

Probabilistic Proximity Searching Algorithms Based on Compact Partitions^{*}

Benjamin Bustos² and Gonzalo Navarro^{1,2}

¹ Center for Web Research

² Departamento de Ciencias de la Computación, Universidad de Chile
Blanco Encalada 2120, Santiago, Chile
{bebustos,gnavarro}@dcc.uchile.cl

Abstract. The main bottleneck of the research in metric space searching is the so-called curse of dimensionality, which makes the task of searching some metric spaces intrinsically difficult, whatever algorithm is used. A recent trend to break this bottleneck resorts to probabilistic algorithms, where it has been shown that one can find 99% of the elements at a fraction of the cost of the exact algorithm. These algorithms are welcome in most applications because resorting to metric space searching already involves a fuzziness in the retrieval requirements. In this paper we push further in this direction by developing probabilistic algorithms on data structures whose exact versions are the best for high dimensions. As a result, we obtain probabilistic algorithms that are better than the previous ones. We also give new insights on the problem and propose a novel view based on time-bounded searching.

1 Introduction

The concept of proximity searching has applications in a vast number of fields, for example: multimedia databases, machine learning and classification, image quantization and compression, text retrieval, computational biology, function prediction, etc. All those applications have in common that the elements of the database form a *metric space* [6], that is, it is possible to define a positive real-valued function d among the elements, called *distance* or *metric*, that satisfies the properties of *strict positiveness* ($d(x, y) = 0 \Leftrightarrow x = y$), *symmetry* ($d(x, y) = d(y, x)$), and *triangle inequality* ($d(x, z) \leq d(x, y) + d(y, z)$). For example, a *vector space* is a particular case of metric space, where the elements are tuples of real numbers and the distance function belongs to the L_s family, defined as $L_s((x_1, \dots, x_k), (y_1, \dots, y_k)) = \left(\sum_{1 \leq i \leq k} |x_i - y_i|^s\right)^{1/s}$. For example, L_2 is the *Euclidean distance*.

One of the typical queries that can be posed to retrieve similar objects from a database is a *range query*, which retrieves all the elements within distance r to a query object q . An easy way to answer range queries is to make an exhaustive

^{*} Work supported by the Millenium Nucleus Center for Web Research, Grant P01-029-F, Mideplan, Chile.

search on the database, but this turns out to be too expensive for real-world applications, because the distance d is considered expensive to compute. Think, for example, of a biometric device that computes the distance between two fingerprints.

Proximity searching algorithms build an *index* of the database and perform range queries using this index, avoiding the exhaustive search. Many of these algorithms are based on dividing the space in *partitions* or *zones* as compact as possible. Each zone stores a representative point, called the *center*, and a few extra data that permit quickly discarding the entire zone at query time, without measuring the actual distance from the elements of the zone to the query object, hence saving distance computations. Other algorithms are based in the use of *pivots*, which are distinguished elements from the database and are used together with the triangle inequality to filter out elements of the database at query time.

An inherent problem of proximity searching in metric spaces is that the search becomes more difficult when the “intrinsic” dimension of the metric space increases, which is known as the *curse of dimensionality*. The intrinsic dimension of a metric space is defined in [6] as $\mu^2/2\sigma^2$, where μ and σ^2 are the mean and the variance of the distance histogram of the metric space. This is coherent with the usual vector space definition. Analytical lower bounds and experiments [6] show that all proximity searching algorithms degrade their performance systematically as the dimension of the space grows. For example, in the case of vector space there is no technique that can reasonably cope with dimension higher than 20 [6]. This problem is due to two possible reasons: high dimensional metric spaces have a very concentrated distance histogram, which gives less information for discarding elements at query time; on the other hand, in order to retrieve a fixed fraction of the elements of the space it is necessary to use a larger search radius, because in high dimensional spaces the elements are “far away” from each other.

Probabilistic algorithms are acceptable in most applications that need to search in metric spaces, because in general the modelization as a metric space already carries some kind of relaxation. In most cases, finding some close elements is as good as finding all of them.

There exists a pivot-based probabilistic proximity searching algorithm which largely improves the search time at the cost of missing few elements [5]. On the other hand, it is known that compact partitioning algorithms perform better than pivot-based algorithms in high dimensional metric spaces [6] and they have lower memory requirements.

In this paper we present several probabilistic algorithms for proximity searching based on compact partitions, which alleviate in some way the curse of the dimensionality. We also present experimental results that show that these algorithms perform better than probabilistic algorithms based on pivots, and the latter needs much more memory space to beat the former when the dimension of the space is very high.

