# Approximate Nearest Neighbors: Towards Removing the Curse of Dimensionality 

Piotr Indyk* Rajeev Motwani ${ }^{\dagger}$<br>Department of Computer Science<br>Stanford University<br>Stanford, CA 94305

\{indyk,rajeev\}@cs.stanford.edu


#### Abstract

The nearest neighbor problem is the following: Given a set of $n$ points $P=\left\{p_{1}, \ldots, p_{n}\right\}$ in some metric space $X$, preprocess $P$ so as to efficiently answer queries which require finding the point in $P$ closest to a query point $q \in X$. We focus on the particularly interesting case of the $d$-dimensional Euclidean space where $X=\Re^{d}$ under some $l_{p}$ norm. Despite decades of effort, the current solutions are far from satisfactory; in fact, for large $d$, in theory or in practice, they provide little improvement over the brute-force algorithm which compares the query point to each data point. Of late, there has been some interest in the approximate nearest neighbors problem, which is: Find a point $p \in P$ that is an $\epsilon$-approximate nearest neighbor of the query $q$ in that for all $p^{\prime} \in P, d(p, q) \leq(1+\epsilon) d\left(p^{\prime}, q\right)$.


We present two algorithmic results for the approsimate version that significantly improve the known bounds: (a) preprocessing cost polynomial in $n$ and $d$, and a truly sublinear query time (for $\epsilon>1$ ); and, (b) query time polynomial in $\log n$ and $d$, and only a mildly exponential preprocessing cost $\tilde{O}(n) \times O(1 / \epsilon)^{d}$. Further, applying a classical geometric lemma on random projections (for which we give a simpler proof), we obtain the first known algorithm with polynomial preprocessing and query time polynomial in $d$ and $\log n$. Unfortunately, for small $\epsilon$, the latter is a purely theoretical result since the exponent depends on $1 / \epsilon$. Experimental results indicate that our first algorithm offers orders of magnitude improvement on running times over real data sets. Its key ingredient is the notion of locality-sensitive hashing which may be of independent interest; here, we give applications to information retrieval, pattern recognition, dynamic closest-pairs, and fast clustering algorithms.

[^0]
## 1 Introduction

The nearest neighbor search (NNS) problem is: Given a set of $n$ points $P=\left\{p_{1}, \ldots, p_{n}\right\}$ in a metric space $X$ with distance function $d$, preprocess $P$ so as to efficiently answer queries for finding the point in $P$ closest to a query point $q \in X$. We focus on the particularly interesting case of the $d$-dimensional Euclidean space where $X=\Re^{d}$ under some $l_{p}$ norm. The low-dimensional case is well-solved [2G], so the main issue is that of dealing with the "curse of dimensionality" [16]. The problem was originally posed in the 19603 by Minsky and Papert [53, pp. 222-225], and despite decades of effort the current solutions are far from satisfactory. In fact, for large $d$, in theory or in practice, they provide little improvement over a brute-force algorithm which compares a query $q$ to each $p \in P$. The known algorithms are of two types: (a) low preprocessing cost but query time linear in $n$ and $d$; and, (b) query time sublinear in $n$ and polynomial in $d$, but with severely exponential preprocessing cost $n^{d}$. This unfortunate situation carries over to average-case analysis, and even to the $\epsilon$-approximate nearest neighbors ( $\epsilon$-NNS) problem: Find a point $p \in P$ that is an $\epsilon$-approximate nearest neighbor of the query $q$, in that for all $p^{\prime} \in P, d(p, q) \leq(1+\epsilon) d\left(p^{\prime}, q\right)$.

We present two algorithms for the approximate version that significantly improve the known bounds: (a) preprocessing cost polynomial in $n$ and $d$, and a truly sublinear query time (for $\epsilon>1$ ); and, (b) query time polynomial in $\log n$ and $d$, and only a mildly exponential preprocessing cost $\tilde{O}(n) \times O(1 / \epsilon)^{d}$. Further, by applying a classical geometric lemma on random projections (for which we give a simpler proof), we obtain the first known algorithm with polynomial preprocessing and query time polynomial in $d$ and $\log n$. Unfortunately, for small $\epsilon$, this is a purely theoretical result as the exponent depends on $1 / \epsilon$. Experimental results [36] indicate that the first algorithm offers orders of magnitude improvement on running times over real data sets. Ito ley ingredient is the notion of locality-sensitive hashing which may be of independent interest; we give applications to information retrieval, pattern recognition, dynamic closest-pairs, and fast clustering.

Motivation. The nearest neighbors problem is of major importance to a variety of applications, usually involving
similarity searching. Some examples are: data compression [30]; databases and data mining [12, 38]; information retrieval [10, 20, 57]; image and video databases [28, 30, 65, 60]; machine learning [18]; pattern recognition [19, 25]; and, statistics and data analysis [21, 44]. Typically, the features of the objects of interest (documents, images, etc) are represented as points in $\Re^{d}$ and a distance metric is used to measure (dis)similarity of objects. The basic problem then is to perform indexing or similarity searching for query objects. The number of features (i.e., the dimensionality) ranger anywhere from tens to thousands. For example, in multimedia applications such as IBM's QBIC (Query by Image Content), the number of features could be several hundreds [28, 30]. In information retrieval for text documents, vector-space representations involve several thousands of dimensions, and it is considered to be a dramatic improvement that dimension-reduction techniques, such as LSI (latent semantic indexing) $[8,10,20]$, principal components analysin [39] or the Karhunen-Loéve transform [43, 49], can reduce the dimensionality to a mere few hundreds!

Of late, there has been an increasing interest in avoiding the curse of dimensionality by resorting to approximatenearest neighbor searching. Since the selection of features and the use of a distance metric in the applications are rather heuristic and merely an attempt to make mathematically precise what is after all an essentially aesthetic notion of similarity, it seems like an overkill to insist on the absolute nearest neighbor; in fact, determining an $\epsilon$-approximate nearest neighbor for a reasonable value of $\epsilon$, say a small constant, mhould suffice for most practical purposes. Unfortunately, even this relaxation of goals has not removed the curse of dimensionality, although the recent results of Kleinberg [45] gives some improvements.

Previous Work. Samet [58] surveys a variety of data structures for nearest neighbors including variants of $k$-d trees, $R$-trees, and structures based on space-filling curves; more recent results are surveyed in [59]. While some perform well in 2-3 dimensions, in high-dimensional spaces they all exhibit poor behavior in the worst case and in typical cases as well (c.g., sce Arya, Mount, and Narayan [3]). Dobkin and Lipton [22] were the first to provide an algorithm for nearest neighbors in $\Re^{d}$, with query time $O\left(2^{d} \log n\right)$ and preprocessing ${ }^{1}$ cost $O\left(\mathbf{n}^{d+1}\right)$. Clarkson [15] reduced the preprocessing to $O\left(n^{[d / 2](1+\delta)}\right)$, while increasing the query time to $O\left(2^{O(d \log d)} \log n\right)$. Later results, e.g., Agarwal and Matoušek [1], Matoušek [50], and Yao and Yao [64], all suffer from a query time that is exponential in d. Meiser [51] obtained query time $O\left(d^{5} \log n\right)$ but after $O\left(n^{d+\delta}\right)$ preprocessing. The so-called "vantage point" technique $[11,12,61,62]$ is a recently popular heuristic, but we are not aware of any analysis for high-dimensional Euclidean spaces. In gencral, even the average-case analysis of heuristics for points distributed over regions in $\Re^{d}$ gives an exponential query time $[0,34,58]$.

The situation is only slightly better for approximate nearest neighbors. Arya and Mount [2] gave an algorithm with

[^1]query time $O(1 / \epsilon)^{d} O(\log n)$ and preprocessing $O(1 / \epsilon)^{d} O(n)$. The dependence on $\epsilon$ was later reduced by Clarkson [16] and Chan [14] to $\epsilon^{-(d-1) / 2}$. Arya, Mount, Netanyahu, Silverman, and $\mathrm{Wu}[4]$ obtained optimal $O(n)$ preprocessing cost, but with query time growing as $O\left(d^{d}\right)$. Bern [7] and Chan [14] considered error $\epsilon$ polynomial in $d$ and managed to avoid exponential dependence in that case. Recently, Kleinberg [45] gave an algorithm with $O(n \log d)^{2 d}$ preprocessing and query time polynomial in $d, \epsilon$, and $\log n$, and another algorithm with preprocessing polynomial in $d, \epsilon$, and $n$ but with query time $O\left(n+d \log ^{3} n\right)$. The latter improves the $O(d n)$ time bound of the brute-force algorithm.

For the Hamming space $\{0,1\}^{d}$, Dolev, Harari, and Parnas [24] and Dolev, Harari, Linial, Nisan, and Parnas [23] gave algorithms for retrieving all points within distance $r$ of the query $q$. Unfortunately, for arbitrary $r$, these algorithms are exponential either in query time or preprocessing. Greene, Parnas, and Yao [37] present a scheme which, for binary data chosen uniformly at random, retrieves all points within distance $r$ of $q$ in time $O\left(d n^{r / d}\right)$, using $O\left(d n^{1+r / d}\right)$ preprocessing.

