

Indexing High-Dimensional Data for Content-Based Retrieval in Large Databases

Manuel J. Fonseca, Joaquim A. Jorge
Department of Information Systems and Computer Science
INESC-ID/IST/Technical University of Lisbon
R. Alves Redol, 9, 1000-029 Lisboa, Portugal
mjf@inesc-id.pt, jorgej@acm.org

Abstract

Many indexing approaches for high-dimensional data points have evolved into very complex and hard to code algorithms. Sometimes this complexity is not matched by increase in performance. Motivated by these ideas, we take a step back and look at simpler approaches to indexing multimedia data. In this paper we propose a simple, (not simplistic) yet efficient indexing structure for high-dimensional data points of variable dimension, using dimension reduction. Our approach maps multidimensional points to a 1D line by computing their Euclidean Norm and use a B⁺-Tree to store data points. We exploit B⁺-Tree efficient sequential search to develop simple, yet performant methods to implement point, range and nearest-neighbor queries.

To evaluate our technique we conducted a set of experiments, using both synthetic and real data. We analyze creation, insertion and query times as a function of data set size and dimension. Results so far show that our simple scheme outperforms current approaches, such as the Pyramid Technique, the A-Tree and the SR-Tree, for many data distributions. Moreover, our approach seems to scale better both with growing dimensionality and data set size, while exhibiting low insertion and search times.

1. Introduction

In recent years, increasing numbers of computer applications in CAD, geography, biology, medical imaging, etc., access data stored on large databases. A common feature to such databases is that objects are described by vectors of numeric values, known as *feature vectors*, which map individual instances to points in a high dimensional vector space. An important functionality that should be present in such applications is similarity search, *i.e.* finding a set of objects similar to a given query. The similarity between

complex objects is not measured on their contents directly, since this tends to be expensive, but using their feature vectors. This way, searching objects by similarity in a database becomes a nearest neighbor search in a high-dimensional vector space.

To support processing large amounts of high-dimensional data, a variety of indexing approaches have been proposed in the past few years. Some of them are structures for low-dimensional data that were adapted to high-dimensional data spaces. However, such methods while providing good results on low-dimensional data, do not scale up well to high-dimensional spaces. Recent studies [13] show that many indexing techniques become less efficient than sequential search, for dimensions higher than ten. Other indexing mechanisms are incremental evolutions from existing approaches, where sometimes, the increased complexity does not yield comparable enhancements in performance. Other indexing techniques based on dimension reduction return only approximate results. Finally, structures combining several of the previous approaches have emerged. Often, the corresponding algorithms are very complex and unwieldy.

Our interest in this problem derives from the development of effective and efficient systems for content-based retrieval of technical drawings. In the approach we presented in [6], technical drawings are described using topology graphs, which will be converted in feature vectors by computing and combining eigenvalues from their adjacency matrix. The dimensionality of the resulting feature vectors will depend of the complexity of the technical drawing. Thus, more complex drawings will produce descriptors of higher dimension.

All indexing structures studied so far, only support data sets of fixed dimension. However, in some application domains, as the one described before, the dimension of feature vectors can vary from object to object and the maximum dimension can not be predicted in advance. In such sce-

narios, current indexing structures have to rebuilt all data structures, every time the maximum dimension changes.

The increasingly complex data structures and specialized approaches to high-dimensional indexing make it difficult to ascertain whether there might be a reasonably fast and general approach to address variable dimension data. We believe there might be some merit in taking a step back and looking at simpler approaches to indexing such data.

Unlike other existing approaches, we use a very simple dimension reduction function to map high-dimensional points into a one-dimensional value. Our approach is motivated by three observations. First, real data sets tend to have a lot of clusters distributed along the data space. Thus, the Euclidean norms of points tend to be "evenly" distributed (see section 4). Second, the set of resulting nearest neighbors will lie inside a narrow range of norm values, thus reducing the number of data points to examine during search. Third, some application domains need to manipulate large amounts of variable dimension data points, which calls for an approach that can handle variable data in a natural manner.

Based on these requirements, we developed the NB-Tree¹, an indexing technique based on a simple, yet efficient algorithm to search points in high-dimensional spaces with variable dimension, using dimension reduction. Multidimensional points are mapped to a 1D line by computing their Euclidean Norm. In a second step we sort these mapped points using a B⁺-Tree on which we perform all subsequent operations. Thus, the NB-Tree can be implemented on existing DBMSs without additional complexity.

