# Approximate Line Nearest Neighbor in High Dimensions 

Alexandr Andoni MIT<br>andoni@mit.edu

Piotr Indyk<br>MIT<br>indyk@mit.edu

Robert Krauthgamer*<br>Weizmann Institute of Science<br>robert.krauthgamer@weizmann.ac.il

Huy L. Nguyen<br>MIT<br>hlnguyen@mit.edu


#### Abstract

We consider the problem of approximate nearest neighbors in high dimensions, when the queries are lines. In this problem, given $n$ points in $\mathbb{R}^{d}$, we want to construct a data structure to support efficiently the following queries: given a line $L$, report the point $p$ closest to $L$. This problem generalizes the more familiar nearest neighbor problem. From a practical perspective, lines, and low-dimensional flats in general, may model data under linear variation, such as physical objects under different lighting.

For approximation $1+\epsilon$, we achieve a query time of $d^{3} n^{0.5+t}$, for arbitrary small $t>0$, with a space of $d^{2} n^{O\left(1 / \epsilon^{2}+1 / t^{2}\right)}$. To the best of our knowledge, this is the first algorithm for this problem with polynomial space and sub-linear query time.


## 1 Introduction

Nearest neighbor problem is one of the fundamental geometric problems, with a wide range of applications. Given a set $S$ of $n$ points in a $d$-dimensional Euclidean space, the goal of the problem is to build a data structure which, given a query point, returns the data point closest to the query. Efficient algorithms for this problem exist if the data dimension is "low" Cla88, Mei93. In the high-dimensional case, efficient solutions exist as well, provided that the data structure is allowed to report an approximate nearest neighbor, whose distance to the query is at most some factor $c=1+\epsilon>1$ larger than the distance from the query to the actual nearest neighbor. In particular, several data structures with query time of $(d+\log n+1 / \epsilon)^{O(1)}$ using $n^{(1 / \epsilon)^{O(1)}}$ space are known IM98, KOR00, HP01].

When the query is more complex than a point, however, the complexity of the problem is much less understood. Perhaps the simplest such case is when the query is allowed to be a $k$-dimensional flat, for some small value of $k$; we call this the Approximate $k$-Flat Near-

[^0]est Neighbor problem. This is a natural generalization of the approximate nearest neighbor problem, where $k=0$. From a practical perspective, low-dimensional flats are often used to model data subject to linear variations. For example, one could capture the appearance of a physical object under different lighting or viewpoint [BHZ07.

It is therefore quite surprising that, for the case of high dimension $d$, there has been no prior work on this problem. For the "dual" version of the problem, where the query is a point but the data set consists of $k$-flats, two results are known BHZ07, Mag07 ${ }^{11}$. The first algorithm is essentially heuristic (although it allows some control of the quality of approximation). The second algorithm provides provable guarantees and very fast query time of $(d+\log n+1 / \epsilon)^{O(1)}$. Unfortunately, the space requirement of the algorithm is super-polynomial, of the form $2^{(\log n)^{O(1)}}$.

In this paper we provide the first non-trivial algorithm for the approximate flat nearest neighbor problem, for the case when the query is a 1 -flat, i.e., a line. Our data structure has query time sublinear in $n$ and uses polynomial space. Thus, it demonstrates that the flat version of the approximate nearest neighbor problem supports efficient solutions, even in high dimensions.

Specifically, for $1+\epsilon$ approximation, we obtain a query time of $O\left(d^{3} n^{0.5+t}\right)$ with a space of $d^{2} n^{O\left(1 / \epsilon^{2}+1 / t^{2}\right)}$ for any desired $t>0$.
1.1 Preliminaries and Definitions A line $\ell$ in $\mathbb{R}^{d}$ is described as $\ell=\{a+\lambda \cdot u \mid \lambda \in \mathbb{R}\}$, where $a$ is a vector in $\mathbb{R}^{d}$ and $u$ is a unit-length vector in $\mathbb{R}^{d}$.

For a point $p \in \mathbb{R}^{d}$ and a line $q$ in $\mathbb{R}^{d}$, we let $d(q, p)$ be the Euclidean distance from $p$ to $q$, which is the distance from $p$ to the projection of $p$ onto $q$.

We use the term cylinder, with axis $\ell$ and radius $R$, to denote the set $\{p \mid d(p, \ell) \leq R\}$. Any plane perpendicular to $\ell$ will be called base (or bottom) of

[^1]the cylinder. Furthermore, abusing notation, we will also use the notion of cylinder to denote the part of the (unbounded) cylinder contained between two hyperplanes perpendicular to $\ell$, or, in other words, a set $\left\{p \mid d(p, \ell) \leq R \wedge x_{1} \leq \operatorname{pr}_{\ell}(p) \leq x_{2}\right\}$ for some reals $x_{1}, x_{2} \in \mathbb{R}$ where $\operatorname{pr}_{\ell}(p)$ is the projection of $p$ into $\ell$. In such case, we say the height of the cylinder is $x_{2}-x_{1}$ (otherwise, the height is infinite).

For $c=1+\epsilon$, we define the problem of $c$ approximate line nearest neighbor (LNN) as follows. Given a dataset $S \subset \mathbb{R}^{d}$ of $n$ points, construct a data structure, which, given a query line $q$, reports a point $\tilde{p}$ such that $d(q, \tilde{p}) \leq c \cdot d(q, p)$ for all $p \in S$.

We use abbreviations LNN for approximate line nearest neighbor and NN for (standard) approximate nearest neighbor, instead of potentially more natural ALNN and ANN, to make the two acronyms easier to distinguish.

## 2 Approximate Line Nearest Neighbor

In this section, we describe our algorithm for solving the approximate line nearest neighbor. Our algorithm is mainly based on two data structures, treating two different cases. We call these data structures Cylinder-DS and Projection-DS , and we describe their implementation in later sections.