2 Basic Concepts

Let (\mathbb{X}, d) be a metric space and $\mathbb{U} \subseteq \mathbb{X}$ the set of objects or database, with $|\mathbb{U}| = n$. There are two typical proximity searching queries:

- *Range query.* A range query (q, r) , $q \in \mathbb{X}$, $r \in \mathbb{R}^+$, reports all elements that are within distance r to q , that is $(q, r) = \{u \in \mathbb{U}, d(u, q) \leq r\}$.
- *k nearest neighbors (k -NN).* Reports the k elements from \mathbb{U} closer to q , that is, returns the set $\mathbb{C} \subseteq \mathbb{U}$ such that $|\mathbb{C}| = k$ and $\forall x \in \mathbb{C}, y \in \mathbb{U} - \mathbb{C}, d(x, q) \leq d(y, q)$.

The volume defined by (q, r) is called the *query ball*, and all the elements inside it are reported. Nearest neighbors queries can be implemented using range queries.

There exist two classes of techniques used to implement proximity searching algorithms: based on pivots and based on compact partitions.

2.1 Pivot-Based Algorithms

These algorithms select a number of “pivots”, which are distinguished elements from the database, and classify all the other elements according to their distance to the pivots.

The canonical pivot-based algorithm is as follows: given a range query (q, r) and a set of k pivots $\{p_1, \dots, p_k\}$, $p_i \in \mathbb{U}$, by the triangle inequality it follows for any $x \in \mathbb{X}$ that $d(p_i, x) \leq d(p_i, q) + d(q, x)$, and also that $d(p_i, q) \leq d(p_i, x) + d(x, q)$. From both inequalities it follows that a lower bound on $d(q, x)$ is $d(q, x) \geq |d(p_i, x) - d(p_i, q)|$. The elements $u \in \mathbb{U}$ of interest are those that satisfy $d(q, u) \leq r$, so one can exclude all the elements that satisfy $|d(p_i, u) - d(p_i, q)| > r$ for some pivot p_i (exclusion condition), without actually evaluating $d(q, u)$.

The index consists of the kn distances $d(u, p_i)$ between every element and every pivot. Therefore, at query time it is necessary to compute the k distances between the pivots and the query q in order to apply the exclusion condition. Those distance calculations are known as the *internal complexity* of the algorithm, and this complexity is fixed if there is a fixed number of pivots. The list of elements $\{u_1, \dots, u_m\} \subseteq \mathbb{U}$ that cannot be excluded by the exclusion condition, known as the *element candidate list*, must be checked directly against the query. Those distance calculations $d(u_i, q)$ are known as the *external complexity* of the algorithm. The total complexity of the search algorithm is the sum of the internal and external complexity, $k + m$. Since one increases and the other decreases with k , it follows that there is an optimum k^* that depends on the tolerance range r of the query. In practice, however, k^* is so large that one cannot store the k^*n distances, and the index uses as many pivots as space permits.

Examples of pivot-based algorithms [6] are *BK-Tree*, *Fixed Queries Tree (FQT)*, *Fixed-Height FQT*, *Fixed Queries Array*, *Vantage Point Tree (VPT)*, *Multi VPT*, *Excluded Middle Vantage Point Forest*, *Approximating Eliminating Search Algorithm (AESA)* and *Linear AESA*.

2.2 Algorithms Based on Compact Partitions

These algorithms are based on dividing the space in *partitions* or *zones* as compact as possible. Each zone stores a representative point, called the *center*, and a few extra data that permit quickly discarding the entire zone at query time, without measuring the actual distance from the elements of the zone to the query object. Each zone can be partitioned recursively into more zones, inducing a *search hierarchy*. There are two general criteria for partitioning the space: *Voronoi partition* and *covering radius*.

Voronoi partition criterion. A set of m centers is selected, and the rest of the elements are assigned to the zone of their closest center. Given a range query (q, r) , the distances between q and the m centers are computed. Let c be the closest center to q . Every zone of center $c_i \neq c$ which satisfies $d(q, c_i) > d(q, c) + 2r$ can be discarded, because its Voronoi area cannot have intersection with the query ball. Figure 1 (left) shows an example of the Voronoi partition criterion. For q_1 the zone of c_4 can be discarded, and for q_2 only the zone of c_3 must be visited.

