlike our proposed index structure, *mmLSH*.

Contributions of this Paper: In this paper, we propose a practical and efficient indexing approach for finding top-k approximate nearest neighbors for multimedia data using LSH, called *mmLSH*. To the best of our knowledge, we are the first work to provide a rigorous theoretical analysis for answering approximate nearest neighbor queries on high-dimensional multimedia data using LSH. Our main contributions are:

- *mmLSH* can efficiently solve approximate nearest neighbor queries for multimedia data while providing rigorous theoretical analysis and guarantees on the accuracy of the query result.
- Additionally, we present an advanced buffer-conscious strategy to speedup the processing of a multimedia query.
- Lastly, we experimentally evaluate *mmLSH*, on diverse real multimedia datasets and show that *mmLSH* can outperform the state-of-the-art solutions in terms of performance efficiency and query accuracy.

2 Related Work

LSH was originally proposed in [8] for the Hamming distance and then later extended to the popular Euclidean distance [6]. C2LSH [7] introduced two main concepts of *Collision Counting* and *Virtual Rehashing* that solved the two main drawbacks of E2LSH [6]. QALSH [9] used these two concepts to build query-aware hash functions such that the hash value of the query object is considered as the anchor bucket during query processing. [19] proposes an efficient distributed LSH implementation which includes a cache-conscious hash table generation

(to avoid cache misses to improve the index construction time). Our proposed cache-conscious optimization is to improve the efficiency of the query processing (and hence very different).

Query Workloads in High-Dimensional Spaces: Until now, only two works [10, 15] have been proposed that focus on efficient execution of query workloads in high-dimensional spaces. Neither of these two works provide a rigorous theoretical guarantees on the accuracy of the final result. In [15], the authors propose to efficiently execute set queries using a two-level index structure. The problem formulation, which is quite restrictive compared to our work, states that a point will be considered in the result set only if it satisfies a certain user-defined percentage of the queries in the query workload. In [10], the authors build a model based on the cardinality and dimensionality of the high-dimensional data to efficiently utilize the cache. The main drawback of these two approaches is that they require prior information that is found by analyzing past datasets. Hence the accuracy and efficiency of the index structures is determined by the accuracy of the models. Our proposed work is very different from these previous works: *mmLSH* does not require any training models and additionally, we provide a theoretical guarantee on the accuracy of the returned results.

3 Key Concepts and Problem Specification

A hash function family H is (R, cR, p_1, p_2) -sensitive if it satisfies the following conditions for any two points x and y in a d -dimensional dataset $D \subset \mathbb{R}^d$: if $|x - y| \leq R$, then $\Pr[h(x) = h(y)] \geq p_1$, and if $|x - y| > cR$, then $\Pr[h(x) = h(y)] \leq p_2$. Here, p_1 and p_2 are probabilities and c is an approximation ratio. LSH requires $c > 1$ and $p_1 > p_2$. In the original LSH scheme for Euclidean distance, each hash function is defined as $h_{a,b}(x) = \lfloor \frac{a \cdot x + b}{w} \rfloor$, where a is a d -dimensional random vector and b is a real number chosen uniformly from $[0, w)$, such that w is the width of the hash bucket [6]. C2LSH [7] showed that two close points x and y collide in at least l hash layers (out of m) with a probability $1 - \delta$.

Given a multidimensional database \mathcal{D} , \mathcal{D} consists of n d -dimensional points that belongs to \mathbb{R}^d . Each d -dimensional point x_i is associated with an object X_j s.t. multiple points are *associated* with a single object. There are S objects in the database ($1 \leq S \leq n$), and for each object X_j , $\text{set}(X_j)$ denotes the set of points that are associated with X_j . Thus, $n = \sum_{j=1}^S |\text{set}(X_j)|$.

Our goal is to provide a k -NN version of the c -approximate nearest neighbor problem for multidimensional objects. For this, we propose a notion of distance between multidimensional objects called Γ -distance (defined in Sect. 4.1) and that depends on a percentage parameter that we denote by Γ .

Let us denote the Γ -distance between two objects X_1 and X_2 by $\Gamma\text{dist}(X_1, X_2)$. For a given query object Q , an object X_j is a Γ - c -approximate nearest neighbor of Q if the Γ -distance between Q and X_j is at most c times the Γ -distance between Q and its true (or exact) nearest neighbor, X_j^* , i.e. $\Gamma\text{dist}(Q, X_j) \leq c \times \Gamma\text{dist}(Q, X_j^*)$, where $c > 1$ is an *approximation ratio*.