We first describe a high-level approach to the problem. For simplicity, suppose we are solving a simpler problem of approximate near-neighbor search ${ }^{[2]}$. In the latter problem, we are given a radius $r>0$ in advance, and the algorithm has to find a point $p$ at distance at most $(1+\epsilon) r$ from the query line $q$, if there exists a point $p^{*}$ at distance at most $r$ from $q$. For this problem, our algorithm is as follows. Let $t>0$ be a small constant. During preprocessing, we repeatedly find (and remove from the data set) cylinders of radius $r n^{t}$ that contain at least $m=\sqrt{n}$ points. We continue doing so until no such cylinder can be found anymore. For each of these cylinders we construct a "cylinder structure", Cylinder-DS . For the remaining points, we construct a "low-dimensional LNN structure", Projection-DS .

The query algorithm then queries each CylinderDS structure, as well as the Projection-DS structure. For each cylinder $C_{i}$, the algorithm tries to find a point $p \in C_{i}$ that is at distance at most $(1+\epsilon) r$ if a point at distance $\leq r$ exists. Each such query will take $\tilde{O}\left(n^{t}\right)$ time. Since there are at most $n / m=\sqrt{n}$ cylinders in total, we obtain a $\tilde{O}\left(n^{0.5+t}\right)$ time bound for this step.

[^2]The query algorithm also queries once the Projection-DS structure. This structure finds an (exact) $r$-near neighbor, but under the condition that there are at most $m=\sqrt{n}$ points at distance $r n^{t}$ from the query line. Note that this condition holds since we removed all cylinders of radius $r n^{t}$ that contain more than $m$ points. The Projection-DS structure has query time of $m \cdot \operatorname{poly}(d, \log n)=\tilde{O}(\sqrt{n})$ and the space requirement is $n^{O\left(1 / t^{2}\right)}$.

We now state precise conditions that we require of the structures Cylinder-DS and Projection-DS . We construct them in later sections.

Lemma 2.1. (Cylinder-DS ) There exists a data structure with the following guarantees. The input to the preprocessing stage is an m-point set $C \subset \mathbb{R}^{d}$, which is contained at distance at most $R$ from an axis line $\ell$. It is also given an approximation factor $1+\epsilon$ for $\epsilon>0$. The data structure supports the following query:

- for a query line $q$, if the closest point $p^{\star}$ is at distance at least $R / 2 n^{2 t}$, then the structure returns a point $p \in C$ at distance at most $(1+\epsilon) d\left(q, p^{\star}\right)$.
CYLINDER-DS has space and preprocessing time bound of $\tilde{O}\left(d^{2} m^{O\left(1 / \epsilon^{2}\right)}\right)$. The query time is $\tilde{O}\left(n^{2 t} d^{3}\right)$ and achieves a constant success probability.

Lemma 2.2. (Projection-DS ) There exists a data structure with the following guarantees. The input to the preprocessing stage is an $n$-point set $S \subset \mathbb{R}^{d}$ and an approximation factor $T=n^{t} / 2$ for constant $t \in(0,1 / 8]$. The data structure supports two types of queries:

- T-approximate line nearest neighbor: for a query line $q$, report a point $p \in S$ that is a $T$-approximate nearest neighbor;
- T-approximate line near neighbor reporting: for a query line $q$ and threshold $R$, report a set $W \subset S$ that includes all points at distance $R$ from $q$ but does not include any point $p \in S$ at distance more than $T \cdot R$ from $q$.

Projection-DS has space and preprocessing time bound of $\tilde{O}\left(n^{O\left(1 / t^{2}\right)}+d n\right)$. The query time is $d \log ^{O(1)} n$ for the query of the first type and $|W| \cdot d \log ^{O(1)} n$ for the query of the second type. Each query algorithm has a constant success probability.
2.1 Full Algorithm Our actual algorithm differs only slightly from the above description. First, we find cylinders as above by considering only cylinders whose axis passes through pairs of points in the dataset $S$. As we will show later, it is enough to consider such cylinders only.

The second difference is that we need to solve nearest neighbor problem and, thus, we do not have an estimate of $r$ in advance. Instead, we use a Projection-DS structure to find a $n^{t}$-approximate nearest neighbor $\tilde{p} \in S$. The distance from $\tilde{p}$ to the query line $q$, denoted $\tilde{r}$, is a very rough estimate for $r$. Additionally, we remove the dependence on $r$ in the construction of cylinder structures. To accomplish this, we construct cylinders as follows: pick the thinnest cylinder that contains $m$ points and remove it from the dataset; then pick the next thinnest cylinder and so forth. Finally, for each "suffix" of cylinders (thickest $x$ cylinders), we construct a Projection-DS structure on the points contained in these cylinders. During the query stage, we use the computed value $\tilde{r}$ to "zoom in" on the correct (prefix of) cylinders and the correct Projection-DS structure containing the rest of the points.

The resulting preprocessing and query algorithms are described in Algorithms 1 and 2 .

Algorithm 1: Preprocessing algorithm. The input is a set $S \subset \mathbb{R}^{d}$ and an approximation ratio $1+\epsilon$.

### 2.2 Correctness and Run-time Analysis

Correctness. We argue that the algorithm returns a $(1+\epsilon)$-approximate nearest neighbor. Let $p^{\star}$ be the closest point to $q$, and let $r^{\star}=d\left(q, p^{\star}\right)$. Then, by the guarantee of Projection-DS, we conclude that $\tilde{r}$ such that $\tilde{r} / n^{t} \leq r^{\star} \leq \tilde{r}$.

Suppose $p^{\star}$ is in a cylinder $C_{i}$ with $i \leq i^{*}-1$. Then $R_{i} \leq \tilde{r} \cdot n^{t}$ and thus $r^{\star} \geq \tilde{r} / n^{t} \geq R_{i} / n^{2 t}$. In this case the structure Cylinder-DS $C_{i}$ guarantees to return a $c$-approximate nearest neighbor.