Covering radius criterion. The covering radius $cr(c)$ is the maximum distance between a center c and an element that belongs to its zone. Given a range query (q, r) , if $d(q, c_i) - r > cr(c_i)$ then zone i cannot have intersection with the query ball and all its elements can be discarded. In Figure 1 (right), the query ball of q_1 does not have intersection with the zone of center c , thus it can be discarded. For the query balls of q_2 and q_3 , the zone cannot be discarded, because it intersects these balls.

Generalized-Hyperplane Tree [14] is an example of an algorithm that uses the Voronoi partition criterion. Examples of algorithms that use the covering radius criterion are *Bisector Trees (BST)* [11], *Monotonous BST* [13], *Voronoi Tree* [8], *M-Tree* [7] and *List of Clusters* [4]. Also, there exist algorithms that use both criteria, for example *Spatial Approximation Tree (SAT)* [12] and *Geometric*

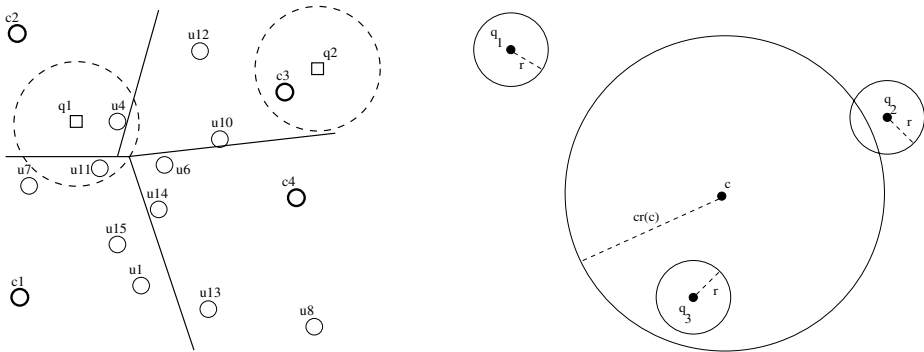


Fig. 1. Voronoi partition criterion (left) and covering radius criterion (right)

Near-neighbor Access Tree [2]. Of all these algorithms, two of the most efficient are SAT and List of Clusters, so now we explain briefly how these algorithms work.

2.3 Spatial Approximation Tree

The SAT [12] is based on approaching the query spatially rather than dividing the search space, that is, start at some point in the space and get closer to the query, which is done only via “neighbors”. The SAT uses both compact partition criteria for discarding zones, it needs $O(n)$ space, reasonable construction time $O(n \log^2(n) / \log(\log(n)))$ and sublinear search time $O(n^{1-\Theta(1/\log(\log(n)))})$ in high dimensional spaces.

Construction of SAT is as follows: an arbitrary object $a \in \mathbb{U}$ is chosen as the root node of the tree (note that since there exists only one object per node, we use both terms interchangeably in this section). Then, we select a suitable set of neighbors $N(a)$ such that $\forall u \in \mathbb{U}, u \in N(a) \Leftrightarrow \forall v \in N(a) - \{u\}, d(u, v) > d(u, a)$. Note that $N(a)$ is defined in terms of itself in a non-trivial way, and that multiple solutions fit the definition. In fact, finding the minimal set of neighbors seems to be a hard combinatorial optimization problem [12]. A simple heuristic that works well in most cases considers the objects in $\mathbb{U} - \{a\}$ in increasing order of their distance from a , and adds an object x to $N(a)$ if x is closer to a than to any object already in $N(a)$. Next, we put each node in $\mathbb{U} - N(a)$ into the bag of its closest element of $N(a)$. Also, for each subtree $u \in N(a)$ we store its covering radius $cr(u)$. The process is repeated recursively in each subtree using the elements of its bag. Figure 2 (left) shows an example of a SAT.

This construction process ensures that if we search for an object $q \in \mathbb{U}$ by spatial approximation, we will find that element in the tree because we are repeating exactly what happened during the construction process, i.e., we enter into the subtree of the neighbor closest to q , until we reach q (in fact, in this case we are doing an exact search because q is present in the tree). For general range queries (q, r) , instead of simply going to the closest neighbor, we first determine the closest neighbor c of q among $\{a\} \cup N(a)$. Then, we enter into all neighbors $b \in N(a)$ such that $d(q, b) \leq d(q, c) + 2r$. During the search process, all the nodes x such that $d(q, x) \leq r$ are reported. The search algorithm can be improved a bit more: when we search for an element $q \in \mathbb{U}$ (exact search), we follow a single path from the root to q . At any node a' in this path, we choose the closest to q among $\{a'\} \cup N(a')$. Therefore, if the search is currently at tree node a , we have that q is closer to a than to any ancestor a' of a and also any neighbor of a' . Hence, if we call $A(a)$ the set of ancestors of a (including a), we have that, at search time, we can avoid entering any element $x \in N(a)$ such that $d(q, x) > 2r + \min\{d(q, c), c \in \{a'\} \cup N(a'), a' \in A(a)\}$. This condition is a stricter version of the original Voronoi partition criterion. The covering radius stored for all nodes during the construction process can be used to prune the search further, by not entering into subtrees such that $d(q, b) - r > cr(b)$.