Similarly, the Γk -NN version of this problem states that we want to find k objects that are respectively the Γ - c -approximate nearest neighbors of the exact k -NN objects of Q .

4 mmLSH

The Borda Count method [17] (along with other aggregation techniques [16]) are popular existing techniques to aggregate results of multiple point queries to find similar objects in multimedia retrieval [2]. In order to find top- k nearest neighbor *objects* of multimedia object query Q , the existing methods find the top- k' nearest neighbor *points* for each query *point* q_i , where $1 \leq i \leq |set(Q)|$, k is the number of desired results by the user, and k' is an arbitrarily chosen number such that $k' \gg k$ [2]. Once the top- k' nearest neighbors of each query point q_i is found, an overall score is assigned to each multimedia object X_j based on the depth of the points (associated with X_j) in the top- k' results for each of the point queries q_i of Q . **Drawbacks of this approach:** 1) there is no theoretical guarantee for the accuracy of the returned top- k result objects, and 2) all query points q_i of the query object Q are executed independently of each other. Hence, if a query point takes too long to execute as compared to others, then the overall processing time is negatively affected. Our proposed method, *mmLSH*, solves both these drawbacks as explained in the next sections.

4.1 Key Definitions of mmLSH

Justification for Using R -Object Similarity and Γ -distance: In order to define two *Nearby Objects*, we first define a similarity/distance measure between two objects in the context of ANN search. Note that, there have been several works that have defined voting-based similarity/distance measures between two multimedia objects, especially images [11, 22, 23]. Also, region-based algorithms have been explored in the past, whose main strategy is to divide the query object in regions and compare these with regions of the dataset objects via some region distance, and then aggregate the resulting distances [1]. In this work we define the Γ -distance as a way to measure distances between objects as a whole. Our definition follows the naive strategy of comparing all pairs of features of the objects but it uses a percentage parameter Γ to ensure two identical objects have a near zero distance. Another key advantage of the proposed distance is that it allows us to provide theoretical guarantees for our results. In order to do so, we leverage the theoretical guarantees of LSH in our design.

Definition 1 (R -Object Similarity). *Given a radius R , the R -Object Similarity between two objects Q and X_j , that consists of $set(Q)$ and $set(X_j)$ d -dimensional feature vectors respectively, is defined as:*

$$sim(Q, X_j, R) = \frac{|\{q \in set(Q), x_i \in set(X_j) : \|q - x_i\| \leq R\}|}{|set(Q)| \cdot |set(X_j)|} \quad (1)$$

Note that, $0 \leq \text{sim}(Q, X_j, R) \leq 1$. $\text{sim}(Q, X_j, R)$ will be equal to 1 if every point of Q is a distance at most R to every point of X_j (e.g. if you are comparing an entirely green image with another green image - and assuming the feature vectors were based on the color of the pixel. But if you are comparing two identical images, then $\text{sim}(Q, X_j, R) < 1$ if R is less than the largest among $\|q - x_i\|$). Since the number of points associated with two objects can be different, we normalize the similarity w.r.t the points associated with Q and X_j .

Definition 2 (Γ -distance). *Given a two objects Q and X_j , the Γ -distance between Q and X_j is defined as:*

$$\Gamma\text{dist}(Q, X_j) = \inf\{R \mid \text{sim}(Q, X_j, R) \geq \Gamma\} \quad (2)$$

In order to find points that are within R distance, we use the *Collision Counting* method that is introduced in C2LSH [7].

We define a *Collision Index* (denoted by $ci(Q, X_j)$) that determines how close two objects are based on the number of points between the two objects that are considered *close* (i.e. the collision counts between the points of the two objects is greater than the collision threshold l).

Definition 3 (Collision Index of Two Objects). *Given two objects Q and X_j , the collision index of X_j with respect to Q is defined as:*

$$ci(Q, X_j) = \frac{|\{q \in \text{set}(Q), x_i \in \text{set}(X_j) : cc(q, x_i) \geq l\}|}{|\text{set}(Q)| \cdot |\text{set}(X_j)|} \quad (3)$$

The *Collision Index* between two objects depends on how many nearby points are considered as candidates between the two objects. Thus, in turn, the accuracy of the collision index depends on the accuracy of the collision counting process (which is shown to be very high [7, 9]). Hence we define an object X_j to be a Γ -candidate if the collision index between them is greater than or equal to $(1 - \varepsilon)\Gamma$, where $\varepsilon > \delta$ is an approximation factor which we set to 2δ .