Query-Processing ( $q$ ) :
1 Query Projection-DS structure $L_{0}$ to find an $n^{t}$-approximate nearest neighbor $\tilde{p}$
2 Let $\tilde{r}$ be the distance from $q$ to $\tilde{p}$
3 Let $i^{*}$ be the minimum $i \in[n / m]$ such that
$R_{i}>\tilde{r} n^{t}$ or set $i^{*}=n / m+1$ if no such $i$ exists
for $i \in\left\{1,2, \ldots, i^{*}-1\right\}$ do
Query Cylinder-DS structure $C_{i}$
end
7 Query Projection-DS structure $L_{i^{*}-1}$ to find all $\frac{n^{t}}{2}$-approximate line near neighbors for threshold $\tilde{r}$
8 Output the point nearest to $q$ from the ones found above
Algorithm 2: Query algorithm. The input is a line $q$ and the output is a point $p$ that is a $(1+\epsilon)$ approximate line nearest neighbor.

Now suppose $p^{\star}$ belongs to a cylinder $C_{i}$ for $i \geq i^{*}$. Then the structure Projection-DS $L_{i^{*}}$ is guaranteed to report the exact nearest neighbor $p^{\star}$.

Runtime. First we prove the following lemma.
Lemma 2.3. If there is a line $\ell$ within distance $R$ from a set $A$ of $m$ points then there is a line through two of the points that is within distance $2 R$ from all $m$ points.

Proof. Consider the cylinder of radius $R$ with axis at $\ell$, and, w.l.o.g., assume $\ell$ coincides with the $x$-axis. Take the point $a \in A$ with the smallest $x$-value and the point $b \in A$ with the largest $x$-value. Let line $\ell^{\prime}$ be the line passing through $a$ and $b$. We claim that any point $c \in A$ is within distance $2 R$ from $\ell^{\prime}$. Consider the point $z$ obtained through the intersection of the line $\ell^{\prime}$ and the plane through $c$ perpendicular to $x$ axis. This point $z$ is inside the cylinder, and thus, $d\left(c, \ell^{\prime}\right) \leq d(c, z) \leq d(c, \ell)+d(z, \ell) \leq 2 R$.

The first query to Projection-DS structure, and the queries to CYLINDER-DS structure, each run in time $\tilde{O}\left(n^{2 t}\right)$. Since there are at most $n / m=\sqrt{n}$ calls to Cylinder-DS structure, these calls take $\tilde{O}\left(d^{3} n^{.5+2 t}\right)$ time. It remains to argue that the last call to Projection-DS takes $\tilde{O}\left(d n^{.5+2 t}\right)$ time. Indeed, consider the structure Projection-DS $L_{i^{*}}$. By the construction of the cylinder $C_{i^{*}}$ and above lemma, we know that there is no line $\ell$ that contains more than $m$ points from $L_{i^{*}}$ within distance $R_{i^{*}} / 2 \geq \tilde{r} \cdot n^{t} / 2$. The guarantee of the Projection-DS structure for $\frac{n^{t}}{2}$ approximate reporting then says that the run-time of the query is at most $\tilde{O}(d m)=\tilde{O}(d \sqrt{n})$.

We note that the space is bounded by the space used by at most $O(\sqrt{n})$ Cylinder-DS and Projection-DS
structures, giving a total of $\tilde{O}\left(d^{2} n^{O\left(1 / \epsilon^{2}+1 / t^{2}\right)}\right)$ space.

## 3 Data Structure Cylinder-DS for Cylinder Case

In this section, we consider the problem of finding a $c$-approximate line nearest neighbor of a query line $q$ among $m$ points inside a cylinder $C_{1}$ of radius $R$ with axis $\ell$ assuming that the nearest neighbor $p^{\star}$ is at least $\frac{R}{2 n^{2 t}}$ away from $q$. We first describe the preprocessing and query answering algorithms in Section 3.1. We then prove their correctness and analyze the query time and space usage in Section 3.2. Finally, we show how the resulting algorithm yields Lemma 2.1.
3.1 The Algorithm In this section, we describe the algorithm for the cylinder case. At a high level, the algorithm works as follows. It first computes the intersection of the query line $q$ and a bigger cylinder containing $C_{1}$ and then extends the intersection region so as to include all $c$-LNNs in the extended region. This extended region is then divided into $\tilde{O}\left(n^{2 t}\right)$ smaller regions so that the part of $q$ in each region is "small" enough to be approximated by a single point with a small additive approximation error. For each smaller region, the algorithm queries a standard data structure for finding approximate nearest neighbors (NN) of points, to find a $c$-approximate nearest neighbor within that region. Finally, it reports the point closest to $q$ among the ones it finds. We choose to use a data structure for finding approximate nearest neighbor of query points with an efficient query time of $O\left(d^{2}\left(\log m+\frac{1}{\epsilon}\right)^{O(1)}\right)$ using $d^{2} m^{O\left(1 / \epsilon^{2}\right)}$ space (e.g., KOR00] ).

There are two additional technical problems that the algorithm needs to take care of. First, it is not known at preprocessing time where the intersection between $q$ and the cylinder is. This problem can be rectified as follows. Assume that the points are sorted by their projections on the cylinder axis and labeled $1,2, \ldots, m$ accordingly. A dyadic interval is an interval of the form $\left[i 2^{l}+1,(i+1) 2^{l}\right]$ for some $i, l$ satisfying $0 \leq l \leq \log m, 0 \leq i \leq\left\lceil m / 2^{l}\right\rceil-1$. It is wellknown that any interval can be divided into $O(\log m)$ dyadic intervals. At preprocessing time, for each dyadic interval, we build a separate data structure for points whose labels constitute that interval. At query time, the algorithm can just query $O(\log m)$ of them to get the answer. Second, there needs to be a special treatment when the nearest neighbor is very far from the query line. It turns out that this case can be reduced to finding the nearest neighbor in one-dimensional space, which can be solved easily.