Our approach supports all typical kinds of queries, such as point, range and nearest neighbor queries (KNN). Also, our method provides fast and accurate results, in contrast to other indexing techniques, which return approximate results. Moreover, for some methods [15] when accuracy increases, their performance decreases.

We implemented the NB-Tree and evaluated its performance against more complex approaches, such as, the Pyramid Technique, the A-Tree and the SR-Tree. We conducted a set of experiments, using synthetic and real data, to analyze creation, insertion and query times as a function of data set size and dimension. Results so far show better performance for our method, for many data distributions. Moreover, our approach seems to scale better both with growing dimensionality and data set size, while exhibiting low insertion and search times.

The rest of the paper is organized as follows. In the next section we give an overview of the related work in high-dimensional indexing structures. Section 3 explains the basic algorithms of the NB-Tree and their complexity. In section 4 we present experimental studies to characterize the different data sets used for evaluation. Section 5 describes

our experimental evaluation and shows performance results. We conclude the paper by discussing the results and presenting further work directions in our algorithm and related areas.

2. Related Work

Indexing techniques developed so far can be classified into two categories. One that includes all structures derived from the K-D-Tree and the other composed by derivatives of the R-Tree. The main difference between these two categories lies in the approach to dividing the data space. Structures in the first category use space-partitioning methods that divide the data space along predefined hyperplanes regardless of data distribution. The resulting regions are mutually disjoint, with their union being the complete space. Structures from the second class use data-partitioning methods, which divide the data space according to their distribution. This can yield possible overlapping regions. Besides these two categories, there are other techniques combining several methodologies to improve performance.

In [7] Henrich proposed the LSD^h-Tree as an improvement to the LSD-Tree. This tree combines the low fanout from K-D-Trees with the advantage that R-Trees have in not covering empty space. This structure is slightly better than the X-Tree.

The Hybrid-Tree [3] was introduced by Chakrabarti in 1999 to combine the advantages from space-partitioning structures and data-partitioning structures. The Hybrid-Tree always splits a node using a single dimension to guarantee that the fanout is independent of data dimension and it allows overlapping regions. This structure outperforms both the SR-Tree and the hB-Tree.

White and Jain presented the SS-Tree [14], an R-Tree-like index structure that uses minimum bounding spheres (MBSs) instead of minimum bounding rectangles (MBRs). Even though the use of spheres reduce the overlapping of regions and consequently the SS-Tree outperforms the R*-Tree, spheres tend to overlap for high-dimensional spaces.

Recently, Katayama proposed the SR-Tree [8], an improvement to the SS-Tree combining concepts from both the R*-Tree and SS-Tree approaches. This structure uses both MBRs and MBSs as an approximation. This way, the region spanned by a node is the intersection of the MBR and MBS associated to that node, thus reducing region overlap between sibling nodes. This structure has been shown to outperform both the R*-Tree and the SS-Tree. In [11] Sakurai presents an experimental evaluation, using non-uniform data, where the SR-Tree outperforms the VA-File [13].

Berchtold, Keim and Kriegel proposed the X-Tree [2], an index structure adapting the algorithms of R*-Tree to high-dimensional data. The X-Tree uses an overlap-free split

¹Norm + B⁺-Tree = NB-Tree

algorithm based on the split history of the tree. When the overlap-free split algorithm would lead to an unbalanced directory, the X-Tree omits the split and the according directory node is enlarged becoming a super-node. The X-Tree outperforms both the R*-Tree and the TV-Tree [9].

Recently, Sakurai proposed the A-Tree [11] that introduced the notion of relative approximation. The basic idea of the A-Tree is the use of Virtual Bounding Rectangles to approximate minimum bounding rectangles or objects. Thus, on one node of the tree we have information about the exact position of the region MBR and an approximation of the relative position of its sons. Authors claim that the A-Tree outperforms the VA-File and the SR-Tree.

In 1998 Weber et al. [13] proposed the VA-File, a method to search points of high dimension. The authors show that existing structures (R*-Tree and X-Tree) are outperformed on average by a simple sequential scan if the dimension exceeds around ten. So, instead of developing another indexing structure they proposed to speed-up the sequential scan. The basic idea of the VA-File is to keep two files; one with an approximate version of data points and other with the exact representation. When searching points, the approximation file is sequentially scanned with some look-ups to the exact file whenever it is necessary. The VA-File outperforms both the R*-Tree and the X-Tree when the dimension is higher than six, but its performance is very sensitive to data distribution.