Very recently, Kushilevitz, Ostrovsky and Rabani [46] obtained a result similar to Proposition 3 below.

Overview of Results and Techniques. Our main results are algorithms ${ }^{2}$ for $\epsilon$-NNS described below. ${ }^{3}$
Proposition 1 For $\epsilon>1$, there is an algorithm for $\epsilon$-NNS in $\Re^{d}$ under the $l_{p}$ norm for $p \in[1,2]$ which uses $\tilde{O}\left(n^{1+1 / c}+\right.$ $d n)$ preprocessing and requires $\tilde{O}\left(d n^{1 / e}\right)$ query time.
Proposition 2 For $0<\epsilon<1$, there is an algorithm for $\epsilon-N N S$ in $\Re^{d}$ under any $l_{p}$ norm which uses $\tilde{O}(n) \times O(1 / \varepsilon)^{d}$ preprocessing and requires $\tilde{O}(d)$ query time.
Proposition 3 For any $\epsilon>0$, there is an algorithm for $\epsilon$-NNS in $\Re^{d}$ under the $l_{p}$ norm for $p \in[1,2]$ which uses $(n d)^{\circ(1)}$ preprocessing and requires $\tilde{O}(d)$ query time.

We obtain these results by reducing $\epsilon$-NNS to a new problem, viz., point location in equal balls. This is achieved by means of a novel data structure called ring-cover trees, described in Section 3. Our technique can be viewed as a variant of parametric search [52], in that they allow us to reduce an optimization problem to its decision version. The main difference is that in our case in answering a query we can only ask for a solution to a decision problem belonging to a prespecified set, since solving the decision problem (i.e., point location in equal balls) requires data structures created during preprocessing. We believe this technique will find further applications to problems where parametric search has been helpful.

In Section 4, we give two solutions to the point location problem. One is based on a method alin to the Elias bucketing algorithm [63] - we decompose each ball into a

[^2]bounded number of cells and store them in a dictionary. This allows us to achieve $\tilde{O}(d)$ query time, while the preprocessing is exponential in $d$, implying Proposition 2. For the second solution, we introduce the technique of localitysensitive hashing. The key idea is to use hash functions such that the probability of collision is much higher for objects that are close to each other than for those that are far apart. We prove that existence of such functions for any domain (not necessarily a metric space) implies the existence of fast $\epsilon$-NNS algorithms for that domain, with preprocessing cost only linear in $d$ and sublinear in $n$ (for $\epsilon>1$ ). We then present two families of such functions - one for a Hamming space and the other for a family of subsets of a set under the resemblance measure used by Broder et al [9] to cluster web documents. The algorithm based on the first family is used to obtain a nearest-neighbor algorithm for data sets from $\Re^{d}$, by embedding the points from $\Re^{d}$ onto a Hamming cube in a distance-preserving manner. The algorithm for the resemblance measure is shown to have several applications to information retrieval and pattern recognition. We also give additional applications of locality-sensitive hashing to dynamic closest-pair problem and fast clustering algorithms. All our algorithms based on this method are easy to implement and have other advantages - they exploit sparsity of data and the running times are much lower in practice [36] than predicted by theoretical analysis. We expect these results will have a significant practical impact.

An elegant technique for reducing complexity owing to dimensionality is to project the points into a random subspace of lower dimension, e.g., by projecting $P$ onto a small collection of random lines through the origin. Specifically, we could employ the result of Frankl and Maehara [32], which improves upon the Johnson-Lindenstrauss Lemma [41], showing that a projection of $P$ onto a subspace defined by roughly $9 \epsilon^{-2} \ln n$ random lines preserves all inter-point distances to within a relative error of $\epsilon$, with high probability. Applying this result to an algorithm with query time $O(1)^{d}$, we obtain an algorithm with query time $n^{9 \epsilon^{-2}}$. Unfortunately, this would lead to a sublinear query time only for large values of $\epsilon$. In Section A of the Appendix, we give a version of the random projection result using a much simpler proof than that of Frankl and Maehara. We also consider the extensions of the random projection approach to $l_{p}$ norms for $p \neq 2$. Using random projections and Proposition 2, we obtain the algorithm described in Proposition 3. Unfortunately, the high preprocessing cost (its exponent grows with $1 / \epsilon$ ) makes this algorithm impractical for small $\epsilon$.

## 2 Preliminaries

We use $l_{p}^{d}$ to denote the space $\Re^{d}$ under the $l_{p}$ norm. For any point $v \in \Re^{d}$, we denote by $\|\vec{v}\|_{p}$ the $l_{p}$ norm of the vector $\vec{v}$; we omit the subscript when $p=2$. Also, $H^{d}=$ $\left(\{0,1\}^{d}, d_{H}\right)$ will denote the Hamming metric space of dimension $d$. Let $\mathcal{M}=(X, d)$ be any metric space, $P \subset$ $X$, and $p \in X$. We will employ the following notation: $d(p, P)=\min _{q \in P} d(p, q)$, and $\Delta(P)=\max _{p, q \in P} d(p, q)$ is the diameter of $P$,

Definition 1 The ball of radius $r$ centered at $p$ is defined as $B(p, r)=\{q \in X \mid d(p, q) \leq r\}$. The ring $R\left(p, r_{1}, r_{2}\right)$ centered at $p$ is defined as $R\left(p, r_{1}, r_{2}\right)=B\left(p, r_{2}\right)-B\left(p, r_{1}\right)=$ $\left\{q \in X \mid r_{1} \leq d(p, q) \leq r_{2}\right\}$.

Let $V_{p}^{d}(r)$ denote the volume of a ball of radius $r$ in $l_{p}^{d}$. The following fact is standard [56, page 11].

Fact 1 Let $\Gamma($.$) denote the gamma function. Then V_{p}^{d}(r)=$ $\frac{(2 \Gamma(1+1 / p))^{d}}{\Gamma(1+n / p)} r^{d}$ and $V_{2}^{d}(r)=\frac{2 \pi^{d / 2}}{d \Gamma(d / 2)} r^{d}$.

## 3 Reduction to Point Location in Equal Balls

The key idea is to reduce the $\epsilon$-NNS to the following problems of point location in equal balls.
Definition 2 (Point Location in Equal Balls (PLEB)) Given $n$ radius-r balls centered at $C=\left\{c_{1}, \ldots, c_{n}\right\}$ in $\mathcal{M}=$ $(X, d)$, devise a data structure which for any query point $q \in X$ does the following: if there exists $c_{i} \in C$ such that $q \in B\left(c_{i}, r\right)$ then return $c_{i}$, else return NO.

Definition 3 ( $\epsilon$-Point Location in Equal Balls ( $\epsilon$-PLEB)) Given $n$ radius-r balls centered at $C=\left\{c_{1}, \ldots, c_{n}\right\}$ in $\mathcal{M}=$ $(X, d)$, devise a data structure which for any query point $q \in X$ does the following:

- if there exists $c_{i} \in C$ with $q \in B\left(c_{i}, r\right)$ then return YES and a point $c_{i}^{\prime}$ such that $q \in B\left(c_{i}^{\prime},(1+\epsilon) r\right)$,
- if $q \notin B\left(c_{i},(1+\epsilon) r\right)$ for all $c_{2} \in C$ then return No,
- if for the point $c_{i}$ closest to $q$ we have $r \leq d\left(q, c_{i}\right) \leq$ $((1+\epsilon) r)$ then return either YES or NO.
Observe that PLEB ( $\epsilon$-PLEB) can be reduced to NNS ( $\epsilon$-NNS), with the same preprocessing and query costs, as follows: it suffices to find an exact ( $\epsilon$-approximate) nearest neighbor and then compare its distance from $q$ with $r$. The main point of this section is to show that there is a reduction in reverse from $\epsilon$-NNS to $\epsilon$-PLEB, with only a small overhead in preprocessing and query costs. This reduction relies on a data structure called a ring-cover tree. This structure exploits the fact that for any point set $P$, we can either find a ring-separator or a cover. Either construct allows us to decompose $P$ into smaller sets $S_{1}, \ldots, S_{l}$ such that for all $i,\left|S_{i}\right| \leq c|P|$ for some $c<1$, and $\sum_{i}\left|S_{i}\right| \leq U|P|$ for $b<1+1 / \log ^{2} n$. This decomposition has the property that while searching $P$ it is possible to quickly restrict the search to one of the sets $S_{1}$.