Definition 4 (Γ -candidate Objects). *Given an object query Q and an object X_j , we say that X_j is a Γ -candidate with respect to Q if $ci(Q, X_j) \geq (1 - \varepsilon)\Gamma$.*

Additionally, we define an object to be a Γ -false positive if it is a Γ -candidate but its Γ -distance to the object query is too high.

Definition 5 (Γ -False Positives). *Given an object query Q and an object X_j , we say X_j is a Γ -false positive with respect to Q if we have $ci(Q, X_j) \geq \Gamma + \frac{\beta}{2}$ but $\Gamma\text{dist}(Q, X_j) > cR$.*

4.2 Design of mmLSH

During query processing, instead of executing the query points of Q independently, we execute them one at a time in each projection (Lines 5–6 in Algorithm 1). The function *CountCollisions*(q_i) (Line 7), an existing function from

Algorithm 1. k-Nearest Neighbor Object

```

1: while TRUE do
2:   if  $|\{X_j | X_j \in \mathcal{CL} \wedge \Gamma \text{dist}(Q, X_j) \leq cR\}| \geq k$  then
3:     return the top- $k$  objects from  $\mathcal{CL}$ ;
4:   end if
5:   for  $g = 1; g \leq m; g++$  do
6:     for  $i = 1; i \leq |\text{set}(Q)|; i++$  do
7:        $\text{CountCollisions}(q_i)$ ;
8:        $\forall_{X_j \in S} \text{Update } ci(Q, X_j)$ ;
9:     end for
10:    if  $|\mathcal{CL}| \geq k + \beta S$  then
11:      return the top- $k$  objects from  $\mathcal{CL}$ ;
12:    end if
13:  end for
14:   $R = c^{\text{numIter}}$ ;
15:   $\text{numIter}++$ ;
16: end while

```

C2LSH, is responsible for counting collisions of query points and points in the database. The *Buffer-conscious Optimizer* module (Sect. 4.3) is responsible for finding an effective strategy to utilize the buffer to speed up the query processing. This module decides which query and the hash bucket should be processed next. The Γ -Analyzer module is in charge of calculating the *collision indexes* (Sect. 4.1) for objects in the database and for checking/terminating the process if the terminating conditions are met.

Terminating Conditions for mmLSH: The existing solution (Sect. 4) finds top- k' candidates for each query point in Q and then terminates. Instead, *mmLSH* stops when top- k objects are found. These conditions guarantee that Γ - c^2 -approximate NN are found with constant probability (Sect. 4.4):

T1) At certain point at level- R , at least $k + \beta S$ Γ -candidates have been found, where βS is the allowed number of false positives. (Line 14, Algorithm 1)

T2) At the end of level- R , there exists at least k Γ -candidates whose Γ -distance to Q is at most R . (Line 6, Algorithm 1)

4.3 Buffer-Conscious Optimization for Faster Query Processing

Another goal of *mmLSH* is to improve the processing speed by efficiently utilizing a given buffer space. In order to explain our strategy, we first analyze the two expensive operations and the two naive strategies for solving the problem. The two main dominant costs in LSH-based techniques are the *Algorithm time* (which is the time required to find the candidate points that collide with the given query point) and the *Index IO time* (which is the time needed to bring the necessary index files from the secondary storage to the buffer).

Due to space limitations, we do not present a formal cost model for this process. Our main focus is on minimizing the above mentioned two dominant

costs: *algorithm time* and *index IO time*. We want to store the most important hash buckets from the cache to maximize total number of buffer hits.

Table 1. Performance comparison of naive strategies NS1 and NS2 (in sec)

	Total	AlgTime	IndexIOTime
NS1: LRU	263.6	117.6	146.0
NS2: Per-Bucket	279.5	260.8	18.7




Fig. 1. (a) Query split strategy, (b) Comparison of exact and approx. Frequencies on a real and random query

Naive Strategy 1: Using LRU Eviction Strategy on a Given Buffer.

Given Q , we first find the hash bucket locations for each of point queries of Q . In order to make the LRU (Least Recently Used) eviction strategy more effective, in each hash function m , we order the execution of point queries of Q according to the hash bucket locations from left to right. During query processing, we evict the LRU index files from the buffer when the buffer gets full.