Berchtold et al.'s Pyramid Technique [1] uses a special partitioning strategy optimized for high-dimensional data. Its basic idea is to perform a dimension reduction allowing the use of efficient uni-dimensional index structures to store and search data. The Pyramid Technique divides the data space into 2D pyramids whose apexes lie at the center point. In a second step, each pyramid is cut into several slices parallel to the basis of the pyramid forming the data pages. The Pyramid Technique associates to each high-dimensional point a single value, which is the distance from the point to the top of the pyramid, according to a specific dimension (j_{max}). This value is later used as a key in the B⁺-Tree. Points in the same slice will have the same identification value. As a result, for skewed distributions, many points in the B⁺-Tree will be indexed by the same key. The Pyramid Technique outperforms the X-Tree and the sequential scan using uniformly distributed data. Although the Pyramid Technique was developed mainly for uniformly distributed data, authors suggested an extended version to handle real (skewed) data. This version works by shifting the center point to the barycenter of the data. However, in a dynamic scenario, this implies the recalculation of all index values, *i.e.* redistribution of points among the pyramids and reconstruction of the B⁺-Tree, each time the center of the cluster changes. We have receded from Berchtold's work by keeping the dimension reduction approach

and focusing on reducing data-structure traversal and book keeping overheads.

More recently, Berchtold et al. proposed the IQ-Tree [10] which combines a compression method with an index structure trying to join the best of both worlds. The IQ-Tree is a three-level structure, with a regular directory level, a compression level and a third level that contains the actual data. While the VA-File uses a fixed number of bits to compress all partitions, the IQ-Tree uses a different compression scheme for each partition depending on the density of data points. This structure outperforms both the X-Tree and the VA-File.

In [12] Shepherd et al. presented an overview of an efficient file-access method for similarity searching in high-dimensional spaces, named CurveIx. The basic idea is to order the d-dimensional space in many ways, with a set of (one-dimensional) space-filling curves. Each curve constitutes a mapping from $R^d \rightarrow R^1$, yielding a linear ordering of all points in the data set. Their approach uses standard one-dimensional indexing schemes such as the B-Tree to perform searching. Close points along the space-filling curve tend to correspond to close points in the d-dimensional feature space. However, some near neighbors may be mapped far apart along a single space-filling curve. As a result, this method does not have an accuracy of 100%, *i.e.* some near neighbors may be ignored.

More recently, Yu et al. proposed a new index scheme, called iMinMax(θ) [15], that maps high-dimensional points to single dimension values determined by their maximum or minimum coordinate values. By varying the θ value they can optimize the iMinMax structure to different distributions of data sets. As other dimension reduction methods, this scheme also uses the B⁺-Tree to index the resulting single dimension points. The iMinMax structure was mainly developed to address window search (range queries) while still supporting approximate KNN search (accuracy less than 100%) at the expense of increasing runtimes for higher accuracy. This approach outperforms both the VA-File and the Pyramid Technique for range queries.

The research work outlined above highlights two main patterns in previous developments. First, there is no total order connecting any two algorithms so far developed. This suggests both the need to do some research work on developing reference data sets representative of significant problem domains to ensure commensurable results. Second, due to the diversity of problem domains, it is difficult to establish emerging techniques as related to previous work. Sometimes, even when reference implementations of some algorithms exist, problems with existing code prevent other researchers from duplicating experimental results. It becomes difficult to compare results either when the reference implementation does not provide correct results or the programs terminate abnormally on large data sets.

3. The NB-Tree Structure

The NB-Tree provides a simple and compact means to indexing high-dimensional data points of variable dimension, using a light mapping function that is computationally inexpensive. The basic idea of the NB-Tree is to use the Euclidean norm value as the index key for high-dimensional points. Thus, values resulting from the dimension reduction can be ordered and later searched in the resulting one-dimensional structure. To index data points sorted by their Euclidean norm we use the B⁺-Tree, since it is the most efficient indexing structure for one dimension and also, because it is supported by all commercial DBMSs.

The use of the Euclidean norm as a mapping function from $R^D \rightarrow R^1$, assures that near high-dimensional points will have near Euclidean norms. Consequently, when performing a query (of any type) the system only examine points whose norm is in the neighborhood of the query point norm. Moreover, the B⁺-Tree has the particularity that its leaves are linked as an ordered list. Thus, walking sequentially through all the elements of the B⁺-Tree is a costless operation.