There is a simpler but much weaker reduction from $\epsilon$ NN to $\epsilon$-PLEB. Let $R$ be the ratio of the smallest and the largest inter-point distances in $P$. For each $l \in\{1+c)^{0},(1+$ $\left.\epsilon)^{1}, \ldots, R\right\}$, generate a sequence of balls $B^{l}=\left\{B_{1}^{l}, \ldots, B_{n}^{l}\right\}$ of radius $l$ centered at $p_{1}, \ldots, p_{n}$. Each sequence $B^{l}$ forms an instance of PLEB. Then, given query $q$, we find via binary search the minimal $l$ for which there exists an $i$ such that $q \in B_{i}^{l}$ and return $p_{i}$ as an approximate nearest neighbor. The overall reduction parameters are: query time overhead factor $O(\log \log R)$ and space overhead factor $O(\log R)$. The simplicity of this reduction is very useful in practice. On the
other hand, the $O(\log R)$ space overhead is unacceptable when $R$ is large; in general, $R$ may be unbounded. In the final version, we will show that by using a variation of this method, storage can be reduced to $O\left(n^{2} \log n\right)$, which still does not give the desired $O(1 / \varepsilon)^{\text {d }} O(n)$ bound.

Definition $4 A$ ring $R\left(p, r_{1}, r_{2}\right)$ is an $\left(\alpha_{1}, \alpha_{2}, \beta\right)$-ring separator for $P$ if $\left|P \cap B\left(p, r_{1}\right)\right| \geq \alpha_{1}|P|$ and $\left|P \backslash B\left(p, r_{2}\right)\right| \geq$ $\alpha_{2}|P|$, where $r_{2} / r_{1}=\beta$.

Deflnition 5 A set $S \subset P$ is a $(\gamma, \delta)$-cluster for $P$ if for cvery $p \in S,|P \cap B(p, \gamma \Delta(S))| \leq \delta|P|$.
Dofinition $6 A$ sequence $A_{1}, \ldots, A_{l}$ of sets $A_{i} \subset P$ is called $a(b, c, d)$-cover for $S \subset P$, if there exists an $r \geq d \Delta(A)$ for $A=U_{i} A_{i}$ such that $S \subset A$ and for $i=1, \ldots, l$,

- $\left|P \cap\left(\cup_{p \in A_{i}} B(p, r)\right)\right| \leq b\left|A_{i}\right|$,
- $\left|A_{i}\right| \leq c|P|$.

Thoorem 1 For any $P, 0<\alpha<1$, and $\beta>1$, one of the following two properties must hold:

1. $P$ has an $(\alpha, \alpha, \beta)$-ring separator, or
2. $P$ contains $a\left(\frac{1}{2 \beta}, \alpha\right)$-cluster of size at least $(1-2 \alpha)|P|$.

Proof Sketch: First note that for $\alpha>1 / 2$, property (1) must be false but then property (2) is trivially true. In general, assume that (1) does not hold. Then, for any point $p$ and radius $r$ define:

- $f_{r}^{\infty}(r)=|P-B(p, \beta r)|$,
- $f_{p}^{0}(r)=|P \cap B(p, r)|$.

Clearly, $f_{p}^{\infty}(0)=n, f_{p}^{\infty}(\infty)=0, f_{p}^{0}(0)=0$, and $f_{p}^{0}(\infty)=n$. Also, notice that $f_{p}^{\infty}(r)$ is monotonically decreasing and $f_{p}^{0}(r)$ is monotonically increasing. It follows that there must exint a choice of $r$ (say $r_{p}$ ) such that $f_{p}^{\infty}\left(r_{p}\right)=f_{p}^{0}\left(r_{p}\right)$. Since (1) does not hold, for any value of $r$ we must have $\min \left(f_{p}^{\infty}(r), f_{p}^{0}(r)\right) \leq \alpha n$, which implies that $f_{p}^{\infty}\left(r_{p}\right)=f_{p}^{0}\left(r_{p}\right) \leq$ $\alpha n$.

Let $q$ be a point such that $r_{q}$ is minimal. Define $S=$ $P \cap R\left(g, r_{q}, \beta r_{q}\right) ;$ it follows that $|S| \geq(1-2 \alpha) n$. Also, notice that for any $s, s^{\prime} \in S, d\left(s, s^{\prime}\right) \leq 2 \beta r_{q}$, implying that $\Delta(S) \leq 2 \beta r_{q}$. Finally, for any $s \in S,\left|P \cap B\left(s, r_{q}\right)\right| \leq$ $\left|P \cap B\left(s, r_{0}\right)\right| \leq \alpha n$.

Theorem 2 Let $S$ be a $(\gamma, \delta)$-cluster for $P$. Then for any $b$, there is an algorithm which produces a sequence of sets $A_{1}, \ldots, A_{k} \subset P$ constituting $a\left(b, \delta, \frac{\gamma}{(1+\gamma) \log _{b} n}\right)$-cover for $S$.

## Proof Skotch:

The algorithm below greedily computes a good cover for $S$.

Algorithm Cover: $S=P \cap R\left(q, r_{q}, \beta r_{q}\right)$;
$r \leftarrow \frac{\pi \Delta(\mathcal{I})}{\log \iota n} ; j \leftarrow 0 ;$
ropeat

$$
\begin{aligned}
& j \leftarrow j+1 ; \text { choose some } p_{j} \in S ; B_{j}^{1} \leftarrow\left\{p_{j}\right\} ; \\
& i \leftarrow 1 ;
\end{aligned}
$$

while $\left|P \cap \cup_{q \in B_{j}^{i}} B(q, r)\right|>b\left|B_{j}^{i}\right|$ do $B_{j}^{i+1} \leftarrow P \cap \cup_{q \in B_{j}^{i}} B(q, r) ;$
$i \leftarrow i+1$
endwhile;
$A_{j} \leftarrow B_{j}^{i} ; S \leftarrow S-A_{j} ; P \leftarrow P-A_{j}$
until $S=\phi$;
$k \leftarrow j$.
In order to prove the correctness of the algorithm, it suffices to make the following four claims.

- $S \subset A=U_{j} A_{j}$ - Follows from the termination condition of the outer loop.
- for all $j \in\{1, \ldots, k\}$ and any $p \in S,\left|P \cap \cup_{q \in A_{j}} B(p, r)\right| \leq$ $b\left|A_{j}\right|$ - Follows from the termination condition of the inner loop.
- for all $j \in\{1, \ldots, k\},\left|A_{j}\right| \leq \delta|P|$ - Clearly, for any $j$, the inner loop is repeated at most $\log _{b} n$ times. Hence, $\max _{q \in A_{j}} d\left(p_{j}, q\right) \leq r \log _{b} n \leq \gamma \Delta(S)$. As $S$ is a $(\gamma, \delta)$-cluster, we have that $\left|B\left(p_{,}, \gamma \Delta(S)\right) \cap P\right| \leq \delta|P|$. Hence, $\left|A_{j}\right| \leq \delta|P|$.
- $r \leq \frac{\gamma \Delta(S)}{(1+\gamma) \log _{b} n}$ - Since $\Delta(A) \leq \Delta(S)+r \log _{b} n=$ $\Delta(S)+\gamma \Delta(S)=(1+\gamma) \Delta(S)$.

Corollary 1 For any $P, 0<\alpha<1, \beta>1, b>1$, one of the following properties must hold:

1. $P$ has an $(\alpha, \alpha, \beta)$-ring separator $R(p, r, \beta r)$, or
2. There is a $(b, \alpha, d)$-cover for some $S \subset P$ such that $|S| \geq(1-2 \alpha) n$ and $d=\frac{1}{(2 \beta+1) \log _{b} n}$.

### 3.1 Constructing Ring-Cover Trees

The construction of a ring-cover tree is recursive. For any given $P$ at the root, we use properties (1) and (2) in Corollary 1 to decompose $P$ into some smaller sets $S_{1}, \ldots, S_{l}$; these sets are assigned to the children of the node for $P$. Note the base case case is when $P$ is sufficiently small and we omit that in this abstract. We also store some additional information at the node for $P$ which enables us to restrict the nearest neighbor search to one of the children of $P$, by using distance computations or point lacation queries. For simplicity, assume that we can invoke an exact PLEB (not $\epsilon$-PLEB); the construction can be easily modified for approximate point location. There are two cases depending on which of the two properties (1) and (2) holds. Let $\beta=2\left(1+\frac{1}{6}\right), b=\frac{1}{\log ^{2} n}$, and $\alpha=\frac{1-1 / \log n}{2}$.

Case 1. In this case, we will call $P$ a ring node. We define its children to be $S_{1}=P \cap B(p, \beta r)$ and $S_{2}=P-$ $B(p, r)$. Also, we store the information about the ring separator $R$ at the node $P$.

Case 2. Here, we call $P$ a cover node. We define $S_{i}=$ $P \cap \cup_{p \in A_{i}} B(p, r)$ and $S_{0}=S-A$. The information stored at $P$ is as follows. Let $r_{0}=(1+1 / \epsilon) \Delta(A)$
and let $r_{i}=r_{0} /(1+\epsilon)^{i}$ for $i \in\{1, \ldots, k\}$, where $k=$ $\log _{1+\epsilon} \frac{(1+1 / c) \log _{b} n}{\gamma}+1$. Notice that $r_{k}=\frac{\gamma \Delta(A)}{\log _{p} p(1+\varepsilon)}=$ $\frac{r}{1+\varepsilon}$. For each $r_{i}$, generate an instance of PLEB with balls $B\left(p, r_{i}\right)$ for $p \in A$; all instances are stored at $P$.