Naive Strategy 2: Using a Per-bucket Execution Strategy. Since one of our goals is to reduce the *indexIOCost*, we also consider a Per-bucket execution strategy. Given a query object Q , we bring each useful hash bucket, hb , into the buffer, and for every q in Q that requires hb , we perform *Collision Counting* to find the candidate nearest neighbor points to the point query. In Fig. 1(a), this strategy would bring in hb_1 (then solve for q_1), then bring hb_2 (and then solve for q_1 and q_2 , since both queries using hb_2) and so on.

As seen from Table 1, NS1, due to its simplicity, has a lot smaller *AlgTime* than NS2, but its *IndexIOTime* is a lot more than NS2. NS2 needs to find the queries that require the particular hash bucket brought into the main memory. While this process can be sped up with more index structures, it is still an expensive operation to check for all queries. In each projection, since a hash bucket is brought into the buffer only once for NS2, *IndexIOTime* is the lowest.

Hence we propose an efficient and effective buffer-conscious strategy that reduces the *IndexIOTime* of NS1 without adding significant overhead to the *AlgTime*. Instead of using LRU, our eviction strategy is to evict a bucket based

on the following three intuitive criteria. **Criterion 1:** if the bucket was not added to the buffer *very recently*. When a bucket is added to the buffer, then there is a high likelihood that another query might use it in the near future. **Criterion 2:** if the bucket is *far away* from the current query. It is more beneficial to evict another bucket that is far apart in the projection than the position of the current query. **Criterion 3:** if the number of queries that still require this bucket (called *frequency* of the bucket) is the lowest *after* the first two criteria are satisfied. Criterion 3 ensures that a bucket needed by a lot of queries is not evicted.

The main challenge is that the main criterion (Criterion 3) requires *mmLSH* to know the frequencies of each bucket to decide which bucket to evict. This is an unfair expensive requirement to have in online query processing. Across different multimedia datasets, we *observed* that the frequencies of buckets on collection of queries associated with an object showed a behavior very similar to a collection of randomly chosen queries. Figure 1(b) shows that the bucket frequencies for a randomly chosen query from the Wang [21] dataset exhibit a similar pattern for a set of randomly generated point queries on a single projection.

Projection-Dividing Strategy: We use the above stated important observation to estimate the frequencies of buckets during offline processing. The following is the overview: 1) *We divide a projection into different regions*. Too few divisions will result in a high error between the estimated and actual frequencies. Too many divisions will also result in a high error because if the frequency behavior is slightly deviated than the random queries' behavior, then we assign same frequencies as that of the random queries. For this paper, we empirically decide the total number of divisions (set to 10). 2) *We calculate the average frequencies for the random point queries for each region*, and assign the region's frequency to each bucket in that particular region. 3) For each projection, *we assign approximate frequencies to all buckets in each projection*.

Query-Splitting Strategy: In order to utilize the buffer more effectively, we split the queries into multiple sub-queries and reorder the execution of the queries based on these new set of queries. In Fig. 1(a), the query execution order will change from q_1, q_2, q_3 to $q_{1a}, q_{2a}, q_{1b}, q_{3a}, q_{2b}, q_{3b}$ to utilize the buffer more effectively. Note that too many splits is still detrimental due to the increase in the overall Algorithm time (like Naive Strategy 2). We also try different number of splits and get the overall times of 173, 169, 171, and 176 s for 5, 10, 15, and 20 splits respectively. In this work, we empirically find a good split (that is found during the indexing phase, and set to 10).

4.4 Theoretical Analysis

Guarantees on the Stopping Conditions. The goal of this section is to prove the following theorem which provides a theoretical guarantee to *mmLSH*. For simplicity we perform the theoretical analysis for the case $k = 1$, the general case follows similarly after simple adaptations.

Theorem 1. *Let Q be a query object and let $L = \min\{|X| \mid X \in \mathcal{D}\}$. If*

$$\Gamma \geq \sqrt{\max \left\{ \frac{\ln \frac{1}{\delta}}{(\varepsilon - \delta)^2 |Q|L}, \frac{2 \ln \frac{2}{\beta}}{\beta^2 |Q|L} \right\}},$$

then mmLSH finds a Γ - c^2 -approximate NN with constant high probability.