2.4 List of Clusters

The *List of Clusters* [4] is a list of “zones”. Each zone has a center and stores its covering radius. A center $c \in \mathbb{U}$ is chosen at random, as well as a radius rp , whose value depends on whether the number of elements per compact partition is fixed or not. The *center ball* of (c, rp) is defined as $(c, rp) = \{x \in \mathbb{X}, d(c, x) \leq rp\}$. We then define $I = \mathbb{U} \cap (c, rp)$ as the bucket of “internal” objects lying inside (c, rp) , and $E = \mathbb{U} - I$ as the rest of the elements (the “external” ones). The process is repeated recursively inside E . The construction process returns a list of triples (c_i, rp_i, I_i) (center, radius, internal bucket), as shown in Figure 2 (right).

This data structure is asymmetric, because the first center chosen has preference over the next centers in case of overlapping balls, as shown in Figure 2 (right). With respect to the value of the radius rp of each compact partition and the selection of the next center in the list, there exist many alternatives. In [4] it is shown experimentally that the best performance is achieved when the compact partition has a fixed number of elements, so rp becomes simply $cr(c)$, and the next center is selected as the element which maximizes the distance sum to the centers previously chosen. The brute force algorithm for constructing the list takes $O(n^2/m)$, where m is the size of the compact partition, but it can be improved using auxiliary data structures to build the partitions. For high dimensional metric spaces, the optimal m is very low (we used $m = 5$ in our experiments).

For a range query (q, r) , $d(q, c)$ is computed, reporting c if it is within the query ball. Then, we search exhaustively inside I only if $d(q, c) - cr(c) \leq r$ (covering radius criterion). E is processed only if $cr(c) - d(q, c) < r$, because of the asymmetry of the data structure. The search cost has a form close to $O(n^\alpha)$ for some $0.5 < \alpha < 1.0$ [4].

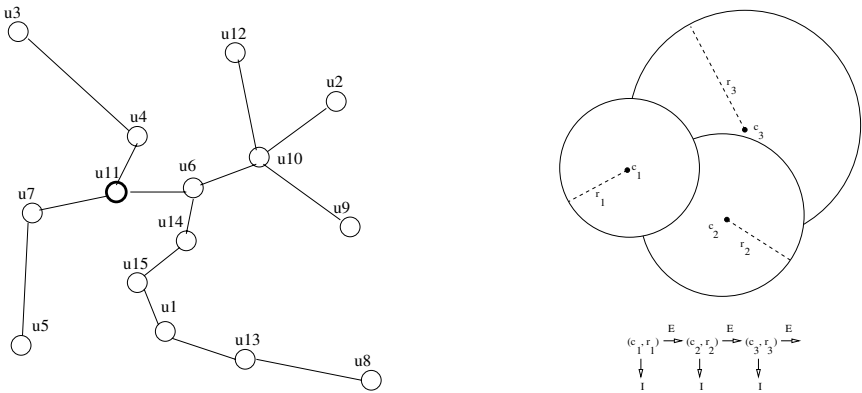


Fig. 2. Example of SAT (left) and List of Clusters (right)

3 Probabilistic Algorithms for Proximity Searching

All the algorithms seen in the previous section are *exact algorithms*, which retrieve exactly the elements of U that are within the query ball of (q, r) . In this work we are interested in *probabilistic algorithms*, which relax the condition of delivering the exact solution. As explained before, this is acceptable in most applications.

In [5] they present a probabilistic algorithm based on “stretching” the triangle inequality. The idea is general, but they applied it to pivot based algorithms. Their analysis shows that the net effect of the technique is to reduce the search radius by a factor β , and that that reduction is larger when the search problem becomes harder, i.e., the intrinsic dimension of the space becomes high. Even with very little stretching, they obtain large improvements in the search time with low error probability. The factor β can be chosen at search time, so the index can be built beforehand and later one can choose the desired level of accurateness and speed of the algorithm. As the factor is used only to discard elements, no element closer to q than r/β can be missed during the search. In practice, all the elements that satisfy $|d(p_i, u) - d(p_i, q)| > r/\beta$ for some p_i are discarded. Figure 3 illustrates how the idea operates. The exact algorithm guarantees that no relevant element is missed, while the probabilistic one stretches both sides of the ring and can miss some elements.