See algorithms 3 and 4, as well as figures 1 and 2, for a detailed description.
3.2 Correctness and Running Time Analysis When $q$ is parallel to $\ell$, the distance between a point and $q$ is exactly the same as the distance between their projections on the subspace orthogonal to $\ell$ so the correctness follows from the correctness of the NN data structure. For the rest of the section, we consider the case when $q$ is not parallel to $\ell$.

Let $r^{\star}$ be the distance between $q$ and the nearest point $p^{\star} \in A$. Recall that we are only finding a $c$-LNN when $r^{\star}>\frac{R}{2 n^{2 t}}$. There are two cases of $r^{\star}$ we need to deal with. In section 3.2.1, we solve the problem when $r^{\star}$ is "small". In section 3.2.2, we solve the problem when $r^{\star}$ is "large". Finally, we give a concluding remark for the cylinder case in section 3.3 .
3.2.1 Small $r^{\star}$ In this section, we consider the case $r^{\star}<\frac{3 R}{\epsilon}$.

Lemma 3.1. Compute $E$ as in algorithm 4. We claim that all c-LNNs of $q$ are contained in $E$.

Proof. We prove by showing the contra-positive. Let $p$ be a point in $C_{1}$ but outside of $E$. The distance between $p$ and the part of $q$ outside of $C_{2}$ (i.e., set $q \backslash C_{2}$ ) is at least $\frac{7 R}{\epsilon}-R>\frac{3 R(1+\epsilon)}{\epsilon}>r^{\star}(1+\epsilon)$ for all $\epsilon<1$. The distance between $p$ and the part of $q$ inside $C_{2}$ is at least $\frac{3 R(1+\epsilon)}{\epsilon}>r^{\star}(1+\epsilon)$ because E was extended on both ends by $\frac{3 R(1+\epsilon)}{\epsilon}$. Thus, $p$ cannot be a $c$-LNN of $q$.

We distinguish two cases for the height of $E: E$ can be either long or short compared to the cylinder radius. Firstly, consider the case $E$ is short i.e. the height of $E$ is $M \leq \frac{252 R \log n}{\epsilon}$. We start by showing that the distance being used in steps 14 and 15 of algorithm 4 is a good approximation to the true distance. See Figure 1 for reference. Let $G$ be defined as in step 14 .

Lemma 3.2. Let $\epsilon$ be a constant smaller than 1. We claim that when $M \leq \frac{252 R \log n}{\epsilon}$, $a\left(1+\epsilon-\frac{1}{\log n}\right)-N N$ of an end point of some sub-segment from $G$ is a $(1+\epsilon)$ $L N N$ of $q$.

Proof. The radius of $C_{2}$ is $\frac{7 R}{\epsilon}$ so the segment of $q$ limited by two hyperplanes containing the two faces of $E$ has length at most $O\left(\frac{R(\log n+d)}{\epsilon}\right)$. This segment of $q$ is divided into $\Theta\left(\frac{d n^{2 t} \log ^{2} n}{\epsilon}\right)$ sub-segments so each subsegment has length at most $\frac{R}{6 n^{2 t} \log n}$. Therefore, using an end point as an approximation for any point on the same sub-segment only distorts distance by an additive term of at most $\frac{R}{6 n^{2 t} \log n}$. Since $r^{\star}>\frac{R}{2 n^{2 t}}$, the distance between any point $p$ in $S$ and the closest sub-segment end point is a $\left(1+\frac{1}{3 \log n}\right)$-approximation of the distance

```
CylinderDS-Preprocessing(A):
Sort the points in A by their projections on \ell and label them 1,2,\ldots,m accordingly
for D\in dyadic intervals do
    Consider the section of C1 that contains the points whose labels constitute D, and partition D into
    \frac{252\mp@subsup{n}{}{2t}\operatorname{log}n}{\epsilon}\mathrm{ sub-cylinders of equal heights}
4 For each sub-cylinder, project all the points in it onto the sub-cylinder's bottom face and build a data
    structure for finding NN of query points among the projections onto the bottom face
end
Build a standard NN data structure, denoted N Nall, on all points in }A\subset\mp@subsup{\mathbb{R}}{}{d
7 Project the set }A\mathrm{ onto the subspace orthogonal to }\ell\mathrm{ , and build a standard NN data structure, denoted
NN Northogonal},\mathrm{ on the resulting points (in a (d-1)-dimensional space)
```

Algorithm 3: Algorithm for building Cylinder-DS . Input is a set of points $A$.

```
CylinderDS-Query-Processing(q) :
if \(q\) and \(\ell\) are parallel then
    Query \(N N_{\text {orthogonal }}\) to find a \(\left(1+\epsilon-\frac{1}{\log n}\right)\)-NN of the projection of \(q\) (a point) in the subspace
    orthogonal to \(\ell\)
else
    Let \(C_{2}\) be an infinite cylinder with the same axis \(\ell\) as \(C_{1}\) and with radius \(\frac{7 R}{\epsilon}\). Compute the section of
    \(C_{2}\) containing the intersection of \(q\) and \(C_{2}\). Extend that section by \(\frac{3 R(1+\epsilon)}{\epsilon}\) on both ends and call the
    resulting extended cylinder \(E\). Let \(M\) be the height of cylinder \(E\)
    if \(M>\frac{252 R \log n}{\epsilon}\) then
            Use binary search to find the interval \(I\) of labels of points inside \(E\)
            Decompose \(I\) into at most \(O(\log m)\) dyadic intervals
            for \(D \in\) above dyadic intervals do
                for \(C \in\) sub-cylinders of the cylinder containing the points whose labels constitute \(D\) do
                        Query the NN structure of the bottom face of \(C\) to find a \(\left(1+\epsilon-\frac{1}{\log n}\right)\)-NN of the
                        intersection of \(q\) and the hyperplane containing the bottom face.
                end
            end
        else
            Divide the segment of \(q\) limited by two hyperplanes containing the two faces of \(E\) into \(\Theta\left(\frac{d n^{2 t} \log ^{2} n}{\epsilon}\right)\)
            sub-segments of equal length, and let \(G\) be the set of the resulting sub-segments
            Query \(N N_{\text {all }}\) to find a \(\left(1+\epsilon-\frac{1}{\log n}\right)\)-NN of each end point of each segment in \(G\)
        end
        Find the point \(p_{1}\) on \(\ell\) that is closest to \(q\). Use binary search to find the point in \(A\) whose projection
        on \(\ell\) is closest to \(p_{1}\)
        Return the point nearest to \(q\) among the points found above
end
```