For the proof of this theorem we need the following lemma. For this, we consider the following two properties for a given query object Q and level R :

- $\mathcal{P}1$) If X is an object such that $\Gamma \text{dist}(Q, X) \leq R$ then X is a Γ -candidate.
- $\mathcal{P}2$) The number of Γ -false positives is at most βS .

In the next lemma, we show that the above properties hold with high probability.

Lemma 1. *Let δ be the probability defined in Sect. 3 and $\varepsilon > \delta$ as defined in Sect. 4.1, then if Γ satisfies the inequality in Theorem 1 we have $\Pr[\mathcal{P}1] \geq 1 - \delta$ and $\Pr[\mathcal{P}2] > \frac{1}{2}$.*

Proof. For $q_i \in \text{set}(Q)$ and $x_j \in \text{set}(X)$, let A be the condition $cc(q, x_j) \geq l$, B be $\|q - x_j\| \leq R$, and C be $\|q - x_j\| > cR$. From the proof of Lemma 1 in [7] we know the following inequalities hold:

$$\Pr[A|B] \geq 1 - \delta \quad \text{and} \quad \Pr[\neg A|C] \geq (1 - \exp(-2(\alpha - p_2)^2 m)) \geq (1 - \frac{\beta}{2}). \quad (4)$$

We proceed to prove inequality $\Pr[\mathcal{P}1] \geq 1 - \delta$. Assume $\Gamma \text{dist}(Q, X) \leq R$, which is equivalent to $\Pr[\|q_i - x_j\| \leq R] \geq \Gamma$, where $q_i \in \text{set}(Q)$ and $x_j \in \text{set}(X)$. Therefore, $p = \Pr[A] \geq \Pr[A \wedge B] = \Pr[A|B]\Pr[B] \geq (1 - \delta)\Gamma$, where the last inequality follows from the left hand side inequality in Eq. (4).

For every $1 \leq i \leq |Q|$ and $1 \leq j \leq |X|$, let $Y_{i,j} \sim \text{Ber}(1 - p)$ be a Bernoulli random variable which is equal to 1 if $cc(q_i, x_j) < l$. Then

$$\begin{aligned} \Pr[ci(Q, X) \geq (1 - \varepsilon)\Gamma] &= 1 - \Pr\left[\sum_{i,j} Y_{i,j} \geq (1 - (1 - \varepsilon)\Gamma)|Q||X|\right] \\ &\geq 1 - \exp(-2(\varepsilon - \delta)^2 \Gamma^2)|Q||X|, \end{aligned}$$

where the inequality follows from Hoeffding's Inequality. Therefore for the given range of Γ we have $\Pr[\mathcal{P}1] = \Pr[ci(Q, X) \geq (1 - \varepsilon)\Gamma] \geq 1 - \delta$.

We continue with the proof of $\Pr[\mathcal{P}2] > \frac{1}{2}$. For this, we assume $\Gamma \text{dist}(Q, X) > cR$. Which is equivalent to $\Pr[\|q_i - x_j\| > cR] \geq 1 - \Gamma$. Then

$$1 - p = \Pr[\neg A] \geq \Pr[\neg A \wedge C] = \Pr[\neg A|C]\Pr[C] \geq (1 - \frac{\beta}{2})(1 - \Gamma)$$

where the last inequality follows from the right hand side inequality in Eq. (4). Therefore, $p \leq \Gamma + \frac{\beta}{2} - \frac{\beta\Gamma}{2}$. For every $1 \leq i \leq |Q|$ and $1 \leq j \leq |X|$, let $Y_{i,j} \sim \text{Ber}(1 - p)$ be a Bernoulli random variable defined as above. Thus

$$\Pr[ci(Q, X) \geq \Gamma + \frac{\beta}{2}] = \Pr\left[\sum_{i,j} Y_{i,j} \leq (1 - \Gamma - \frac{\beta}{2} - \Delta)|Q||X|\right]$$

for some $\Delta > 0$. Thus, from Hoeffding's Inequality it follows that

$$q = \Pr[ci(Q, X) \geq \Gamma + \frac{\beta}{2}] < \exp(-2(\Gamma + \frac{\beta}{2} - p)^2 |Q||X|) \leq \exp(-2(\frac{\beta}{2})^2 \Gamma^2 |Q||X|).$$

Let FP be the set of false positives, that is $FP = \{X \in \mathcal{D} \mid ci(Q, X) \geq \Gamma + \frac{\beta}{2} \text{ and } \Gamma dist(Q, X) > cR\}$, then $\Pr[\mathcal{P}2] = \Pr[|FP| \leq \beta S]$. Therefore, it suffices to show the latter is larger than $\frac{1}{2}$.