4 Our Approach

We focus in probabilistic algorithms for high dimensional metric spaces, where for exact searching it is very difficult to avoid the exhaustive search regardless of the index and search algorithm used.

It is well known that compact partition algorithms perform better than pivot-based algorithms in high dimensional metric spaces [6], and that the latter need

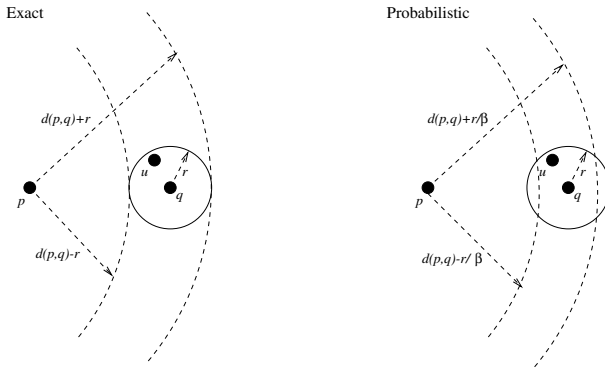


Fig. 3. How the probabilistic algorithm based on pivots works

more space requirements, i.e., many pivots, to reach the performance of the former. For this reason, it is interesting to develop probabilistic algorithms based on compact partitions, with the hope that these algorithms could have at least the same performance than pivot-based probabilistic algorithms, with less memory requirements. It is worth noting that the index data structure used with the probabilistic search algorithm is the same used with the exact search algorithm.

We propose two techniques: the first based on incremental searching and the last based on ranking zones.

4.1 Probabilistic Incremental Search

This technique is an adaptation of the *incremental nearest neighbor search* algorithm [10]. This incremental search traverses the search hierarchy defined by the index (whatever it be) in a “best-first” manner. At any step of the algorithm, it visits the “element” (zone or object) with the smallest distance from the query object among all unvisited elements in the search hierarchy. This can be done by maintaining a priority queue of elements organized by their maximum lower bound distance known to the query object at any time.

In [10] is proved that this search is *range-optimal*, that is, it obtains the k^{th} nearest neighbor, o_k , after visiting the same search hierarchy elements as would a range query with radius $d(q, o_k)$ implemented with a top-down traversal of the search hierarchy.

The incremental nearest neighbor search can be adapted to answer range queries. We report all objects u that satisfy $d(q, u) \leq r$, but we stop when it is dequeued an element with lower bound $l > r$ (*global stopping criterion*). It is not possible to find another object within the query ball among the unexplored elements, because we have retrieved them ordered by their lower bounded distances to q . An equivalent method is to enqueue elements only if they have a lower bound $l \leq r$, in which case the queue must be processed until it gets empty.

The idea of the probabilistic technique based on the incremental search is to fix in advance the number of distance computations allowed to answer a range query. Using the adapted incremental search for range queries, if the search is pruned after we make the maximum number of distance computations allowed, then we obtain a probabilistic algorithm in the sense that some relevant elements can be missed. However, as the search is performed range-optimally, one can presume that the allotted distance computations are used in an efficient way.

Figure 4 depicts the general form of the probabilistic incremental search. *Index* is the data structure that indexes \mathbb{U} , q is the query object, e is an element of the index and $d_{LB}(q, e)$ is a lower bound of the real distance between q and all the elements rooted in the search hierarchy of e , where $d_{LB}(q, e) = d(q, e)$ if e is an object of \mathbb{U} , and $d_{LB}(q, e) \geq d_{LB}(q, e')$ if e' is an ancestor of e in the hierarchy. For example, in the List of Clusters, if e is a child of a and belongs to the zone of center c then $d_{LB}(q, e) = \max(d(q, c) - cr(c), d_{LB}(q, a))$; in SAT if e is a child of a then $d_{LB}(q, e) = \max(d(q, a) - cr(a), (d(q, e) - \min\{d(q, c), c \in \{a'\} \cup N(a'), a' \in A(a)\})/2, d_{LB}(q, a))$. The maximum number of distance computations allowed

to perform the search is denoted by *quota*. Once *quota* has been reached, no more elements are enqueued. Note that the only stopping criterion of the algorithm is that the queue gets empty, even if the work quota has been reached, because for all the objects enqueued their distance to q are already known. The syntax of the enqueue procedure is `Enqueue(queue, element, lower bound distance)`. The dequeue procedure recovers the element e and its lower bound distance. Variable *cost* indicates the number of distance computations needed to process the children of element e in the search hierarchy. In SAT, *cost* is equal to $N(e)$; in List of Clusters, *cost* is equal to m .