To create an NB-Tree we start by computing the Euclidean norm of each D-dimensional point from the data set, using the Formula: $\|P\| = \sqrt{p_0^2 + p_1^2 + \dots + p_{D-1}^2}$. D-dimensional points are then inserted into a B⁺-Tree, using the norm as key. After inserting all points we get a set of D-dimensional data ordered by norm value. Figure 1 shows an example of dimension reduction for 2D data points. As we can see, with this mapping function, different points can have the same key value. On the other hand near data points have similar keys (norms).

Any searching process in the NB-Tree starts by computing the norm of the query point. Then we perform a set of searching steps in the 1-dimensional B⁺-Tree, which will depend of the query type. Current indexing structures usually support three types of queries. The first is Point Query, which checks if a specific point belongs or not to the database. The second type, Range Query, returns the points inside a specific range of values. In our case that range will be specified by an hyper-ball. Finally, the *k*-NN Query (*k* Nearest Neighbors) returns the *k* nearest neighbors of the query point. Since this is the most often-used query in content-based retrieval and due to space limitations, we only describe this type of query.

Before we start describing our *k*-NN algorithm we have to define some relevant concepts: **higherLimit** and **lowerLimit** are the norm values used to define the upper and lower bounds to search for nearest neighbors; **delta** - value by which the previous limits are incremented at each step; **priority list** - list with *k* nodes, where we store *k* points that temporarily store for intermediate results. The last element of this list is the farthest nearest neighbor so far. Thus, when




Figure 1. NB-Tree dimension reduction.

we want to insert a new neighbor into the list, we just have to compare its distance against the last.

We start the nearest neighbor search by locating a point in the B⁺-Tree with a key equal (or near) the norm of the query point. After this positioning in the 1D line, we execute a set of iterative steps until we get all the desired results. Since the B⁺-Tree has its leaves linked sequentially, we only navigate at the leaves level, as if we were using a double linked ordered list. Below we describe all the steps to achieve the final *k*-NN points.

After positioning in the 1D line we define the *upperLimit* and the *lowerLimit* based on the *delta* value. Next we go through all points until we reach the *upperLimit* or the point whose norm follows the limit. In this case the *lowerLimit* is changed to keep the symmetry to the query. After, we do the same for the *lowerLimit*. During this forward and backward advance we compute the distance from each point to the query point. As described before, if the distance to the query point is smaller than the distance of the farthest current near neighbor, we store that point in the list. After going up and down, we check whether there are enough points inside the hypersphere of diameter equal to the difference between limits. This iterative process happens until all *k* neighbors from the list are contained in the previously defined hypersphere.

In summary, we compute the *k*-NN through an iterative process, where the size of the searching ball increases gradually until we get the number of points specified by the query.

3.1. Computational Complexity

In this subsection, we present a short analysis of the computational complexity of the NB-Tree as a function of data point dimension and data set size. For a more detailed analysis see [5].

It is important for our complexity study to recall that the B⁺-Tree has a computational complexity of $O(\lg N)$ for insertion and searching algorithms.

NB-Tree Algorithm	N	D
Insertion	$O(\lg N)$	$O(D)$
Point Query	$O(\lg N)$	$O(D)$
Range Query	$O(N)$	$O(D)$
k -NN Query	$O(N)$	$O(D)$

Table 1. NB-Tree complexity analysis.

In Table 1 we present a summary of the complexity analysis. As we can see, considering the size of the data set (N), the NB-Tree presents a logarithmic running time for insertion and point query, since in this two algorithms the heaviest operation is the access to the B^+ -Tree. For range queries and k -NN queries we have linear running time because computation of distances is the more time consuming operation.

If we now consider the dimension (D), all NB-Tree algorithms have a worst-case running time linear with the dimension.

4. Data Set Characterization

Before describing our experimental evaluation, we present an experimental study characterizing the different data distributions we used to evaluate our NB-Tree. The first data sets analyzed were of synthetic uniformly distributed data, with independent attributes. After, we analyzed a set of real data from the NIST database, where data points have a dimension of 256 and coordinate values can each take the value zero or one. Finally, we examined a synthetic data set, where data points have different dimensions, simulating a data space of variable dimension.

We start our data analysis by computing statistical moments, such as minimum, maximum, average, median and standard deviation, for the Euclidean norm of all uniformly distributed data sets. Results revealed that average and median of all data sets have the same value (for all dimensions), meaning that the Euclidean norm has a normal distribution. We also observed that the amplitude (maximum - minimum) of the norm values is always constant for all dimensions and that the minimum, the maximum, the average and the median present the same evolution along the dimension. A more detailed analysis of this type of data set is described in [5].