Let X_1, \dots, X_S denote the elements of \mathcal{D} . For every $1 \leq i \leq S$ let $Z_i \sim Ber(q)$ be the Bernoulli random variable which is equal to one if $X_i \in FP$. Then the expected value of the size of FP satisfies

$$E(|FP|) = E(\sum_i Z_i) = \sum_i E(Z_i) = S \cdot q < S \cdot \exp(-2(\frac{\beta}{2})^2 \Gamma^2 |Q||X|).$$

Therefore, from Markov's Inequality it follows that

$$\Pr[|FP| \leq \beta S] \geq 1 - \frac{E[|FP|]}{\beta S} > 1 - \frac{1}{\beta} \exp(-2(\frac{\beta}{2})^2 \Gamma^2 |Q||X|) \geq \frac{1}{2},$$

where the last inequality holds by the assumption on Γ . This finishes the proof.

We are now ready to prove the theorem.

Proof (of Theorem 1). By Lemma 1 properties $\mathcal{P}1$ and $\mathcal{P}2$ hold with constant high probability. Therefore, we may assume these properties hold simultaneously.

Let r be the smallest Γ -distance between Q and an object of \mathcal{D} . Set $t = \lceil \log_c r \rceil$ and $R = c^t$.

Assume first that the algorithm finishes with terminating condition $\mathcal{T}1$, that is at level R at least $1 + \beta S$ Γ -candidates have been found. By property $\mathcal{P}2$ at most βS of these are false positives. Let X be the object returned by the algorithm, then we have $\Gamma dist(Q, X) \leq cR \leq c^2 r$.

Now, if the algorithm does not finish with $\mathcal{T}1$, then property $\mathcal{P}1$ guarantees it finishes with $\mathcal{T}2$ at the end of level R . Let X be the object returned by the algorithm, then we have $\Gamma dist(Q, X) \leq R \leq cr < c^2 r$. This finishes the proof.

5 Experimental Evaluation

In this section, we evaluate the effectiveness of *mmLSH* on four real multimedia data sets. All experiments were run on the nodes of the Bigdata cluster¹ with: two Intel Xeon E5-2695, 256 GB RAM, and CentOS 6.5 operating system. We used the state-of-the-art C2LSH [7] as our base implementation.² All codes were written in C++11 and compiled with gcc v4.7.2 with the -O3 optimization flag. For existing state-of-the-art algorithms (C2LSH and QALSH), we used the Borda Count process (Sect. 4) to aggregate the results of the point queries to find the

¹ Supported by NSF Award #1337884.

² *mmLSH* can be implemented over any state-of-the-art LSH technique.

nearest neighbor objects. Additionally, since the accuracy and the performance of the aggregation is affected by the chosen number of top- k' results of the point queries, we choose a varying k' for Linear, C2LSH, and QALSH for a fair comparison: $k' = 25, 50, 100$. We also implement an LRU buffer for the indexes in C2LSH and QALSH to show a fair comparison with our results. We compare our work with the following alternatives:

- **LinearSearch-Borda:** In this alternative, the top- k' results of the point queries are found using a brute-force linear search. This method does not utilize the buffer since it does not have any indexes.
- **C2LSH-Borda:** top- k' results of point queries are found using C2LSH [7].
- **QALSH-Borda:** top- k' results of point queries are found using QALSH [9].

5.1 Datasets

We use the following four real multimedia datasets to evaluate *mmLSH*:

- **Caltech** [3] This dataset consists of 3,767,761 32-dimensional points that were created using BRIEF on 28,049 images belonging to 256 categories.
- **Corel** [5] This dataset consists of 1,710,725 64-dimensional points that were created using SURF on 9,994 images belonging to 100 categories.
- **MirFlicker** [14] This dataset consists of 12,004,143 32-dimensional points that were created using ORB on 24,980 images.
- **Wang** [21] This dataset consists of 695,672 128-dimensional SIFT descriptors belonging to 1000 images. These images belong to 10 different categories.