4.2 Ranking of Zones

The probabilistic incremental search aims at quickly finding elements within the query ball, before the work quota gets exhausted. As the maximum number of distance computations is fixed, the total search time is also bounded. This technique can be generalized to what we call *ranking of zones*, where the idea is to sort the zones in order to favor the most promising and then to traverse the list until we use up the quota. The probabilistic incremental search can be seen as a ranking method, where we first rank all the zones using $d_{LB}(q, e)$ and then work until we use up the quota. However, this ranking does not have to be the best zone ranking criterion.

The sorting criterion must aim at quickly finding elements that are close to the query object. As the space is partitioned into zones, we must sort these zones in a promising search order using the information given by the index data structure. For example, in List of Clusters the only information we have is the distances from q to each center and the covering radius of each zone. One not

```

ProbIncrSearch( $q$ , Index, quota)

1. Queue  $\leftarrow \emptyset$  // Priority queue
2.  $e \leftarrow$  root of Index
3. counter  $\leftarrow 0$  // Number of distances computed
4. Enqueue(Queue,  $e$ , 0)
5. while not IsEmpty(Queue) do
6.   ( $e, d_{LB}(q, e)$ )  $\leftarrow$  Dequeue(Queue)
7.   if  $e$  is an object then report  $e$ 
8.   else
9.     cost  $\leftarrow$  cost to process children of  $e$ 
10.    if counter + cost  $\leq$  quota
11.      for each child element  $e'$  of  $e$  do
12.        Compute  $d_{LB}(q, e')$ 
13.        if  $d_{LB}(q, e') \leq r$  then
14.          Enqueue(Queue,  $e'$ ,  $\max(d_{LB}(q, e), d_{LB}(q, e'))$ )
15.        counter  $\leftarrow$  counter + cost

```

Fig. 4. Probabilistic incremental search algorithm

only would like to search first the zones closer to the query, but also to search first the zones that are more compact, that is, the zones which have “higher element density”. In spite of the fact that it is very difficult to define the volume of a zone in a general metric space, we assume that if the zones have the same number of elements, as in the best implementation of List of Clusters, then the zones with smaller covering radii have higher element density than those with larger covering radii.

We have tested several zone ranking criteria:

- the distance from q to each zone center, $d(q, c)$, closest first.
- the covering radius of each zone, $cr(c)$, in increasing order.
- $d(q, c) + cr(c)$, the distance from q to the farthest element in the zone.
- $d(q, c) \cdot cr(c)$.
- $d(q, c) - cr(c)$, the distance from q to the closest element in the zone.
- $\beta(d(q, c) - cr(c))$.

The first two techniques are the simplest ranking criteria. The next two techniques aim to search first in those zones that are closer to q and also are compact. The next technique, $d(q, c) - cr(c)$, is equivalent to the incremental search technique. The last technique is equivalent to reducing the search radius by a factor β as in [4], where $1/\beta \in [0..1]$.

If factor β is fixed, then this technique is equivalent to the probabilistic incremental search, because the ordering is the same in both cases. However, instead of using a constant factor $\beta \in [0..1]$, we can use a *dynamic factor* of the form $\beta = 1/(1.0 - \frac{cr(c)}{mcr})$, where mcr is the maximum size of the covering radius of all zones. This implies that we reduce more the search radius as the covering radius of a particular zone is greater. A special case is when $cr(c) = mcr$. In this case we define $d_{LB}(q, e) = \infty$ for all objects in the zone of center c .

Note that $d(q, c) - cr(c)$ is the only criterion that can be used with the incremental search technique, because only with this criterion is guaranteed that $d_{LB}(q, e) \geq d_{LB}(q, e')$ for any element e' ancestor of e .