Our next step was to study the distribution of points according to the Euclidean norm. To that end we compute the mode of the norm for each distribution, using intervals of 0.1 and considering data points of dimension 256^2 . The main objective of this study was to achieve some conclusions about the behavior of the Euclidean norm when we move from uniformly distributed data to real data and see

²For the data set of variable dimension 256 is the maximum dimension.

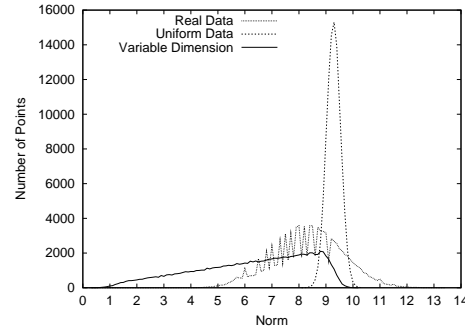


Figure 2. Distribution of norms.

what kind of data sets yield worst performance for the NB-Tree approach.

As we can see, in Figure 2, for uniformly distributed data the values of the norm form a "tight" cluster around a single value. In the case of the real data and of points with variable dimensions, the values of the norm are more evenly distributed. This wide distribution of the norm values seems to validate the NB-Tree approach of using the Euclidean norm as a method to index high-dimensional real data sets of fixed or variable dimensions.

5. Experimental Evaluation

While our approach seems to yield commensurable times to other structures tested, we had difficulties comparing it to other approaches, since some of them crashed on data sets of significant size, preventing comparison. We chose the SR-Tree, the A-Tree and the Pyramid Technique as benchmarks because they are the more recent indexing structures and because there are reliable and stable implementations, which provide correct results and scale up to our intended test data sizes.

We conducted a set of experiments to analyze final tree size (on disk), creation times and query times as a function of data set dimension and size. All experiments were performed on a PC Pentium II @ 233 MHz running Linux 2.4.8, with 384 MB of RAM and 15GB of disk. Our NB-Tree algorithms were developed using the *dbm* implementation of the B^+ -Tree.

First, we present the results of the evaluation for several data sets of uniformly distributed data. Then, we introduce our observations for some real data sets and for a data set of variable dimension. While data points generated from a uniformly distributed set of coordinates arguably represent any "real" distribution, most published approaches provide performance measurements and comparisons based on such data. Thus it becomes logical to evaluate our indexing structure using these data sets. Moreover, as we will show in the present section, many algorithms (including ours) seem to perform poorly with these data, which makes better suited

for the task at hand than most empirical data sets.

We did not measure I/O accesses for four reasons: First, nowadays computers can handle databases of millions of high-dimensional points in main memory. Second, after a dozen queries all the indexing structure is loaded into main memory. Third, users of content-based retrieval systems do not care about the number of I/O operations, but they do care about response time. Fourth, in many indexing schemes, data structure overhead dominates I/O, especially after most of the database is memory-resident. Finally, we assume that locality of reference holds for well-behaved features. Thus, during our experimental evaluation we start by loading all the structures into main memory and after that we measure search times. We think that the time taken to load the NB-Tree into main memory is not significant, taking into account that it is done just one time. The NB-Tree takes less than one minute to load, for dimensions up to 100 and less than 4 minutes for dimensions up to 256.

5.1. Uniformly Distributed Data Sets

We evaluated the structures using data sets of randomly generated uniformly distributed data points of fixed size (100,000) and various dimensions (10, 20, 30, 40, 60, 80, 100, 150 and 256). We also created data sets with fixed dimension (20) and various sizes (250,000, 500,000, 750,000 and 1,000,000). Additionally, we randomly generated a set of 100 queries for each dimension, which we later used to evaluate the searching performance of each approach. We selected the number of nearest neighbors to search for to be always ten.

Due to space limitations, we only present results for creation, k -NN search and final tree size. Experimental results for range queries can be found in [5].

Creation Times Most published work tends to ignore tree insertion times. This is because conventional scenarios focus on large static databases which are far more often queried upon than updated. However, there are many applications requiring frequent updating of data sets. For these applications, low insertion times are an important usability factor.