We now describe how to efficiently search a ring-cover tree. It suffices to show that for any node $P$ we can restrict the search to one of its children using a small number of tests. Let $\min _{q}\left(p, p^{\prime}\right)$ denote the point out of $p$ and $p^{\prime}$ that is closer to $q$. The search procedure is as follows; we omit the obvious base case.

## Procedure Search:

1. if $P$ is a ring node with an ( $\alpha, \alpha, \beta$ )-ring separator $R(p, r, \beta r)$ then:
(a) if $q \in B(p, r(1+1 / \epsilon))$ then return $\operatorname{Search}\left(q, S_{1}\right)$;
(b) else compute $p^{\prime}=\operatorname{Search}\left(q, S_{2}\right)$; return $\min _{q}\left(p, p^{\prime}\right)$.
2. if $P$ is a cover node with a ( $b, c, d$ )-cover $A_{1}, \ldots, A_{l}$ of radius $r$ for $S \subset P$ then:
(a) if $q \notin B\left(a, r_{0}\right)$ then for all $a \in A$ then compute $p=\operatorname{Search}(q, P-A)$, choose any $a \in A$, and return $\min _{q}(p, a)$;
(b) else if $q \in B\left(a, r_{0}\right)$ for some $a \in A$ but $q \notin$ $B\left(a^{\prime}, r_{k}\right)$ for all $a^{\prime} \in A$ then using binary search on $r_{i} s$, find an $\epsilon$-NN $p$ of $q$ in $A$, compute $p^{\prime}=$ $\operatorname{Search}(q, P-A)$, and return $\min _{q}\left(p, p^{\prime}\right)$;
(c) else if $q \in B\left(a, r_{k}\right)$ for some $a \in A_{i}$ then return $\operatorname{Search}\left(q, S_{i}\right)$.

### 3.2 Analysis of Ring-Cover Trees

We begin the analysis of the ring-cover tree construction by establishing the validity of the search procedure.
Lemma 1 Procedure Search $(q, P)$ produces an $\epsilon$-nearest neighbor for $q$ in $P$.
Proof Sketch: Consider the two cases:

1. $P$ is a ring node.
(a) Consider any $s \in P-S_{1}$. Then $d(s, p) \leq d(s, q)+$ $d(q, p)$, implying that $d(s, q) \geq d(s, p)-d(q, p)$. Since $s \notin S_{1}$, we know that $d(s, p) \geq \beta r=2(1+$ $1 / \epsilon) r$, while $d(p, q) \leq r(1+1 / \epsilon)$. Then, $d(3, q) \geq$ $(1+1 / \epsilon) r \geq d(q, p)$.
(b) For any $s \in B(p, r), d(q, p) \leq d(q, s)+d(s, p)$, implying that $d(q, s) \geq d(q, p)-d(s, p) \geq d(q, p)-r$. It follows that $\frac{d(q, p)}{d(q, s)} \leq \frac{d(q, p)}{d(q, p)-r}=1+\frac{r}{d(q, p)-r} \leq$ $1+c$.
2. $P$ is a cover node.
(a) Similar to Case 1(b),
(b) Obvious.
(c) For any $p \in P-S_{i}, d(p, a) \geq r$. Since $q \in$ $B\left(a, r_{k}\right)$, we have $d(q, a) \leq r_{k}=\frac{r}{1+\epsilon} \leq \frac{d(p, q)}{1+\epsilon}$.

Lemma 2 The depth of a ring-cover tree is $O\left(\log _{1 / 2 a} n\right)=$ $O\left(\log ^{2} n\right)$.

Lemma 3 Procedure Search requires $O\left(\log ^{2} n \times \log l\right)$ distance computations or PLEB queries.

Lemma $4 A$ ring-cover tree requires space at most $O\left(k n b^{\log _{1 / 2 \alpha} n}(1+2(1-2 \alpha))^{\log _{n}}\right)=O($ npoly $\log n)$ not counting the additional non-data storage used by algorithms implementing PLEBs.

Proof Sketch: Let $S(n)$ be an upper bound on the space requirement for a ring-cover tree for point-set $P$ of size $n$. Then for a cover node:

$$
\begin{aligned}
S(n) \leq & \max _{l} A_{1} \ldots A_{i}, A_{i} \text { disjoint, }\left|A_{i}\right| \leq a n,|A| \geq(1-2 a) n \\
& {\left[\sum_{i=1}^{l} S\left(b\left|A_{i}\right|\right)\right]+S(n-|A|)+|A| b l }
\end{aligned}
$$

For a ring node:

$$
S(n) \leq 2 S\left(\frac{n}{2}(1+2(1-2 \alpha))\right)+1
$$

The bound follows by solving this recurrence.

Corollary 2 Given an algorithm for $P L E B$ which uscs $f(n)$ space on an instance of size $n$ where $f(n)$ is convex, a ringcover tree for an $n$-point set $P$ requires total space
$O(f($ npolylog $n))$.
Fact 2 For any PLEB instance ( $C, r$ ) generated by a ringcover tree, $\frac{\Delta(C)}{r}=O\left(\frac{1+\epsilon}{\gamma} \log _{b} n\right)$.

## 4 Point Location in Equal Balls

We present two techniques for solving the $\epsilon$-PLEB problem. The first is based on a method similar to the Elios bucketing algorithm [63] and works for any $l_{p}$ norm, establishing Proposition 2. The second uses locality-sensitive hashing and applies directly only to Hamming spaces (this bears some similarity to the indexing technique introduced by Greene, Parnas, and Yao [37] and the algorithm for allpairs vector intersection of Karp, Waarts, and Zweig [47], although the technical development is very different). However, by exploiting Facts 2 and 6 (Appendix A), the instances of $\epsilon$-PLEB generated while solving $\epsilon-$ NN for $l_{1}^{d}$ can be reduced to $\epsilon$-PLEB in $H^{m}$, where $m=d \log _{\iota} n \times \max (1 / \epsilon, \epsilon)$. Also, by Fact 5 (Appendix A), we can reduce $l_{p}^{d}$ to $l_{1}^{O(d)}$ for any $p \in[1,2]$. Hence, locality-sensitive hashing can be used for any $l_{p}$ norm where $p \in[1,2]$, establishing Proposition 1. It can also be used for the set resemblance measure used by Broder et al [9] to cluster web documents. We assume, without loss of generality, that all balls are of radius 1.

The proofs of Lemmas 2 and 3 are omitted.

### 4.1 The Bucketing Method

Annume for now that $p=2$. Impose a uniform grid of spacing $a=\epsilon / \sqrt{d}$ on $R^{d}$. Clearly, the distance between any two points belonging to one grid cuboid is at most $\epsilon$. By Fact 2, each side of the smallest cuboid containing balls from $C$ is of length at most $O\left(\sqrt{d} \log _{b} n \max (1 / \epsilon, \epsilon)\right)$ times the sidelength of a grid cell. For each ball $B_{i}$, define $\bar{B}_{i}$ to be the net of grid cells intersecting $B_{i}$. Store all elements from $U_{i} \bar{B}_{i}$ in a hash table [33,54], together with the information about the corresponding ball(s). (We can use hashing since by the preceding discussion the universe is of bounded size.) After preprocessing, to answer a query $q$ it suffices to compute the cell which contains $q$ and check if it is stored in the table.

We claim that for $0<\epsilon<1,|\bar{B}|=O(1 / \epsilon)^{d}$. To see this, observe that $|\vec{B}|$ is bounded by the volume of a $d$ dimensional ball of radius $r=2 / \epsilon \sqrt{d}$, which by Fact 1 is $\left.2^{O(d)}\right)^{d} / d^{d / 2} \leq O(1 / \epsilon)^{d}$. Hence, the total space required is $O(n) \times O(1 / \epsilon)^{d}$. The query time is the time to compute the hanh function. We use hash functions of the form:

$$
h\left(\left(x_{1}, \ldots, x_{d}\right)\right)=\left(a_{1} x_{1}+\ldots+a_{d} x_{d} \bmod P\right) \bmod M
$$

where $P$ is a prime, $M$ is the hash table size, and $a_{1}, \ldots, a_{d} \in$ $\mathcal{Z}_{p}^{*}$, This family gives a static dictionary with $O(1)$ access time [33]. The hash functions can be evaluated using $O(d)$ arithmetic operations. For general $l_{p}$ norms, we modify $s$ to $\epsilon / d^{1 / p}$. The bound on $|\bar{B}|$ applies unchanged.