Algorithm 4: Query algorithm for Cylinder-DS of a cylinder $C_{1}$ with axis $\ell$ and radius $R$. The input is a query line $q$ and the output is a $c$-LNN of $q$ among the points inside $C_{1}$.
between $p$ and $q$. Therefore, a $\left(1+\epsilon-\frac{1}{\log n}\right)$-NN of the end point closest to $p^{\star}$ is a $(1+\epsilon)$-LNN of $q$.

We can now show that steps 14 and 15 of algorithm 4 finds a $c$-LNN efficiently when $M \leq \frac{252 R \log n}{\epsilon}$.

Lemma 3.3. When $M \leq \frac{252 R \log n}{\epsilon}$, a $c-L N N$ of $q$ can
be found in $\tilde{O}\left(d^{3} n^{2 t}\right)$ time.
Proof. By lemma 3.2, one can find a $c$-LNN of $q$ by finding a $\left(1+\epsilon-\frac{1}{\log n}\right)$-NN of each end point of subsegments of the segment of $q$ limited by two hyperplanes containing the two faces of $E$ and returning the point closest to $q$. Finding each $\left(1+\epsilon-\frac{1}{\log n}\right)$-NN requires


Figure 1: Query algorithm for Cylinder-DS when $r^{\star}<\frac{3 R}{\epsilon}$ and $M \leq \frac{252 R \log n}{\epsilon}$. $C_{1}$ is the unbounded cylinder with axis $\ell$ and radius $R . C_{2}$ is the unbounded cylinder of radius $\frac{7 R}{\epsilon} . E$ is a cylinder of radius $\frac{7 R}{\epsilon}$ and height $M$. The dots on the query line $q$ are the endpoints of the segments in $G$ for which we query the $N N_{\text {all }}$ data structure.
one query to the $N N_{\text {all }}$ data structure, which takes $O\left(d^{2}\left(\log n+\frac{1}{\epsilon}\right)^{O(1)}\right)$ time. There are $\tilde{O}\left(d n^{2 t}\right)$ end points so the total running time is $\tilde{O}\left(d^{3} n^{2 t}\right)$.

Now consider the case $E$ is long i.e. $M>\frac{252 R \log n}{\epsilon}$. We first show the distance being used in step 10 of algorithm 4 is a good approximation of the true distance. See Figure 2 for reference.

Lemma 3.4. Consider the case $M>\underline{252 R \log n}$. As in algorithm 4, we project the points in $S^{\epsilon}$ in a subcylinder and the segment of $q$ limited by two hyperplanes containing the two faces of the sub-cylinder onto the sub-cylinder's bottom face. We claim that the distance between a point $p$ and an end point of the segment of $q$ in the projection space is a $\left(1+\frac{1}{3 \log n}\right)$-approximation of the distance between $p$ and $q$ in the original space.

Proof. Let $P$ be the hyperplane containing $p$ and orthogonal to $\ell$. Let $v$ be the vector from $p$ to the intersection of hyperplane $P$ and line $q$. Decompose $v$ into $v^{\|}$ and $v^{\perp}$, where $v^{\|}$is the component parallel to $q$ and $v^{\perp}$ is perpendicular to $q$. Note that $v^{\perp}$ is the distance between $p$ and $q$. Let $u$ and $i$ be the unit vectors on $\ell$ and $q$, respectively. Decompose $u$ into $u^{\|}$and $u^{\perp}$ in the same way as $v$. Let $\alpha$ be the angle between $u$ and $i$, where $0 \leq \alpha \leq \frac{\pi}{2}$. By the definition of $\alpha$, we have $u^{\|}=\cos \alpha \cdot i$, and $\left\|u^{\perp}\right\|_{2}=\sin \alpha$. Because $v$ is orthogonal to $\ell$, we have $0=u \cdot v=u^{\|} \cdot v^{\|}+u^{\perp} \cdot v^{\perp}$. By the Cauchy-Schwarz inequality, $\left|u^{\perp} \cdot v^{\perp}\right| \leq\left\|u^{\perp}\right\|_{2}\left\|v^{\perp}\right\|_{2}=\sin \alpha\left\|v^{\perp}\right\|_{2}$.

Therefore, $\cos \alpha\left\|v^{\|}\right\|_{2}=\left|u^{\|} \cdot v^{\|}\right| \leq \sin \alpha\left\|v^{\perp}\right\|_{2}$. In other words, $\frac{\|v\| \|_{2}}{\left\|v^{\perp}\right\|_{2}} \leq \tan \alpha$.

