

High-Dimensional Shape Fitting in Linear Time*

Sariel Har-Peled¹ and Kasturi R. Varadarajan²

¹Department of Computer Science, University of Illinois at Urbana-Champaign,
201 N. Goodwin Avenue, Urbana, IL 61801-2302, USA
sariel@uiuc.edu
<http://www.uiuc.edu/~sariel/>

²Department of Computer Science, University of Iowa,
Iowa City, IA 52242-1419, USA
kvaradar@cs.uiowa.edu

Abstract. Let P be a set of n points in \mathbb{R}^d . The *radius* of a k -dimensional flat \mathcal{F} with respect to P , which we denote by $\mathcal{RD}(\mathcal{F}, P)$, is defined to be $\max_{p \in P} \text{dist}(\mathcal{F}, p)$, where $\text{dist}(\mathcal{F}, p)$ denotes the Euclidean distance between p and its projection onto \mathcal{F} . The k -flat radius of P , which we denote by $R_k^{\text{opt}}(P)$, is the minimum, over all k -dimensional flats \mathcal{F} , of $\mathcal{RD}(\mathcal{F}, P)$. We consider the problem of computing $R_k^{\text{opt}}(P)$ for a given set of points P . We are interested in the high-dimensional case where d is a part of the input and not a constant. This problem is NP-hard even for $k = 1$. We present an algorithm that, given P and a parameter $0 < \varepsilon \leq 1$, returns a k -flat \mathcal{F} such that $\mathcal{RD}(\mathcal{F}, P) \leq (1 + \varepsilon)R_k^{\text{opt}}(P)$. The algorithm runs in $O(ndC_{\varepsilon,k})$ time, where $C_{\varepsilon,k}$ is a constant that depends only on ε and k . Thus the algorithm runs in time linear in the size of the point set and is a substantial improvement over previous known algorithms, whose running time is of the order of $dn^{O(k/\varepsilon^c)}$, where c is an appropriate constant.

1. Introduction

Let P be a set of n points in \mathbb{R}^d . The *radius* of a k -dimensional flat \mathcal{F} with respect to P , which we denote by $\mathcal{RD}(\mathcal{F}, P)$, is defined to be $\max_{p \in P} \text{dist}(\mathcal{F}, p)$, where $\text{dist}(\mathcal{F}, p)$ denotes the Euclidean distance between p and its projection onto \mathcal{F} . The k -flat radius of P , which we denote by $R_k^{\text{opt}}(P)$, is the minimum, over all k -dimensional flats \mathcal{F} ,

* A preliminary version of this paper appeared in *Proc. 19th Annual ACM Symposium on Computational Geometry*, pages 39–47, 2003. Work on this paper by the first author was partially supported by NSF CAREER Award CCR-0132901, and that by the second author was partially supported by NSF CAREER Award CCR-0237431.

of $\mathcal{RD}(\mathcal{F}, P)$. (See Section 2 for detailed definitions.) Thus $R_0^{\text{opt}}(P)$ is the radius of the min-enclosing ball of P , $R_1^{\text{opt}}(P)$ is the radius of the min-enclosing cylinder of P , and so on. Informally, the k -flat radius of P measures how well the point set P can be approximated by an affine subspace of dimension k . Computing the k -flat radius of a point set is a fundamental problem in computational convexity and has applications in data mining, statistics, and clustering [GK1], [GK2], [HV2].

The problem of computing the k -flat radius of a point set has received considerable attention in the computational geometry literature. A classic result is that the min-enclosing ball (i.e., $R_0^{\text{opt}}(P)$) of a point set P can be computed in linear time when the dimension is fixed. For $k \geq 1$, the k -flat radius of a set P of n points can be computed exactly in polynomial time in fixed dimension [FKS]. It can also be approximated to within a factor of $(1 + \varepsilon)$, for any $\varepsilon > 0$, in $O(n + (1/\varepsilon)^c)$ time, where c is a constant that depends only on d and k [BH], [HV1]. Thus the problem is reasonably well understood when the dimension d is taken to be a fixed constant. These algorithms are not satisfactory when the dimension is large as their running times depend exponentially on the dimension. In the rest of this paper we are interested in efficient algorithms when the dimension d can be as large as n .