Thoorem 3 For $0<\epsilon<1$, there is an algorithm for $\epsilon$ PLEB in $l_{p}^{d} u s i n g ~ O(n) \times O(1 / \epsilon)^{d}$ preprocessing and $O(1)$ ovaluations of a hash function for each query.

### 4.2 Locality-Sensitive Hashing

We introduce the notion of locality-sensitive hashing and apply it to sublinear-time similarity searching. The definition makes no assumptions about the object similarity measure. In fact, it is applicable to both similarity and dissimilarity measures; an example of the former is dot product, while any distance metric is an instance of the latter. To unify notation, we define a ball for a similarity measure $D$ as $B(q, r)=\{p ; D(q, p) \geq r\}$. We also generalize the notion of $c$-PLEB to ( $r_{1}, r_{2}$ )-PLEB where for any query point $q$ we require the answer to be YEs if $P \cap B\left(q, r_{2}\right) \neq \emptyset$ and No otherwise.

Deflnition 7 A family $\mathcal{H}=\{h: S \rightarrow U\}$ is called $\left(r_{1}, r_{2}, p_{1}, p_{2}\right)-$ nensitive for $D$ if for any $q, p, p^{\prime} \in S$

- if $p \in B\left(q, r_{1}\right)$ then $\operatorname{Pr}_{\mathcal{H}}[h(q)=h(p)] \geq p_{1}$,
- if $p \neq B\left(q, r_{2}\right)$ then $\operatorname{Pr}{ }_{H}\left[h(q)=h\left(p^{\prime}\right)\right] \leq p_{2}$.

In order for a locality-sensitive family to be useful, it has to aatisfy inequalities $p_{1}>p_{2}$ and $r_{1}<r_{2}$ when $D$ is a dissimilarity measure, or $p_{1}>p_{2}$ and $r_{1}>r_{2}$ when $D$ is a similarity measure.

For $l \mathrm{~s}$ specified later, define a function family $\mathcal{G}=\{g$ : $\left.S \rightarrow U^{k}\right\}$ such that $g(p)=\left(h_{1}(p), \ldots, h_{k}(p)\right)$, where $h_{i} \in \mathcal{H}$.

The algorithm is as follows. For an integer $l$ we choose $l$ functions $g_{1}, \ldots, g_{l}$ from $\mathcal{G}$ independently and uniformly at random. During preprocessing, we store each $p \in P$ in the bucket $g_{j}(p)$, for $j=1, \ldots, l$. Since the total number of buckets may be large, we retain only the non-empty buckets by resorting to hashing [ 33,54 ]. If any bucket contains more than one element, we retain an arbitrary one. To process a query $q$, we search all buckets $g_{1}(p), \ldots, g_{l}(p)$. Let $p_{1}, \ldots, p_{t}$ be the points encountered therein. For each $p$, if $p, \in$ $B\left(q, r_{2}\right)$ then we return YES and $p_{j}$, else we return No.

Let $W_{b}(q)=P-B(q, b)$, and $p^{*}$ be the point in $P$ closest to $q$. The parameters $k$ and $l$ are chosen so as to ensure that with a constant probability there exists $g_{j}$ such that the following properties hold:

1. $g_{j}\left(p^{\prime}\right) \neq g_{j}(q)$, for all $p^{\prime} \in W_{r_{3}}(q)$, and
2. if $p^{*} \in B\left(q, r_{1}\right)$ then $g_{j}\left(p^{*}\right)=g_{j}(q)$.

Lemma 5 If properties (1) and (2) hold for some $g_{3}$, the search procedure works correctly.

## Proof Sketch:

Case $1\left\{p^{*} \in B\left(q, r_{1}\right)\right\}$ : By property (1), the bucket $B=$ $g_{j}(q)$ cannot contain any points from $W_{r_{2}}$. By property (2), $p^{*}$ is contained in $B$. Therefore, $B$ is nonempty and contains only elements $p$ such that $D(q, p) \leq(1+$ $\epsilon) r$, and our algorithm will pick one such element and answer Yes.

Case $2\left\{p^{*} \notin B\left(q, r_{2}\right)\right\}$ : There are no points belonging to $B\left(q_{1} r_{2}\right)$, thus the algorithm answers No.

Theorem 4 Suppose there is $a\left(r_{1}, r_{2}, p_{1}, p_{2}\right)$-sensitive family $\mathcal{H}$ for $D$. Then there exists an algorithm for ( $r_{1}, r_{2}$ )PLEB under measure $D$ which uses $O\left(d n+n^{1+\rho}\right)$ space and $O\left(n^{\rho}\right)$ evaluations of the hash function for each query, where $\rho=-\frac{\ln p_{1}}{\ln p_{2} / p_{2}}$.
Proof Sketch: It suffices to ensure properties (1) and (2) for some $g_{j}$ with a constant probability. Assume that $p^{*} \in$ $B\left(q, r_{1}\right)$; the proof is similar when $p^{*} \notin B\left(q_{1}, r_{2}\right)$. Consider any point $p^{\prime} \in W_{r_{2}}(q)$. Clearly

$$
\begin{aligned}
& P_{1}=\operatorname{Pr}\left[g\left(p^{*}\right)=g(q)\right] \geq p_{1}^{k} \\
P_{2} & =\operatorname{Pr}\left[g\left(p^{\prime}\right)=g(q)\left[g\left(p^{*}\right)=g(q)\right]\right. \\
& =\frac{\operatorname{Pr}\left[g\left(p^{\prime}\right)=g(q) \wedge g\left(p^{*}\right)=g(q)\right]}{\operatorname{Pr}\left[g\left(p^{*}\right)=g(q)\right]} \\
& \leq \frac{\operatorname{Pr}\left[g\left(p^{\prime}\right)=g(q)\right]}{\operatorname{Pr}\left(\left[\left(p^{*}\right)=g(q)\right]\right.} \\
& \leq\left(\frac{p_{2}}{p_{1}}\right)^{*}
\end{aligned}
$$

Setting $k=\log _{\frac{p_{1}}{p_{2}}} 2 n$, we can bound $P_{2}$ by

$$
\left(\frac{p_{2}}{p_{1}}\right)^{\log \frac{p_{1}}{p_{2}} 2 n}=\frac{1}{2 n}
$$

Therefore, if $g_{j}$ satisfies property (1), then it also satisfies property (2) with probability at least

$$
1-\sum_{p^{\prime} \in W_{r_{2}}(q)} \frac{1}{2 n} \geq \frac{1}{2}
$$

It is sufficient to bound $P_{1}$ from below. By substituting for $k$ we obtain that

$$
P_{1} \geq p_{1}^{\log \frac{p_{1}}{p_{2}}}{ }^{n+1}=n^{\frac{p_{1}}{\log p_{1} / p_{2}}}=n^{-\rho}
$$

Choosing $l=n^{\rho}$ functions $g_{j}$, we ensure that with constant probability at least one function satisfies both properties (1) and (2).

We apply Theorem 4 to two measures: the Hamming metric and set resemblance [9]; the latter is a similarity measure defined for any pair of sets $A$ and $B$ as $D(A, B)=$ $|A \cap B|$. For the first measure, we apply a family of projections for fast hashing with $A C^{0}$ operations [5]. For the second measure, we use sketch functions used earlier [ 9 ] for estimation of the resemblance between given sets $A$ and $B$.

Proposition 4 ([5]) Let $S=\mathcal{H}^{d}$ and $D(p, q)$ be the Hamming metric for $p, q \in \mathcal{H}$. Then for any $r, \epsilon>0$, the family $\mathcal{H}=\left\{h_{i}: h_{i}\left(\left(b_{1}, \ldots b_{d}\right)\right)=b_{i}, i=1 \ldots n\right\}$ is
$\left(r, r(1+\epsilon), 1-\frac{r}{d}, 1-\frac{r(1+\epsilon)}{d}\right)$-sensitive.
Corollary 3 For any $\epsilon>1$, there exists an algorithm for $\epsilon$ $P L E B$ in $H^{d}$ (or, $l_{p}^{d}$ for any $p \in[1,2]$ ) using $O\left(d n+n^{1+1 / c}\right)$ space and $O\left(n^{1 / e}\right)$ hash function evaluations for each query. The hash function can be evaluated using $O(d)$ operations.