Since $M>\frac{252 R \log n}{\epsilon}$ and the diameter of the cylinder $C_{2}$ is $\frac{14 R}{\epsilon}$, we have that $\tan \alpha<\frac{1}{9 \log n}$. Thus, $\frac{\left\|v^{\|}\right\|_{2}}{\left\|v^{\perp}\right\|_{2}}<\frac{1}{9 \log n}$ and $\|v\|_{2}$ is a $1 \pm \frac{1}{9 \log n}$-approximation of $\left\|v^{\perp}\right\|_{2}$. The length of the segment of $q$ in the projection space is at most $\frac{1}{\frac{252 n^{2 t} \log n}{\epsilon}} \cdot \frac{14 R}{\epsilon}=\frac{R}{18 n^{2 t} \log n}<\frac{r^{\star}}{9 \log n}$. Therefore, using the projection of an end point of the segment of $q$ as an approximation for the projection of the intersection of $q$ and hyperplane $P$ further distorts distance by at most a factor of $1+\frac{1}{9 \log n}$. Combining the two approximation steps, we get the desired claim.

We can now show that step 10 of algorithm 4 can find a $c$-LNN of $q$ efficiently when $M>\frac{252 R \log n}{\epsilon}$.

Lemma 3.5. When $M>\frac{252 R \log n}{\epsilon}$, a $c-L N N$ of $q$ can be found in $\tilde{O}\left(d^{2} n^{2 t}\right)$ time.

Proof. By lemma 3.4, one can find a $c$-LNN of $q$ inside each sub-cylinder by finding a $\left(1+\epsilon-\frac{1}{\log n}\right)$-NN of the intersection of $q$ and the hyperplane containing the bottom face of the sub-cylinder. This can be done by making one query to the NN data structure, which takes $O\left(d^{2}\left(\log n+\frac{1}{\epsilon}\right)^{O(1)}\right)$ time. There are $\tilde{O}\left(n^{2 t}\right)$ subcylinders so the total running time is $\tilde{O}\left(d^{2} n^{2 t}\right)$.


Figure 2: Query algorithm for Cylinder-DS when $r^{\star}<\frac{3 R}{\epsilon}$ and $M>\frac{252 R \log n}{\epsilon}$. While we partition $E$ into up to $O(\log m)$ dyadic intervals, we illustrate just one such cylinder, $D$. The cylinder $D$ is further partitioned in equal height cylinders. For one such sub-cylinder $C$, the algorithm projects the points $S \cap C$ onto the bottom face of $C$, depicted as the hyperplane $F_{C}$, and constructs a nearest neighbor data structure on the projections. For example, the point $\operatorname{pr}(p)$ is the projection of a point $p \in S$ onto $F_{C}$. In Step 10, the algorithm queries this NN data structure for the query point $q \cap F_{C}$.
3.2.2 Large $r^{\star}$ In this section, we consider the case $r^{\star} \geq \frac{3 R}{\epsilon}$. We will show that in this case, step 17 of algorithm 4 finds a $c$-LNN of $q$ efficiently.

Lemma 3.6. If $r^{\star} \geq \frac{3 R}{\epsilon}$, a $c-L N N$ of $q$ can be found in $O(\log n)$ time.

Proof. The distance between a point and the cylinder axis is at most $R$ so replacing a point by its projection on $\ell$ only distorts distance by at most a factor of $1+\frac{\epsilon}{3}$. Assume that the parametric equation of $\ell$ and $q$ are $u_{\ell}+t_{\ell} v_{\ell}$ and $u_{q}+t_{q} v_{q}$, respectively, where $t_{\ell}, t_{q} \in \mathbb{R}$ and $u_{\ell}, v_{\ell}, u_{q}, v_{q}$ are vectors in $\mathbb{R}^{d}$. For any value of $t_{\ell}$, we can find the value of $t_{q}$ corresponding to the point on $q$ that is nearest to the point $u_{\ell}+t_{\ell} v_{\ell}$. The solution is $t_{q}=\frac{v_{q} \cdot\left(u_{\ell}+t_{\ell} v_{\ell}-u_{q}\right)}{\left\|v_{q}\right\|^{2}}$, which is linear in $t_{\ell}$. Thus, the square of the distance between $u_{\ell}+t_{\ell} v_{\ell}$ and $q$ is a quadratic function in $t_{\ell}$ and the minimum is achieved at some $t_{\ell}^{\star}$. Therefore, we can find the nearest neighbor of $q$ among points on $\ell$ by finding the point closest to the point $u_{\ell}+t_{\ell}^{\star} v_{\ell}$. This can be done by a simple binary search, as in step 17 .
3.3 Summary of the data structure for Cylinder-DS Combining two cases, we have an algorithm for $r^{\star}>\frac{R}{2 n^{2 t}}$. We can now finalize the proof of Lemma 2.1

Proof. [Proof of Lemma 2.1] This lemma follows directly from combining Lemmas 3.3, 3.5, and 3.6. The running time is dominated by the case $\frac{R}{2 n^{2 t}}<r^{\star}<\frac{3 R}{\epsilon}$ so the total running time is $\tilde{O}\left(d^{3} n^{2 t}\right)$. As described in algorithm 3, we have $\tilde{O}\left(n^{2 t}\right)$ NN data structures so the total space is $O\left(d^{2} n^{2 t+O\left(1 / \epsilon^{2}\right)}\right)$.

## 4 Data Structure Projection-DS

In this section, we describe the structure ProjectionDS , proving Lemma 2.2. First we give the algorithm and then prove its correctness.

The outline of the data structure is as follows. First we perform a random projection of the set $S$ into a space of constant dimension, namely dimension $k=2 / t+1$. Then, we solve the exact version of line nearest neighbor problem in the obtained $k$-dimensional space. For the second step, we follow the approach given by Magen in Mag07, Section 5].