We have compared the creation times to these of the SR-Tree, A-Tree and Pyramid Technique. Figure 3.a shows the time spent to create each structure when the dimension of the data changes. We do not present the creation times for the A-Tree because they are very high, even for low dimensions (around 300 minutes for dimensions 10 and 20). As we can see, the NB-Tree and the Pyramid Technique largely outperform the SR-Tree. While the NB-Tree takes 24 seconds to insert 100,000 points of dimension 10, the SR-Tree takes 23 minutes. If we now consider higher dimensions, such as 80, the difference increases even more with the NB-

Tree taking 2 minutes and the SR-Tree taking 40 minutes. The Pyramid Technique reveals creation times below the two minutes, for dimensions up to 256.

Figure 3.b shows insertion times for data sets of varying size. Although, all structures seem to exhibit linear growth with dimension, SR-Tree creation times grow faster than those of the NB-Tree and of the Pyramid Technique. While the NB-Tree requires no more than 15 minutes to create a tree with one million data points, the SR-Tree takes around six hours and the Pyramid Technique takes only two minutes. From this observation it is clear that the Pyramid Technique and the NB-Tree are more suited to large data sets than the SR-Tree.

We were not able to create the SR-Tree for data sets of dimension bigger than 100, in our system. Thus, in the following subsections we do not display values for dimensions higher than 100 corresponding to the SR-Tree.

k -NN Query Times Nearest-neighbor queries are useful when we want to look at the point in the data set which most closely matches the query.

Figure 3.c depicts the performance of nearest neighbor searches when data dimension increases. We can see that the NB-Tree outperforms all the structures evaluated, for any characteristic dimension of the data set. Our approach computes the ten nearest neighbors in less than one second for dimensions up to 40, less than two seconds for dimensions up to 100 and less than five seconds for dimensions up to 256. Moreover, we can notice that the NB-Tree shows linear behavior with the dimension while the SR-tree and the A-Tree seems to exhibit at least a quadratic growth or worse. The Pyramid Technique performs better than the A-Tree and the SR-Tree, but it shows worst search times than the NB-Tree, for all dimensions. Figure 3.d also shows that the NB-Tree outperforms all structures for k -NN queries when the size of the data set increases.

Storage Requirements From our experimental evaluation we realized that the NB-Tree requires less storage space than any other structure. The SR-Tree uses at least three times more storage space than the NB-Tree, while the Pyramid Technique requires more than twice the storage space used by our structure. The A-Tree uses a storage space of the same size of the Pyramid Technique. Furthermore, the SR-Tree size seems to increase quadratically with the data dimension while the NB-Tree and the Pyramid Technique grow linearly. This linear growth is inherited from the B^+ -Tree. We also observed that all structures presented a linear growth with data set size. However, the storage requirements from our approach grow slower than those from the SR-Tree and the Pyramid Technique. Even though, the NB-Tree and the Pyramid Technique share a B^+ -Tree to store information and even though they store the same informa-