It is well known that the minimum enclosing ball of a set of points can be computed in polynomial time; see for instance the paper by Gritzmann and Klee [GK1]. Megiddo [Me] shows that the problem of determining whether there is a line that intersects a set of balls is NP-hard. In his reduction, the balls have the same radius, which implies that computing the radius $R_1^{\text{opt}}(P)$ of the min-enclosing cylinder of a set of points P is NP-hard. Bădoiu et al. [BHI] give a poly-time algorithm that computes a $(1 + \varepsilon)$ -approximation, for any $\varepsilon > 0$, of the minimum enclosing cylinder (i.e., $R_1^{\text{opt}}(P)$) of a set of points. The authors [HV2] gave a poly-time algorithm that computes a $(1 + \varepsilon)$ -approximation, for any $\varepsilon > 0$, to $R_k^{\text{opt}}(P)$ whenever k is a fixed constant. This algorithm runs in, roughly, $dn^{O(k/\varepsilon^5)}$ time. These results show that the k -flat radius can be efficiently approximated for small k .

The problem seems to become harder when k becomes large. Bodlaender et al. [BGKvL] show that computing the width, $R_{d-1}^{\text{opt}}(P)$, of a point set is NP-hard. Gritzmann and Klee [GK1] show that it is NP-hard to compute the width of even a d -dimensional simplex (that is, a set of $d + 1$ points). They also show that it is NP-hard to compute $R_k^{\text{opt}}(P)$ for small point sets (consisting of $2d$ points) as long as $k \geq c \cdot d$, for any fixed $0 < c < 1$. Recently, Brieden [Br] showed that it is NP-hard to approximate the width of a point set to within *any* constant factor. Varadarajan et al. [VVZ] show that Brieden's result can be strengthened as follows: there exists a constant $\delta > 0$ such that it is NP-hard to approximate the k -flat radius of a set of n points to within a factor of $(\log n)^\delta$ whenever $k \geq d^\varepsilon$, for any constant $\varepsilon > 0$. Turning to upper bounds, the papers of Nesterov [Ne] and Nemirovski et al. [NRT] imply an $O(\sqrt{\log n})$ -approximation for the width (i.e., $R_{d-1}^{\text{opt}}(P)$) of a point set in polynomial time. Varadarajan et al. [VVZ] describe a poly-time algorithm that approximates the k -flat radius of a point set to within a factor of $O(\sqrt{\log n \cdot \log d})$ (for any k). We refer the reader to this paper for further details on approximating the k -flat radius for large k . Recently, Ye and Zhang [YZ] improved the result of Varadarajan et al. [VVZ] by presenting an algorithm that gives an $O(\sqrt{\log n})$ approximation.

Our current paper focuses on the case when k is small and significantly improves upon our previous work [HV2]. In that paper we presented an algorithm that, given

a point set P with n points and a parameter $0 < \varepsilon \leq 1$, computes a k -flat whose radius with respect to P is at most $(1 + \varepsilon)R_k^{\text{opt}}(P)$ in $dn^{O(k/\varepsilon^5 \log(1/\varepsilon))}$ time. The basic algorithm of [HV2] is based on proving the existence of a small set of points, called a coresset, such that the affine subspace spanned by the points contains a k -flat that is near-optimal. The size of the coresset depends on k and ε but not on the dimension d . Now, by brute-force enumeration on all such small sets, the algorithm generates all possible candidate coresets. For each candidate coresset, the algorithm finds the best k -flat in the affine subspace spanned by the coresset; this, being a “low-dimensional” problem, is feasible. The algorithm then returns the best k -flat overall. The nice property of coresets is that they are *deus ex machina*—their existence immediately implies relatively efficient approximation algorithms that extends immediately to fitting multiple-flats with outliers. On the negative side, these techniques have relatively bad dependency on n . In this paper we develop a new technique that bypasses the enumeration of coresets altogether, and present direct algorithms for computing a good k -flat. In particular, our new algorithms have linear dependency on the number of input points and the dimension.

Given a set P of n points in d dimensions (d might be as large as n) and a parameter $0 < \varepsilon \leq 1$, we present efficient algorithms for the following problems:

- **Minimum Radius Line.** In Section 3.2 we present an approximation algorithm that returns a line ℓ such that $\mathcal{RD}(\ell, P) \leq (1 + \varepsilon)R_1^{\text{opt}}(P)$. The running time of the new algorithm is $O(ndC_\varepsilon)$, where $C_\varepsilon = \exp((1/\varepsilon^3) \log^2(1/\varepsilon))$.