For completeness, we describe Magen's approach for solving low-dimensional (exact) line nearest neighbor problem. (We note that Magen's approach addresses the dual version of this sub-problem, but it adapts to our case nonetheless.) The idea is to reduce the problem to a ray-shooting problem in some dimension $k^{\prime}=O\left(k^{2}\right)$. Namely, we transform each point $p \in$ $S$ into a hyperplane $H_{p} \subset \mathbb{R}^{k^{\prime}}$ such that the first
$k^{\prime}-1$ coordinates encode "all possible" lines and the last coordinate encodes the distance-squared from the corresponding line to the point $p$. The hyperplane $H_{p}$ is obtained as follows. Fix the point $p$ and consider a line $\ell=\{a+u \lambda \mid \lambda \in \mathbb{R}\}$ for some vector $a \in \mathbb{R}^{k}$ and unit-length vector $u \in \mathbb{R}^{k}$. Then, the distance-squared from $p$ to $\ell$ is equal to

$$
\begin{aligned}
d^{2}(p, \ell)= & \sum_{i}\left(a_{i}-p_{i}+u_{i} \sum_{j}\left(p_{j}-a_{j}\right) u_{j}\right)^{2} \\
= & \left(\sum_{j}\left(p_{j}-a_{j}\right) u_{j}\right)^{2}+\sum_{i}\left(a_{i}-p_{i}\right)^{2}+ \\
& +2 \sum_{i}\left(a_{i}-p_{i}\right) u_{i} \sum_{j}\left(p_{j}-a_{j}\right) u_{j}
\end{aligned}
$$

which is a degree- 6 polynomial in variables $a_{i}, u_{i}, i=$ $1,2, \ldots k$, with $O\left(k^{2}\right)$ monomials. We can linearize this polynomial by assigning a variable to each monomial, and thus view $d^{2}(p, \ell)$ as a linear function $f_{p}$ of $O\left(k^{2}\right)$ monomials (of degree at most 6 ) of $u_{i}, a_{i}$ 's. The hyperplane $H_{p}$ is simply defined by the equation $f_{p}\left(x_{1}, \ldots x_{k^{\prime}-1}\right)=x_{k^{\prime}}$. Now to find the closest point $p$ to a query line $\ell$ (given by $a, u \in \mathbb{R}^{k}$ ), we first compute the corresponding monomials to obtain a point $q_{\ell} \in \mathbb{R}^{k^{\prime}-1}$, and then shoot a ray from $\left(q_{\ell},-\infty\right)$ towards $\left(q_{\ell},+\infty\right)$. Then the first intersected hyperplane $H_{p}$ gives the closest point $p$ to $\ell$. Similarly, to obtain all $R$-near neighbors of $\ell$, we report all hyperplanes intersected by the semi-line from $\left(q_{\ell},-\infty\right)$ to $\left(q_{\ell}, R^{2}\right)$. To solve the ray-shooting problem in $\mathbb{R}^{k^{\prime}}$ itself, we use a data structure of Meiser Mei93 that gives $O\left(n^{k^{\prime}+1}\right)=O\left(n^{O\left(1 / t^{2}\right)}\right)$ space and $\operatorname{poly}(k, \log n)$ query time.

It remains to argue that the initial random projection step preserves the distances sufficiently well. We note that we do a dimensionality reduction into a very low dimensional space, and thus we will obtain some high distortions. Nonetheless, we show that this distortion is still bounded by $T=n^{t} / 2$, a factor polynomial in $n$.

The main part of the analysis for this case is thus the following lemma. We note that Magen Mag07 proves a similar lemma for small approximations, $c=$ $1+\epsilon$.

Lemma 4.1. Consider a point $p$ and a line $l$ in $R^{d}$. For $k \geq 5$, let $P \in M_{k, d}(\mathbb{R})$ be a random projection from $\mathbb{R}^{d}$ to $\mathbb{R}^{k}$, scaled by $\frac{1}{2} \sqrt{\frac{d}{k}}$. Let $p^{\prime}=P p$ be the projection of $p$ and line $l^{\prime} \subset \mathbb{R}^{k}$ is the projection of $l$ under $P$. Let $\Delta$ be the distance between $p$ and $q$ and $\Delta^{\prime}$ be the distance between $p^{\prime}$ and $l^{\prime}$. Then, for any $c>10$, we have that:

- $\operatorname{Pr}\left[\Delta^{\prime} \leq \Delta\right] \geq 1-e^{-k / 2}$, and
- $\operatorname{Pr}\left[\Delta / c \leq \Delta^{\prime}\right] \geq 1-(c / 15)^{-k / 2}$.

Before proving the lemma, we complete the correctness analysis of the algorithm. Indeed, the above lemma implies that:

- if $p^{*}$ is the distance to nearest point to $q$, then the distance to $p^{*}$ does not increase under the projection with constant probability;
- all points $p$ at distance $>T R$ from the line $q$ remain at distance at least $R$ in the projected space since each one remains at distance at least $R$ with probability $\geq 1-\left(n^{t} / 15\right)^{-k / 2} \geq 1-n^{-1-t / 4}$.

The above allows us to conclude Lemma 2.2. It only remains to prove Lemma 4.1.

Proof. [Proof of Lemma 4.1] To simplify the proof, we assume here that the random projection is scaled by $\sqrt{d / k}$ (without the $1 / 2$ factor). Then we need to prove that $\operatorname{Pr}\left[\Delta^{\prime} \leq 2 \Delta\right] \geq 1-e^{-k / 2}$ and $\operatorname{Pr}\left[\Delta^{\prime} \geq \Delta / c\right] \geq$ $1-(c / 7.5)^{-k / 2}$.

We use standard Johnson-Lindenstrauss Lemma to prove our claim. We note that we use JL lemma for high-distortion regime, as opposed to the more usual $1+\epsilon$ approximation regime. We use the following form of JL lemma (see, e.g., [DG99] or [IM98]):

Lemma 4.2. (JL Lemma) Let $P \in M_{k, d}$ be a random projection matrix, scaled by $\sqrt{d / k}$. Then, for every unit-length vector $x$, and every scalar $\beta \geq 2$, we have that:

- $\operatorname{Pr}\left[\|P x\|_{2} \geq \beta\right] \leq O(k) \cdot \exp [-k / 2 \cdot(\beta-(1+\ln \beta))]$ and
- $\operatorname{Pr}\left[\|P x\|_{2} \leq 1 / \beta\right] \leq O(k) \cdot \exp [-k / 2 \cdot(\ln \beta-1)]$.