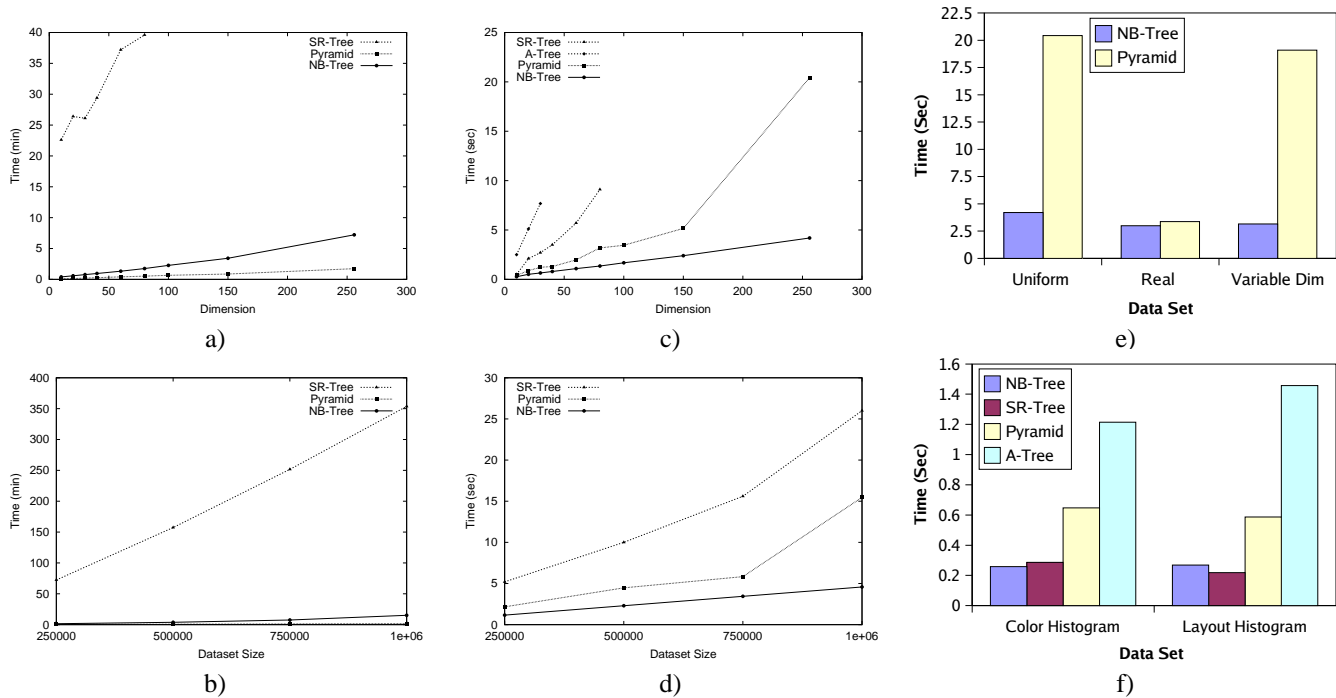


Figure 3. Experimental results for the NB-Tree.

tion (high-dimensional point + 1D value) the files produced by the Pyramid Technique are around twice the size of the ones from the NB-Tree. The only explanation that we encounter is the different implementations used. We used the `dbm` version, while the Pyramid Technique uses their own.

5.2. Variable Dimension Data Sets

After our experimental evaluation with synthetic uniformly distributed data sets, we tested the indexing structures with other synthetic data set that try to simulate data points of variable dimensions. To that end, we filled the first N coordinates of points with values and the rest $(D - N)$ with zeros. We generate points with dimensions between 4 and 256 (but all points have 256 coordinates). We had to do this, because the other structures do not support points of different dimensions at the same time.

From Figure 3.e (right columns) we can see that the NB-Tree is six times faster than the Pyramid Technique finding the ten nearest neighbors. We do not present values for the A-Tree and the SR-Tree, because these structures can not handle data sets of this dimension (256).

5.3. Empirical Data Sets

Finally, we evaluate the indexing structures using three real data sets. Two of them had dimension 32 and contain

image features extracted from Corel image collection. One contains 68,000 color histograms, while the other contains 66,000 color histogram layouts. The last real data set has 100,000 points of dimension 256, each one representing an instance of a hand-drawn digit. This data set is part of the NIST database.

From Figure 3.e we can see that for any kind of data set with dimension 256, the NB-Tree is always faster than the Pyramid Technique and that the NB-Tree takes less time for the real data set.

Figure 3.f presents times for the two real data sets of dimension 32. As we can see the NB-Tree and the SR-Tree have the best performance, with query times around 270 milliseconds. The A-Tree once again presents the worst results, taking five times longer than our approach.

Contrary to published [11] experimental values, we have found better query times for the SR-Tree over the A-Tree. Maybe a fine tuning would yield a different outcome. This did not succeed for the many combinations we have tried.

6. Conclusions

In this paper we proposed a simple, yet efficient structure for data of highly-variable dimensionality, the NB-Tree, using the Euclidean norm to reduce dimension and a B^+ -Tree to store the 1D values. We have conducting experimental evaluations using either uniformly distributed data

sets, variable dimension and real data. Results show our method to outperform the Pyramid Technique, the SR-Tree and the A-Tree.

Our experimental evaluation indicates that the NB-Tree can efficiently support a wide type of queries, including point, range and the most used (at least in multimedia databases) nearest neighbor queries. This is significant because most indexing structures proposed for range queries are not designed to efficiently support similarity search queries (KNN) and metric-based approaches proposed for similarity queries are considerably more difficult to apply to range queries. In contrast to other approaches which use complex algorithms (combination of MBRs and MBSs to describe partitions, division of the space using hyperpyramids, etc.) ours relies on very simple (and therefore practical) methods.

Tree-based techniques do not seem to scale up well with growing data set sizes typical of multimedia databases, where it is relatively easy to assemble large collections of documents. The overhead introduced by data-structure traversal and book keeping in these approaches far outweighs the potential advantages of sophisticated spatial indexing structures.