Proof Sketch: We use Proposition 4 and Theorem 4. First, we need to estimate the value of $\rho=-\frac{\ln p_{1}}{\ln p_{1} / p_{2}}$, where $p_{1}=$ $1-\frac{r}{d}$ and $p_{2}=1-\frac{r(1+e)}{d}$. Without loss of generality, we assume that $r<\frac{d}{\ln n}$, since we can increase dimensionality by adding a sufficiently long string of 0 s at the end of each point. Observe that

$$
\frac{p_{1}}{p_{2}}=\frac{1-r / d}{1-r(1+\epsilon) / d}>\frac{1}{1-\epsilon r / d}
$$

Thus,

$$
\rho=-\frac{\ln p_{1}}{\ln p_{1} / p_{2}}<-\frac{\ln (1-r / d)}{\ln \frac{1}{1-\epsilon r / d}}=\frac{\ln (1-r / d)}{\ln (1-\epsilon r / d)}
$$

Multiplying both the numerator and the denominator by $\frac{d}{r}$ we obtain that:

$$
\rho=\frac{\frac{d}{r} \ln (1-r / d)}{\frac{d}{r} \ln (1-\epsilon r / d)}=\frac{\ln (1-r / d)^{d / r}}{\ln (1-\epsilon r / d)^{d / r}}=\frac{U}{L}
$$

In order to upper bound $\rho$, we need to bound $U$ from below and $L$ from above; note that both $U$ and $L$ are negative. To this end we use the following inequalities [54]:

$$
(1-\epsilon r / d)^{d / r}<e^{-\epsilon} \quad \text { and } \quad(1-r / d)^{d / r}>e^{-1}\left(1-\frac{1}{d / r}\right)
$$

Therefore,

$$
\begin{aligned}
\frac{U}{L} & <\frac{\ln \left(e^{-1}\left(1-\frac{1}{d / r}\right)\right)}{\ln e^{-\epsilon}} \\
& =\frac{-1+\ln \left(1-\frac{1}{d / r}\right)}{-\epsilon} \\
& =1 / \epsilon-\frac{\ln \left(1-\frac{1}{d / r}\right)}{\epsilon} \\
& <1 / \epsilon-\ln (1-1 / \ln n)
\end{aligned}
$$

where the last step uses the assumptions that $\varepsilon>1$ and $r<\frac{d}{\ln n}$. We conclude that
$n^{\rho}<n^{1 / e} n^{-\ln (1-1 / \ln n)}=n^{1 / c}(1-1 / \ln n)^{-\ln n}=O\left(n^{1 / c}\right)$.

The hash function evaluation can be made faster than $O(d)$ for sparse data, i.e., when the number of non-zero coordinates of a query point is small. It suffices to sample the bits from the non-zero entries of the vectors; a similar method works for the functions used to build a static dictionary. Moreover, our experience is that the preprocessing space and query time are much lower than the above bound indicates. In particular, we have implemented a variant of the above data structure for the case when data is stored on disk [36]. For a data set of $20,000 \mathrm{~d}$-color histograms for images (with $d$ ranging up to 64) only 3-9 disk accesses were required in order to achieve small average error.

Proposition 5 ([9]) Let $S$ be the set of all subsets of $X=$ $\{1 \ldots x\}$ and let $D$ be the set resemblance measure. Then, for $1>r_{1}>r_{2}>0$, the following hash family is $\left(r_{1}, r_{2}, r_{1}, r_{2}\right)$ sensitive:

$$
\mathcal{H}=\left\{h_{\pi}: h_{\pi}(A)=\max _{a \in A} \pi(a), \pi \text { is a permutation of } X\right\}
$$

Corollary 4 For $0<\epsilon<r<1$, there exists an algorithm for $(r, \epsilon r)$-PLEB under set resemblance measure $D$ using $O\left(d n+n^{1+\rho}\right)$ space and $O\left(n^{\rho}\right)$ evaluations of the hash function for each query, where $\rho=-\frac{\ln r}{\ln e}$.

We now discuss further applications of the above corollary. For any pair of points $p, q \in \mathcal{H}^{d}$, consider the similarity measure $D(p, q)$ defined as the dot product $p \cdot q$. The dot product is a common measure used in information retrieval applications [31]; it is also of use in molecular clustering [13]. By using techniques by Indyk, Motwani, and Venkatasubramanian [40] it can also be used for solving the approximate largest common point set problem, which has many applications in image retrieval and pattern recognition. By a simple substitution of parameters, we can prove that for a set of binary vectors of approximately the same weight, PLEB under dot product measure (for queries of a fixed weight) can be reduced to PLEB under set resemblance measure. The fixed weight assumption can be easily satisfied by splitting the data points into $O(\log d)$ groups of approximately the same weight, and then making the same partition for weights of potential queries.

## 4．3 Further Applications of PLEB Algorithms

The PLEB procedures described above can also be used in casen where points are being inserted and deleted over time． In the randomized indexing method，insertion can be per－ formed by adding the point to all indices，and deletion can bo performed by deleting the point from all indices．In the bucketing method，insertion and deletion can be performed by adding or deleting all elements of $\bar{B}$ in the hash table． However，in order to apply these methods，we have to as－ sume that the points have integer coordinates with absolute value bounded by，say，$M$ ．Let $n$ be the maximum number of points present at any time．

Corollary 5 There is a data structure for $\epsilon-P L E B$ in $\{1 \ldots M\}^{d}$ which performs insertions，deletions，and queries in time $O(1 / \varepsilon)^{d}$ polv $(\log M, \log n)$ using storage $O(1 / \epsilon)^{d} n$ ．
Corollary 6 There is a data structure for $\epsilon$－PLEB in $\{1 \ldots M\}^{d}$ which performs insertions，deletions，and queries in time $O\left(M d n^{1 / c}\right)$ using storage $O\left(d n+n^{1+1 / c}\right)$ ．

By keeping several copies of PLEB as in the simple method described at the beginning of Section 3，we can answer ap－ proximate closest－pair queries．It is sufficient to check for every radius whether any cell（in the bucketing method）or any bucket（in the randomized indexing method）contains two different points；the smallest radius having this property gives an approximation to the closest－pair distance．The time bounds for all operations are as in the above corollar－ ien，but multiplied by a factor $O\left(\log _{\log }^{1+c} 1 M\right)$ ．

Combining both techniques，we obtain a method for dy－ namic estimation of closest pair．Eppstein［27］showed re－ cently that dynamic closest－pair problem has many applica－ tion to hierarchical agglomerative clustering，greedy match－ ing and other problems，and provided a data structure mak－ ing $O(n)$ distance computations per update operation．Our acheme gives an approximate answer in sublinear time．