5 Experimental Results

We use the SAT and List of Clusters to implement the probabilistic techniques described in Section 4, but with SAT we only implement the probabilistic incremental search because in this data structure every node is a center, so it takes $O(n)$ time to compute the distances between the query and every center. We have tested the probabilistic techniques on a synthetic set of random points in a k -dimensional vector space treated as a metric space, that is, we have not used the fact that the space has coordinates, but treated the points as abstract objects in an unknown metric space. The advantage of this choice is that it allows us to control the exact dimensionality we are working with, which is very difficult to do in general metric spaces. The points are uniformly distributed in the unitary cube, our tests use the L_2 (Euclidean) distance, the database size is $n = 10,000$

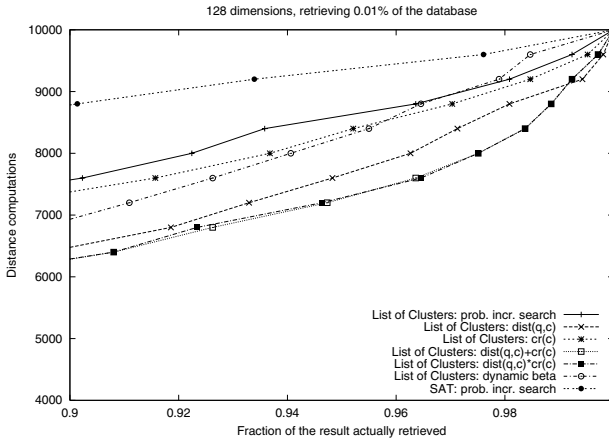


Fig. 5. Probabilistic List of Clusters and SAT in a vector space of dimension 128

and we perform range queries returning 0.01% of the total database size, taking an average from 1,000 queries. The techniques were tested using a space of dimension 128, where no known exact algorithm can avoid an exhaustive search to answer useful range queries.

Figure 5 shows the results of the probabilistic List of Clusters and SAT. The best technique, in this case, is the ranking zone method with criterion $d(q, c) + cr(c)$.

Figure 6 shows a comparison of the probabilistic List of Clusters and the probabilistic pivot-based algorithm, implemented in its canonical form (see Section 2.1 and 3). In this experiment, the probabilistic List of Clusters performs almost equal than the pivot-based algorithm with 256 pivots when more than 97% of the result is actually retrieved. The pivot-based techniques are slightly better when the pivots are selected using the “good pivots” criterion [3]. However, the size of the List of Clusters index (0.12 Mb) is about 82 times less than the size of the pivot-based index with 256 pivots (9.78 Mb) and about 5 times less than the size of the pivot-based index with 16 pivots (0.62 Mb). Experiments with different search radius and database size obtained similar results to those presented here.

One of the most clear applications of metric space techniques to Information Retrieval is the task of finding documents relevant to a query (which can be a set of terms or a whole document itself) [1]. Documents (and queries) are seen as vectors, where every term is a coordinate whose value is the weight of the term in that document. The distance between two documents is the angle between their vectors, so documents sharing important terms are seen as more similar. Documents closer to a query are considered to be more relevant to the query. Hence the task is to find the elements of this metric space of documents which are closest to a given query.

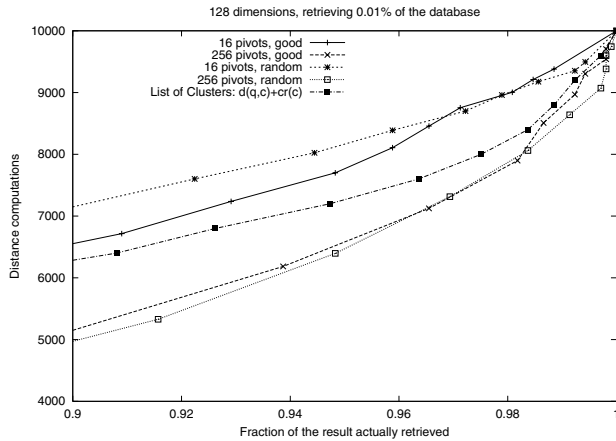


Fig. 6. Comparison among probabilistic algorithms in a vector space of dimension 128

Despite of this clear link, metric space techniques have seldom been used for this purpose. One reason is that the metric space of documents has a very high dimension, which makes any exact search approach unaffordable. This is a case where probabilistic algorithms would be of great value, since the definition of relevance is fuzzy and it is customary to permit approximations. Figure 7 shows a result on a subset of the TREC-3 collection [9], comparing the pivot-based algorithm with the ranking zone method using the dynamic beta criterion ($m = 10$ for the List of Clusters, retrieving on average 0.035% of the database per query). The result shows that our probabilistic algorithms can handle better this