5.2 Evaluation Criteria and Parameters

We evaluate the execution time and accuracy using the following criteria:

- **Time:** Two main dominant costs in LSH-based techniques are the algorithm time and the index IO time. We observed that index IO times were not consistent (i.e. running the same query multiple times, would return drastically different results, mainly because of disk cache and instruction cache issues). Thus, the overall execution time is modeled for an HDD where an average disk seek requires 8.5 ms and an average data read rate is 0.156 MB/ms [18].
- **Accuracy:** Similar to the ratio defined in earlier works [7, 9], we define an object ratio to calculate the accuracy of the returned top- k objects as following: $OR_{\Gamma}(Q) = \frac{1}{k} \sum_{i=1}^k \frac{\Gamma dist(Q, X_i)}{\Gamma dist(Q, X_i^*)}$ where X_1, \dots, X_k denote the top- k objects returned from the algorithm and X_1^*, \dots, X_k^* denote the real objects found from the ground truth. $\Gamma dist$ is computed using Eq. 2. Object Ratio of 1 means 100% accuracy and as it increases, the accuracy decreases.

We do not report the index size or the index construction cost, since they would be the same as the underlying LSH implementation that we use (C2LSH [7]).

We choose $\delta = 0.1$, $\beta = \frac{25}{S}$, $\varepsilon = 0.2$, $w = 2.184$ [9] for C2LSH and mmLSH, $w = 2.7191$ [9] for QALSH. We randomly chose 10 multimedia objects as queries from each dataset and report the average of the results.

5.3 Discussion of the Results

In this section, we analyze the execution time and accuracy of *mmLSH* using the criteria explained in Sect. 5.2 against its alternatives. We note that QALSH gives drastically worse times than C2LSH, and hence when comparing the effectiveness for varying parameters, we only compare with C2LSH.

Effect of Buffer Size: Figure 2(a) shows the benefit of our eviction strategy (Sect. 4.3) when compared with C2LSH + LRU for varying buffer sizes. It is evident from this figure that our three criterion are helpful in evicting less useful




Fig. 2. Effect of (a) Buffer size on time, (b) Varying k on time and accuracy




Fig. 3. Comparison of time of *mmLSH* against alternatives




Fig. 4. Comparison of accuracy of *mmLSH* against alternatives

index files. The small overhead in Algorithm time is offset by a significant reduction in the number of index IOs. We found out that the highest L3 cache size is 24.75 MB for desktop processors and 60MB for server processors. Therefore, we decided to choose the buffer sizes between 20 MB and 50 MB. We use 30 MB as the default cache size in the following experiments.

Effect of Number of Desired Objects: Figure 2(b) shows the execution time and accuracy of *mmLSH* against C2LSH for varying number of desired objects (k). This figure shows that *mmLSH* has better time and object ratio for different k values. Additionally, it shows that *mmLSH* is scalable for a large number of desired objects as well. Moreover, although the object ratio of *mmLSH* stays the same by increasing k , the object ratio of C2LSH increases. We use $k = 25$ as the default for the following experiments.

Comparison of mmLSH vs. State-of-the-art Methods. Figures 3 and 4 show the time and accuracy of *mmLSH*, LinearSearch-Borda, C2LSH-Borda, QALSH-Borda for 4 multimedia datasets with varying characteristics. The Borda count process is done after query processing and takes very negligible time. Note that, In our work, we consider all feature-vectors that are extracted by a feature-extraction algorithm. Several works [12, 20] have been proposed that cluster these points with the purpose of finding a representative point to reduce the complexity and overall processing time of the problem. Our work is orthogonal to those approaches and hence are not included in this paper.

For the Caltech and MirFlickr datasets, QALSH did not finish the experiments due to their slow execution and hence are not included in the charts. The slow execution is mainly due to the use of the B+-tree index structures to find the nearest neighbors in the hash functions. *mmLSH* always returns a higher accuracy than the alternatives while being much faster than all three alternatives. This is because *mmLSH* is able to leverage the common elements between queries and improve cache utilization along with being able to stop earlier than the state-of-the-art algorithms. For future work, we plan on investigating the application of *mmLSH* to other distance measures and compare with other feature vector aggregation techniques [11].