## References

［1］P，K．Agarwal and J．Matoušek．Ray shooting and parametric search．In：Procecdings of the Twenty－ Fourth Annual ACM Symposium on Theory of Com－ puting，1992，pp．517－526．
［2］S．Arya and D．Mount．Approximate nearest neigh－ bor searching．In：Proceedings of the Fourth Annual ACM－SIAM Symposium on Discrete Algorithms，1993， pp．271－280．
［3］S．Arya，D．M．Mount，and O．Narayan，Accounting for boundary effects in nearest－neighbor searching．Dis－ crete and Computational Geometry，16（1996）：155－176．
［4］S．Arya，D．M．Mount，N．S．Netanyahu，R．Silverman， and A．Wu．An optimal algorithm for approximate nearest neighbor searching．In：Proceedings of the Fiflh Annual ACM－SIAM Symposium on Discrete Al－ gorithms，1994，pp．573－582．
［5］A．Andersson，P．B．Miltersen，S．Riis，M．Tho－ rup，Static dictionaries on $A C^{0}$ RAMs：Query time $\Theta(\sqrt{\log n / \log \log n})$ is necessary and sufficient．In： Proceedings of the 37th Annual IEEE Symposium on Foundations of Computer Science，1996，pp．441－450．
［6］J．L．Bentley．Multidimensional binary search trees used for associative searching．Communications of the ACM， 18（1975）：509－517．
［7］M．Bern．Approximate closest－point queries in high di－ mensions．Information Processing Letters，45（1993）：95－ 99.
［8］M．W．Berry，S．T．Dumais，and A．T．Shippy．A case study of latent semantic indexing．U．T．Knoxville Tech－ nical Report CS－95－271，January 1995.
［9］A．Broder，S．Glassman，M．Manasse，and G．Zweig． Syntactic clustering of the Web．In：Proceedings of the Sixth International 「＂orld 「ide 「eb Conference， pp．391－404， 1997.
［10］C．Buckley，A．Singhal，M．Mitra，and G．Salton．New Retrieval Approaches Using SMART：TREC 4．In：Pro－ ceedings of the Fourth Text Retrieval Conference，Na－ tional Institute of Standards and Technology， 1995.
［11］W．A．Burkhard and R．M．Keller．Some approaches to Best－Match File Searching．Communications of the ACM，16（1973）：230－236．
［12］T．Bozkaya and M．Ozsoyoglu．Distance－Based Index－ ing for High－Dimensional Metric Spaces，In：Proceed－ ings of the ACM SIGMOD International Conference on Management of Data（SIGMOD）， 1997.
［13］F．Cazals．Effective Nearest Neighbours Searching on the Hyper－Cube，with Applications to Molecular Clus－ tering．In Proceedings of the 14th Annual ACM Sym－ posium on Computational Geometry， 1998.
［14］T．M．Chan．Approximate Nearest Neighbor Queries Revisited．In：Proceedings of the 13 th Annual ACM Symposium on Computational Geometry，1997， pp．352－358．
［15］K．Clarkson．A randomized algorithm for closest－point queries．SIAM Journal on Computing，17（1988）：830－ 847.
［16］K．Clarkson．An algorithm for approximate closest－ point queries．In：Proceedings of the Tenth Annual ACM Symposium on Computational Geometry，1994， pp．160－164．
［17］K．Clarkson．Nearest Neighbor Queries in Metric Spaces．In：Proceedings of the Twenty－Ninth An－ nual ACM Symposium on Theory of Computing，1997， pp．609－617．
［18］S．Cost and S．Salzberg．A weighted nearest neighbor algorithm for learning with symbolic features．Machine Learning，10（1993）：57－67．
［19］T．M．Cover and P．E．Hart，Nearest neighbor pattern classification．IEEE Transactions on Information The－ ory，13（1967）：21－27．
［20］S．Deerwester，S．T．Dumais，T．K．Landauer，G．W．Fur－ nas，and R．A．Harshman．Indexing by latent semantic analysis．Journal of the Society fo4 Information Sci－ ences，41（1990）：391－407．
［21］L．Devroye and T．J．Wagner，Nearest neighbor methods in discrimination．In：Handbook of Statistics，vol．2， P．R．Krishnaiah and L．N．Kanal，eds．，North－Holland， 1982.
［22］D．Dobkin and R．Lipton．Multidimensional search problems．SIAM Journal on Computing，5（1976）：181－ 186.
［23］D．Dolev，Y．Harari，N．Linial，N．Nisan，and M．Par－ nas．Neighborhood preserving hashing and approx－ imate queries．In：Proceedings of the Fifth Annual ACM－SIAM Symposium on Discrete Algorithms，1994， pp．251－259．
[24] D. Dolev, Y. Harari, and M. Parnas. Finding the neighborhood of a query in a dictionary. In: Proceedings of the 2nd Israel Symposium on Theory and Computing Systems, 1993, pp. 33-42.
[25] R.O. Duda and P.E. Hart. Pattern Classification and Scene Analysis. John Wiley \&\% Sons, NY, 1973.
[26] H. Edelsbrumner. Algorithms in Combinatorial Geomctry. Springer-Verlag, 1987.
[27] D. Eppstein, Fast hierarchical clustering and other applications of dynamic closest pairs. In: Proceedings of the Ninth ACM-SLAM Symposium on Discrete Algorithms, 1998.
[28] C. Faloutsos, R. Barber, M. Flickner, W. Niblack, D. Petkovic, and W. Equitz. Efficient and effective querying by image content. Journal of Intelligent Information Systems, 3(1994):231-262.
[29] W. Feller. An Introduction to Probability Theory and its Applications. John Wiley \& Sons, NY, 1991.
[30] M. Flickner, H. Sawhney, W. Niblack, J. Ashley, Q. Huang, B. Dom, M. Gorkani, J. Hafner, D. Lee, D. Petkovic, D. Steele, and P. Yanker. Query by image and video content: the QBIC system. IEEE Computer, 28(1995):23-32.
[31] W. Frakes and R. Baeza-Yates, editors. Information Retrieval: Data Structures and Algorithms. PrenticeHall, 1992.
[32] P. Frankl and H. Maehara. The Johnson-Lindenstrauss Lemma and the Sphericity of Some Graphs. Journal of Combinatorial Theory B, 44(1988):355-362.
[33] M.L. Fredman, J. Komlós, and E. Szemerédi. Storing a sparse table with $\mathrm{O}(1)$ worst case access time. Journal of the ACM, 31(1984):538-544.
[34] J.K. Friedraan, J.L. Bentley, and R.A. Finkel. An algorithm for finding best matches in logarithmic expected time. ACM Transactions on Mathematical Software, 3(1977):209-226.
[35] A. Gersho and R.M. Gray. Vector Quantization and Data Compression. Kluwer, 1991.
[36] A. Gionis, P. Indyk, and R. Motwani. Similarity Search in High Dimensions via Hashing. Manuscript, 1997.
[37] D. Greene, M. Parnas, and F. Yao. Multi-index hashing for information retrieval. In: Proceedings of the 35th Annual IEEE Symposium on Foundations of Computer Scicnce, 1994, pp. 722-731.
[38] T. Hastie and R. Tibshirani. Discriminant adaptive nearest neighbor classification. In: First Internaîional Conference on Knowledge Discovery $\&$ Data Mining, 1995, pp. 142-149.
[39] H. Hotelling. Analysis of a complex of statistical variables into principal components. Journal of Educational Psychology, 27(1933):417-441.
[40] P. Indyk, R. Motwani, and S. Venkatasubramanian. Geometric Matching Under Noise - Combinatorial Bounds and Algorithms. Manuscript, 1997.
[41] W.B. Johnson and J. Lindenstrauss. Extensions of Lipshitz mapping into Hilbert space. Contemporary Mathcmatics, 26(1984):189-206.
[42] W.B. Johnson and G. Schechtman. Embedding $l_{p}^{m}$ into $l_{1}^{n}$. Acta Mathematica, 149(1982):71-85.
[43] K. Karhunen. Über lineare Methoden in der Walrscheinlichkeitsrechnung. Ann. Acad. Sci. Fennicae, Ser. A137, 1947.
[44] V. Koivure and S. Kassam. Nearest neighbor filters for multivariate data. IEEE 「örlshop on Nonlinear Signal and Image Processing, 1995.
[45] J. Kleinberg. Two Algorithms for Nearest-Neighbor Search in High Dimensions. In: Procecdings of the Twenty-Ninth Annual ACM Symposium on Thcory of Computing, 1997.
[46] E. Kushilevitz, R. Ostrovsky, and Y. Rabani. Efficient search for approximate nearest neighbor in high dimensional spaces. These proceedings.
[47] R.M. Karp, O. Waarts, and G. Zweig. The bit vector intersection problem. In: Procecdings of 36 th Annual IEEE Symposium on Foundations of Computer Science, 1995, pp. 621-630.
[48] N. Linial, E. London, and Y. Rabinovich. The geomctry of graphs and some of its algorithmic applicationa. In: Proceedings of 35th Annual IEEE Symposium on Foundations of Computer Science, 1994, pp. 577-591.
[49] M. Loéve. Fonctions aleastoires de second ordere. Processus Stochastiques et mouvement Brownian, Hermann, Paris, 1948.
[50] J. Matoušek. Reporting points in halfspaces. In: Computational Geometry: Theory and Applications, 2(1992):169-186.
[51] S. Meiser. Point location in arrangements of hyperplanes. Information and Computation, 106(1993):286303.
[52] N. Megiddo. Applying parallel computation algorithme in the design of serial algorithms. Journal of the ACM 31(1983), pp. 852-865.
[53] M. Minsky and S. Papert. Perceptrons. MIT Press, Cambridge, MA, 1969.
[54] R. Motwani and P. Raghavan. Randomized Algorithma. Cambridge University Press, 1995.
[55] A. Pentland, R.W. Picard, and S. Sclaroff. Photobook: tools for content-based manipulation of image databases. In Proceedings of the SPIE Confercince on Storage and Retrieval of Image and Vidco Datalases II, 1994.
[56] G. Pisier. The volume of convex bodies and Banach space geometry. Cambridge University Press, 1989.
[57] G. Salton and M.J. McGill. Introduction to Modern Information Retrieval. McGraw-Hill Book Company, New York, NY, 1983.
[58] H. Samet. The Design and Analysis of Spatial Data Structures. Addison-Wesley, Reading, MA, 1980.
[59] T. Sellis, N. Roussopoulos and C. Faloutsos. Multidimensional Access Methods: Trees Have Grown Everywhere. Proceedings of the 23rd International Conference on Very Large Data Bases (VLDB), 1997, pp, 13-15.
[60] A.W.M. Smeulders and R. Jain, eds. Image Databases and Multi-media Search. Proceedings of the First International Workshop, IDB-MMS '96, Amsterdam University Press, Amsterdam, 1996.
[61] J.K. Uhlmann. Satisfying General Proximity/Similarity Queries with Metric Trees. Information Processing Letters, 40(1991):175-179.
[62] P.N. Yiannilos. Data Structures and Algorithms for Nearest Neighbor Search in General Metric Spaces. In: Proceedings of the Second Annual ACM-SIAM Symposium on Discrete Algorithms, 1993, pp. 311-321.
[63] T. Welch. Bounds on the information retrieval efficiency of static file structures. Technical Report 88, MIT, June 1971.
[64] A.C. Yao and F.F. Yao, A general approach to $d$ dimensional geometric queries. In: Proceedings of the Seventeenth Annual ACM Symposium on Theory of Computing, 1985, pp. 163-168.

## A The Dimension Reduction Technique

We firat outline our proof for the random projections technique for dimension reduction. Combining this with Propoaition 2, we obtain the result given in Proposition 3.