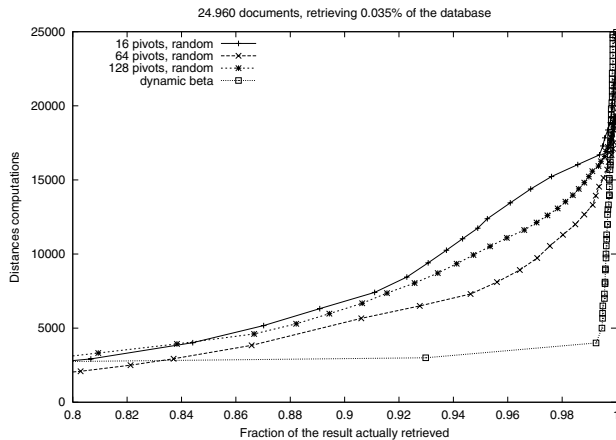


Fig. 7. Comparison among probabilistic algorithms in a document space

space, retrieving more than 99% of the relevant objects and traversing merely a 17% of the database, using much less memory, approximately 16 times less than the index with 64 pivots, hence becoming for the first time a feasible metric space approach to this long standing problem.

6 Conclusions

We have defined a general probabilistic technique based on the incremental nearest search, that allows us to perform time-bounded range search queries in metric spaces with a high probability of finding all the relevant elements. Our experimental results show in both synthetic and real-world examples that our technique performs better than the pivot-based probabilistic algorithm in high dimensional metric spaces, as the latter needs much more memory space to be competitive.

Future work involves testing more zone ranking criteria. Also, we are interested in finding a formal model that allows us to predict how well will perform an arbitrary index with our probabilistic techniques.

References

- [1] R. Baeza-Yates and B. Ribeiro-Neto. *Modern Information Retrieval*. Addison-Wesley, 1999. 294
- [2] S. Brin. Near neighbor search in large metric spaces. In *Proc. 21st Conference on Very Large Databases (VLDB'95)*, pages 574–584, 1995. 288
- [3] B. Bustos, G. Navarro, and E. Chávez. Pivot selection techniques for proximity searching in metric spaces. In *Proc. of the XXI Conference of the Chilean Computer Science Society (SCCC'01)*, pages 33–40. IEEE CS Press, 2001. 294
- [4] E. Chávez and G. Navarro. An effective clustering algorithm to index high dimensional metric spaces. In *Proc. 7th South American Symposium on String Processing and Information Retrieval (SPIRE'00)*, pages 75–86. IEEE CS Press, 2000. 287, 289, 293
- [5] E. Chávez and G. Navarro. A probabilistic spell for the curse of dimensionality. In *Proc. 3rd Workshop on Algorithm Engineering and Experiments (ALENEX'01)*, LNCS 2153, pages 147–160, 2001. 285, 290
- [6] E. Chávez, G. Navarro, R. Baeza-Yates, and J. Marroquín. Proximity searching in metric spaces. *ACM Computing Surveys*, 33(3):273–321, 2001. 284, 285, 286, 290
- [7] P. Ciaccia, M. Patella, and P. Zezula. M-tree: an efficient access method for similarity search in metric spaces. In *Proc. of the 23rd Conference on Very Large Databases (VLDB'97)*, pages 426–435, 1997. 287
- [8] F. Dehne and H. Noltemeier. Voronoi trees and clustering problems. *Information Systems*, 12(2):171–175, 1987. 287
- [9] D. Harman. Overview of the Third Text REtrieval Conference. In *Proc. Third Text REtrieval Conference (TREC-3)*, pages 1–19, 1995. NIST Special Publication 500-207. 295
- [10] G. Hjaltason and H. Samet. Incremental similarity search in multimedia databases. Technical Report TR 4199, Department of Computer Science, University of Maryland, November 2000. 291

- [11] I. Kalantari and G. McDonald. A data structure and an algorithm for the nearest point problem. *IEEE Transactions on Software Engineering*, 9(5):631–634, 1983. 287
- [12] G. Navarro. Searching in metric spaces by spatial approximation. *The Very Large Databases Journal (VLDBJ)*, 2002. To appear. Earlier version in SPIRE'99, IEEE CS Press. 287, 288
- [13] H. Noltemeier, K. Verbag, and C. Zirkelbach. Monotonous Bisector* Trees – a tool for efficient partitioning of complex schemes of geometric objects. In *Data Structures and Efficient Algorithms*, LNCS 594, pages 186–203. Springer-Verlag, 1992. 287
- [14] J. Uhlmann. Satisfying general proximity/similarity queries with metric trees. *Information Processing Letters*, 40:175–179, 1991. 287