We prove that $\Delta^{\prime} \leq 2 \Delta$ with the desired probability. Let $\hat{p}$ be the projection of $p$ onto $l$ (in $\mathbb{R}^{d}$ ), and let $\hat{p}^{\prime}=P \hat{p}$ be the projection of $\hat{p}$ under $P$. Then, by the above lemma, $d\left(p^{\prime}, l^{\prime}\right) \leq d\left(p^{\prime}, \hat{p}^{\prime}\right) \leq 2 d(p, l)$ with probability at least $1-e^{-\Omega(k)}$.

Now we prove that $\Delta^{\prime} / c \leq \Delta$ with the desired probability. By rescaling, we can assume, wlog, that $d(p, l)=1$. Take $q \in l$ to be a point on $l$ at distance $\frac{1}{c^{2}}$ from $\hat{p}$. Also, let $q^{\prime}=P q$. Note that $q^{\prime}$ and $\hat{p}^{\prime}$ completely define $l^{\prime}$ in $\mathbb{R}^{k}$. Finally, let $u=q-\hat{p}$ and $u^{\prime}=P u=q^{\prime}-\hat{p}^{\prime}$. Then, any point on $l^{\prime}$ is equal to $\hat{p}^{\prime}+\lambda u^{\prime}$ for some real $\lambda$. Note that $\left\|u^{\prime}\right\| \leq 1 / c$ with probability at least $1-e^{-k / 2(c-1-\ln c)} \geq 1-e^{-k c / 4}$.

The proof plan is to first consider a (infinite) number of dense, equally spaced points on the line $l^{\prime}$
and prove that none is close to $p^{\prime}$. Then we prove that $p^{\prime}$ cannot be close to a point in between the considered points on the line $l^{\prime}$.

In particular, for $i \in \mathbb{Z}$, let $q_{i}=\hat{p}+i u$ and $q_{i}^{\prime}=P q_{i}=p^{\prime}+i u^{\prime}$. By JL lemma from above, we have that $d\left(p^{\prime}, q_{0}^{\prime}\right)=d\left(\hat{p}^{\prime}, p^{\prime}\right) \geq 1 / c$ with probability at least $1-(e / c)^{k / 2}$. For $i \neq 0$, we have that $d\left(p, q_{i}\right) \geq$ $i$, and thus $d\left(p^{\prime}, q_{i}^{\prime}\right) \geq 1 / c$ with probability at least $1-\left(\frac{e}{c|i|}\right)^{k / 2}$. By union bound over all points $q_{i}^{\prime}$ for $i \in \mathbb{Z}$, we have that $d\left(p^{\prime}, q_{i}^{\prime}\right) \geq 2 / c$ with probability at least $1-(e / c)^{k / 2} \cdot\left(1+\sum_{i \in \mathbb{Z}} 1 /|i|^{k / 2}\right) \geq 1-5(e / c)^{k / 2}$.

Now suppose none of the "bad" (low-probability) events from above holds. Then for any point $w \in l^{\prime}$, we have that $d\left(p^{\prime}, w\right) \geq d\left(p^{\prime}, q_{i}^{\prime}\right)-d\left(q_{i}^{\prime}, w\right)$, where $i$ is such that $q_{i}^{\prime}$ is the closest to $w$. In such case, we have that $d\left(p^{\prime}, w\right) \geq 2 / c-1 / c=1 / c$. This concludes the proof of Lemma 4.1.

## References

[BHZ07] Ronen Basri, Tal Hassner, and Lihi Zelnik-Manor. Approximate nearest subspace search with applications to pattern recognition. In Computer Vision and Pattern Recognition (CVPR'07), pages 1-8, June 2007.
[Cla88] K. Clarkson. A randomized algorithm for closestpoint queries. SIAM Journal on Computing, 17:830847, 1988.
[DG99] S. Dasgupta and A. Gupta. An elementary proof of the johnson-lindenstrauss lemma. ICSI technical report TR-99-006, Berkeley, CA, 1999.
[HP01] S. Har-Peled. A replacement for voronoi diagrams of near linear size. In Proc. of FOCS, pages 94-103, 2001.
[IM98] P. Indyk and R. Motwani. Approximate nearest neighbor: towards removing the curse of dimensionality. Proc. of STOC, pages 604-613, 1998.
[KOR00] E. Kushilevitz, R. Ostrovsky, and Y. Rabani. Efficient search for approximate nearest neighbor in high dimensional spaces. SIAM J. Comput., 30(2):457474, 2000. Preliminary version appeared in STOC'98.
[Mag07] A. Magen. Dimensionality reductions in $\ell_{2}$ that preserve volumes and distance to affine spaces. Discrete and Computational Geometry, 38(1):139-153, July 2007. Preliminary version appeared in RANDOM'02.
[Mei93] S. Meiser. Point location in arrangements of hyperplanes. Information and Computation, 106:286-303, 1993.


[^0]:    *Part of this work was done while at IBM Almaden. This research was supported in part by The Israel Science Foundation (grant \#452/08).

[^1]:    ${ }^{1}$ It is plausible that both algorithms could be adapted to our problem without sacrificing their complexity.

[^2]:    ${ }^{2}$ It is known that if the query is a point, then the approximate near and nearest neighbor problems are equivalent up to polylogarithmic factors in the query time and space complexity IM98 HP01]. Unfortunately, the reduction does not generalize to higher-dimensional queries.