Definition 8 Let $\mathcal{M}=(X, d)$ and $\mathcal{M}^{\prime}=\left(X^{\prime}, d^{\prime}\right)$ be two metric spaces, The space $\mathcal{M}$ is said to have a $c$-isometric ombedding, or simply a c-embedding, in $\mathcal{M}^{\prime}$ if there exists a map $f: \mathcal{M} \rightarrow \mathcal{M}^{\prime}$ such that

$$
(1-\epsilon) d(p, q)<d^{\prime}(f(p), f(q))<(1+\epsilon) d(p, q)
$$

for all $p_{1} q \in X$. $\overline{\text { Fe call }} \boldsymbol{c}$ the distortion of the embedding; if $c=1$, we call the embedding isometric.

Frankl and Maehara [32] gave the following improvement to the Johnson-Lindenstrauss Lemma [41] on ( $1+\epsilon$ )embedding of any $S \subset l_{2}^{d}$ in $l_{2}^{O(\log |S|)}$.

Lemma 0 (Frankl-Maehara [32]) For any $0<\epsilon<\frac{1}{2}$, any (sufficiently large) set $S$ of points in $\Re^{d}$, and $k=\left\lceil 9\left(\epsilon^{2}-\right.\right.$ $\left.2 \epsilon^{3} / 3\right)^{-1} \ln |S| \mid+1$, there exists a map $f: S \rightarrow \Re^{k}$ such that for all $u, v \in S$,

$$
(1-\epsilon)\|u-v\|^{2}<\|f(u)-f(v)\|^{2}<(1+\epsilon)\|u-v\|^{2}
$$

The proof proceeds by showing that the square of the length of a projection of any unit vector $v$ on a random $k$ dimensional hyperplane is sharply concentrated around $\frac{\mathrm{k}}{\mathrm{d}}$. Below we prove an analogous fact. However, thanks to the une of a different distribution, we are able to give a much simpler proof and also improve the constants. Note that the constants are important as they appear in the exponent of the time bounds of the resulting algorithm described in Proposition 3.

Lemma 7 Let $u$ be a unit vector in $\Re^{d}$. For any even positive integer $k$, let $U_{1}, \ldots, U_{k}$ be random vectors chosen independenlly from the d-dimensional Gaussian distribution ${ }^{4}$ $N^{d}(0,1)$. For $X_{i}=u \cdot U_{i}$, define $W=W(u)=\left(X_{1}, \ldots, X_{k}\right)$ and $L=L(u)=\|W\|^{2}$. Then, for any $\beta>1$,

$$
\begin{aligned}
& \text { 1. } E(L)=k \\
& \text { 2. } \operatorname{Pr}[L \geq \beta k]<O(k) \times \exp \left(-\frac{k}{2}(\beta-(1+\ln \beta))\right) \\
& \text { 3. } \operatorname{Pr}[L \leq k / \beta]<O(k) \times \exp \left(-\frac{k}{2}\left(\beta^{-1}-(1-\ln \beta)\right)\right) .
\end{aligned}
$$

Proof Sketch: By the spherical symmetry of $N^{d}(0,1)$ each $X_{i}$ is distributed as $N(0,1)$ [29, page 77]. Define $Y_{i}=$ $X_{2 i m 1}^{2}+X_{2 i}^{2}$, for $i=1, \ldots, k / 2$. Then, $Y_{i}$ follows the Exponential distribution with parameter $\lambda=\frac{1}{2}$ (see [29, page 47]), Thus $E(L)=\sum_{i=1}^{f / 2} E\left(Y_{i}\right)=(k / 2) \times 2=k$; also one can sec that $L$ follows the Gamma distribution with paramcters $\alpha=\frac{1}{2}$ and $v=k / 2$ (see [29, page 46]). Since this distribution is a dual of the Poisson distribution, we obtain that

$$
\operatorname{Pr}[L \geq \beta k]=\operatorname{Pr}\left[P_{\beta k}^{1 / 2} \leq v-1\right]
$$

[^3]where $P_{t}^{\alpha}$ is a random variable following the Poisson distribution with parameter $\alpha t$. Bounding the latter quantity is a matter of simple calculation.

An interesting question is if the Johnson-Lindenstrauss Lemma holds for other $l_{p}$ norms. A partial answer is provided by the following two results.

Theorem 5 For any $p \in[1,2]$, any $n$-point set $S \subset l_{p}^{d}$, and any $\epsilon>0$, there exist a map $f: S \rightarrow l_{2}^{k}$ with $k=O(\log n)$ such that for all $u, v \in S$.

$$
(1-\epsilon)\|u-v\|_{p}<\|f(u)-f(v)\|^{2}<(1+\epsilon)\|u-v\|_{p}
$$

Theorem 6 The Johnson-Lindenstrauss Lemma does not hold for $l_{\infty}$. More specifically, there is a set $S$ of $n$ points in $\Re^{d}$ for some $d$ such that any embedding of $S$ in $\Re^{\prime}$ has distortion $\Omega\left(\frac{\log n}{\sqrt{f}}\right)$.

Proof Sketch: We give a sketch of the proof of Theorem 6 based on the following two known facts.

Fact 3 (Linial, London, and Rabinovich [48]) Everynpoint metric $\mathcal{M}$ can be isometrically embedded in $l_{\infty}^{n}$.

Fact 4 (Linial, London, and Rabinovich [48]) There are graphs with $n$ vertices which for any d cannot be embedded in $l_{2}^{d}$ with distortion $0(\log n)$.

Assume for contradiction that the Johnson-Lindenstrauss Lemma holds for $l_{\infty}$ with distortion $t=o(\log n / \sqrt{f})$. Then for any graph $G$ with $n$ vertices: we embed $G$ in $l_{\infty}^{n}$ using Fact 3 ; by the assumption, we reduce the dimension to $f$ with distortion $t$; finally, we observe that as the norms $l_{\infty}^{f}$ and $l_{2}^{f}$ differ by at most a factor of $\sqrt{f}$, we have an embedding of $G$ in $l_{2}$ with distortion $t \sqrt{f}=o(\log n)$, which contradicts Fact 4.
Proof Sketch: We give a sketch of the proof of Theorem 5 based on the following two known facts.

Fact 5 (Johnson-Schechtman [42]) For any $1 \leq p<2$ and $\epsilon>0$, there exists a constant $\beta \geq 1$ such that for all $d \geq 1$, the space $l_{p}^{d}$ has $a(1+\epsilon)$-embedding in $l_{1}^{\beta d}$.

Fact 6 (Linial, London, and Rabinovich [48]) For any $\epsilon>0$ and every $n$-point metric space $\mathcal{M}=(X, d)$ induced by a set of $n$ points in $l_{1}^{d}$, there exists $m$ such that $\mathcal{M}$ has a $(1+\epsilon)$-embedding in $H^{m}$. If all points have coordinates from the set $\{1 \ldots R\}$, then $\mathcal{M}$ can be embedded isometrically in $H^{m}$ for $m=R d$.

The function $f$ is constructed implicitly by a sequence of reductions: Find an $\left(1+\epsilon_{1}\right)$-isometric embedding of $l^{p}$ into $l_{1}$ (using Fact 5) and let $S_{1}$ be the image of $S$ under this mapping. Find an ( $1+\epsilon_{2}$ )-isometric embedding of $S_{1}$ into $H^{m}$ (using Fact 6) and let $S_{2}$ be the image of $S_{1}$ under this mapping. Notice that for any $s, s^{\prime} \in S_{2}, d_{H}\left(s, s^{\prime}\right)=\| s-$ $s^{\prime} \|_{2}^{2}$, hence we may assume that $S_{2}$ resides in $l_{2}^{m}$. Finally, find an $\left(1+\epsilon_{3}\right)$-isometric embedding of $S_{2}$ into $l_{2}^{k}$ (using Lemma 6). It is now possible to choose suitable values for $\epsilon_{1}, \epsilon_{2}$, and $\epsilon_{3}$ to obtain the desired result.


[^0]:    *Supported by a Stanford Graduate Fellowship and NSF Award CCR-9357849.
    tSupported by a Sloan Fellowship, an IBM Faculty Fartnership Award, an ARO MURI Grant DAAH04-96-1-0007, and NSF Young Investigator Award CCR-9357849.
    Pemmiven to mahe digital or hard copiov of all or part of this work for prowal or ctassoom use is granted wathen fec proveded that copies we not made or dretributed lor profit or commercial athentage and that coper bear tha botace and the full catatom on the lint page To cops
    
    
    
    

[^1]:    ${ }^{1}$ Throughout, preprocessing cost refers to the space requirement; typically, the preprocessing time is roughly the same.

[^2]:    ${ }^{2}$ Our algorithms are randomized and return an approximate nearest neighbor with constant probability. To reduce the arror probability to $\alpha$, we can use several data structures in parallel and return the best result, increasing complexity by a factor $O(\log \alpha)$,
    ${ }^{3}$ For the sake of clarity, the $\tilde{O}$ notation is used to hide terms that are poly-logarithmic in $n$.

[^3]:    ${ }^{4}$ Each component is chosen independently from the standard normal distribution $N(0,1)$.