For a seemingly "bad" point set, generated from uniformly distributed data, both the range and k -NN search algorithms can degenerate into a sequential search, where we have to compute the distance from the query point to all the points in the data set. However, as we have seen in section 4 real data and data of variable dimension both exhibit Euclidean norm distributions that better match our selection and indexing methods. We conjecture that our approach tends to excel for sparsely-distributed data where clusters tend to be evenly distributed throughout the hyperspace. Moreover, since our algorithms are very simple, response times remain reasonable for data sets of considerable size and dimension

We are now using the NB-Tree in a trainable shape recognizer [4] as KNN mechanism. Results so far show that query times are very low, making it a good choice for interactive applications where immediate feedback is required.

In summary our work presents a simple method which is able to handle data of variable dimensions in an efficient and scalable manner suitable for interactive use.

Acknowledgments

This work was funded in part by the Portuguese Foundation for Science and Technology, project 34672/99 and the European Commission, project SmartSketches IST-2000-28169.

References

- [1] S. Berchtold, C. Böhm, and H.-P. Kriegel. The Pyramid-Technique: Towards Breaking the Curse of Dimensionality. In *Proc. of the Int. Conf. on Management of Data (SIGMOD'98)*. ACM Press, 1998.
- [2] S. Berchtold, D. Keim, and H.-P. Kriegel. The X-tree: An Index Structure for High-Dimensional Data. In *Proc. of the 22nd Int. Conf. on Very Large Data Bases (VLDB'96)*, Mumbai(Bombay), India, 1996.
- [3] K. Chakrabarti and S. Mehrotra. The Hybrid Tree: An index structure for high dimensional feature spaces. In *Proc. of the 15th Int. Conf. on Data Engineering (ICDE'99)*, pages 440–447, 1999.
- [4] M. J. Fonseca and J. A. Jorge. CALI : A Software Library for Calligraphic Interfaces. INESC-ID, available at <http://immi.inesc-id.pt/cali/>, 2000.
- [5] M. J. Fonseca and J. A. Jorge. Indexing high-dimensional data for content-based retrieval in large databases. Technical report, INESC-ID, <http://virtual.inesc-id.pt/tr/mjf-jaj-TR-01-03.pdf>, 2003.
- [6] M. J. Fonseca and J. A. Jorge. Towards content-based retrieval of technical drawings through high-dimensional indexing. *Computers and Graphics*, To appear, 2003.
- [7] A. Henrich. The LSDh-tree: An access structure for feature vectors. In *Proc. of the 14th Int. Conf. on Data Engineering (ICDE'98)*, pages 362–369, 1998.
- [8] N. Katayama and S. Satoh. The SR-tree: An Index Structure for High-Dimensional Nearest Neighbor Queries. In *Proc. of the Int. Conf. on Management of Data (SIGMOD'97)*, pages 369–380. ACM Press, 1997.
- [9] K.-I. Lin, H. Jagadashi, and C. Faloutsos. The TV-tree: An index structure for high-dimensional data. *The VLDB Journal*, 3(4):517–543, 1994.
- [10] S., C. Bohm, H. V. Jagadish, H.-P. Kriegel, and J. Sander. Independent quantization: An index compression technique for high-dimensional data spaces. In *Proc. of the 16th Int. Conf. on Data Engineering (ICDE'00)*, pages 577–588, San Diego, USA, 2000.
- [11] Y. Sakurai, M. Yoshikawa, S. Uemura, and H. Kojima. The A-tree: An Index Structure for High-Dimensional Spaces Using Relative Approximation. In *Proc. of the 26th Int. Conf. on Very Large Data Bases (VLDB'00)*, pages 516–526, Cairo, Egypt, 2000.
- [12] J. Shepherd, X. Zhu, and N. Megiddo. A fast indexing method for multidimensional nearest neighbor search. In *SPIE Conf. on Storage and Retrieval for Image and Video Databases*, pages 350–355, 1999.
- [13] R. Weber, H.-J. Schek, and S. Blott. A quantitative analysis and performance study for similarity-search methods in high-dimensional spaces. In *Proc. of the 24th Int. Conf. on Very Large Data Bases (VLDB'98)*, pages 194–205, New York, USA, 1998.
- [14] D. White and R. Jain. Similarity indexing with the SS-tree. In *Proc. of the 12th IEEE Int. Conf. on Data Engineering*, pages 516–523, 1996.
- [15] C. Yu, S. Bressan, B. C. Ooi, and K.-L. Tan. Querying high-dimensional data in single dimensional space. *VLDB Journal*, To Appear, 2002.