6 Conclusion

In this paper, we presented a novel index structure for efficiently finding top- k approximate nearest neighbors for multimedia data using LSH, called *mmLSH*. Existing LSH-based techniques can give theoretical guarantees on these individual high-dimensional feature vector queries, but not on the multimedia object query. These techniques also treat each individual feature vector belonging to the object as independent of one another. In *mmLSH*, novel strategies are used that improve execution time and accuracy of a multimedia object query. Additionally, we provide rigorous theoretical analysis and guarantees on our returned results. Experimental evaluation shows the benefit of *mmLSH* in terms of execution time and accuracy compared to state-of-the-art algorithms. Additionally,

mmLSH can give theoretical guarantees on the final results instead of the individual point queries.

References

1. Bartolini, I., Ciaccia, P., Patella, M.: Query processing issues in region-based image databases. *Knowl. Inf. Syst.* **25**, 389–420 (2010). <https://doi.org/10.1007/s10115-009-0257-4>
2. Arora, A., Sinha, S., Kumar, P., Bhattacharya, A.: Hd-index: pushing the scalability-accuracy boundary for approximate kNN search. In: VLDB (2018)
3. Caltech dataset. http://www.vision.caltech.edu/Image_Datasets/Caltech256
4. Christiani, T.: Fast Locality-sensitive hashing frameworks for approximate near neighbor search. In: Amato, G., Gennaro, C., Oria, V., Radovanović, M. (eds.) SISAP 2019. LNCS, vol. 11807, pp. 3–17. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32047-8_1
5. Corel dataset. <http://www.ci.gxnu.edu.cn/cbir/Dataset.aspx>
6. Datar, M., Immorlica, N., Indyk, P., Mirrokni, V.S.: Locality-sensitive hashing scheme based on p-stable distributions. In: SOCG (2004)
7. Gan, J., Feng, J., Fang, Q., Ng, W.: Locality-sensitive hashing scheme based on dynamic collision counting. In: SIGMOD (2012)
8. Gionis, A., Indyk, P., Motwani, R.: Similarity search in high dimensions via hashing. In: VLDB (1999)
9. Huang, Q., Feng, J., Zhang, Y., Fang, Q., Ng, W.: Query-aware locality-sensitive hashing for approximate nearest neighbor search. In: VLDB (2015)
10. Jafari, O., Ossorgin, J., Nagarkar, P.: qwLSH: cache-conscious indexing for processing similarity search query workloads in high-dimensional spaces. In: ICMR (2019)
11. Jégou, H., Douze, M., Schmid, C.: Improving bag-of-features for large scale image search. *Int. J. Comput. Vis.* **87**, 316–336 (2010). <https://doi.org/10.1007/s11263-009-0285-2>
12. Krizaj, J., Štruc, V., Pavešić, N.: Adaptation of SIFT features for robust face recognition. In: Campilho, A., Kamel, M. (eds.) ICIAR 2010. LNCS, vol. 6111. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-13772-3_40
13. Liu, W., Wang, H., Zhang, Y., Wang, W., Qin, L.: I-LSH: I/O efficient c-approximate nearest neighbor search in high-dimensional space. In: ICDE (2019)
14. MirFlicker dataset. <http://press.liacs.nl/mirflickr>
15. Nagarkar, P., Candan, K.S.: PSLSH: an index structure for efficient execution of set queries in high-dimensional spaces. In: CIKM (2018)
16. Perez, C.A., Cament, L.A., Castillo, L.E.: Methodological improvement on local Gabor face recognition based on feature selection and enhanced Borda count. *Pattern Recogn.* **44**, 951–963 (2011)
17. Reilly, B.: Social choice in the south seas: electoral innovation and the Borda count in the Pacific Island countries. *IPSR* **23**, 355–372+467 (2002)
18. Seagate ST2000DM001 Manual. <https://www.seagate.com/files/staticfiles/docs/pdf/datasheet/disc/barracuda-ds1737-1-1111us.pdf>
19. Sundaram, N., et al.: Streaming similarity search over one billion tweets using parallel locality-sensitive hashing. In: VLDB (2013)
20. Tao, C., Tan, Y., Cai, H., Tian, J.: Airport detection from large IKONOS images using clustered SIFT keypoints and region information. In: GRSL (2011)

21. Wang, J.Z., Li, J., Wiederhold, G.: Simplicity semantics-sensitive integrated matching for picture libraries. *TPAMI* **23**, 947–963 (2001)
22. Wu, Z., Ke, Q., Isard, M., Sun, J.: Bundling features for large scale partial-duplicate web image search. In: *CVPR* (2009)
23. Zhou, W., Li, H., Lu, Y., Tian, Q.: Large scale image search with geometric coding. In: *MM 2011* (2011)