This is a substantial improvement over the previously fastest algorithms of [BC], [BHI], [KMY], and [HV2], which all had running times of the form $dn^{O(1/\varepsilon^c)}$, where c is an appropriate constant.

It is natural to ask if there is an algorithm for solving this problem with a running time that depends polynomially on n , d , and $1/\varepsilon$. We prove in Section 4 that such an algorithm does not exist unless $P = NP$. We do this by showing that the NP-hardness reduction of Megiddo [Me] can be modified so that it yields an appropriate hardness of approximation result.

- **Minimum Radius k -Flat.** In Section 5 we generalize the algorithm for a line to compute a k -flat \mathcal{F} such that $\mathcal{RD}(\mathcal{F}, P) \leq (1 + \varepsilon)R_k^{\text{opt}}(P)$. The running time of the new algorithm is $n \cdot d \cdot \exp((e^{O(k^2)}/\varepsilon^{2k+1}) \log^2(1/\varepsilon))$. This substantially improves over the previous fastest algorithm [BC], [KMY], [HV2] that has running time $dn^{O(k/\varepsilon^c)}$, where c is an appropriate constant. It is clear that the hardness result for $k = 1$ implies a hardness result for any $k \geq 1$, but we have not attempted to prove this rigorously.

The main result of this paper is the algorithm for the case $k = 1$, which extends naturally to the case $k > 1$. We describe the ideas behind this algorithm. Let ℓ_{opt} denote an optimal line for the point set P , and let ℓ be any line. If ℓ is not nearly as good as ℓ_{opt} , then the 2-flat h containing ℓ and the point in P maximizing $\text{dist}(\ell, p)$ has a line ℓ' that is “significantly closer” to ℓ_{opt} than ℓ . Moreover, we can compute a small number of candidate lines on h , one of which is guaranteed to contain ℓ' . By trying all the candidates, we are guaranteed to have a line that is significantly closer to ℓ_{opt} than ℓ . By repeating this process a small number of times, we either stumble upon a line that is near-optimal, or we end up with a line that is very close to ℓ_{opt} and is therefore near-

optimal. To realize this vision, we need an appropriate notion of “closeness” between lines and corresponding machinery to argue about convergence; the provision of these is the most interesting contribution of this paper.

2. Preliminaries

Definition 2.1. Given j points $v_1, \dots, v_j \in \mathbb{R}^d$, the linear subspace they span is denoted by $\text{span}(v_1, \dots, v_j) = \{v \mid v = \sum_{i=1}^j a_i v_i \text{ for any } a_1, \dots, a_j \in \mathbb{R}\}$.

Given $j + 1$ points p_1, \dots, p_{j+1} in \mathbb{R}^d , the *affine space* spanned by them is

$$\begin{aligned} & \text{affine}(p_1, \dots, p_{j+1}) \\ &= \left\{ v \mid v = \sum_{i=1}^{j+1} a_i p_i, \text{ for any } a_1, \dots, a_{j+1} \in \mathbb{R} \text{ such that } \sum_{i=1}^{j+1} a_i = 1 \right\}. \end{aligned}$$

Alternatively, $\text{affine}(p_1, \dots, p_{j+1}) = p_1 + \text{span}(p_2 - p_1, p_3 - p_1, \dots, p_{j+1} - p_1)$.

We refer to an affine subspace of dimension k as a k -flat. When the dimension is not relevant, we simply refer to an affine subspace as a flat. See [Ed] for more on these definitions.

Definition 2.2. For a flat \mathcal{F} in \mathbb{R}^d and a point $y \in \mathbb{R}^d$, let $\text{proj}(\mathcal{F}, y)$ denote the projection of y onto \mathcal{F} . Namely, $\text{proj}(\mathcal{F}, y) = \text{argmin}_{x \in \mathcal{F}} \|xy\|$, where $\|xy\|$ denotes the *Euclidean distance* between x and y . For a set point $Y \in \mathbb{R}^d$, let $\text{proj}(\mathcal{F}, Y) = \{\text{proj}(\mathcal{F}, y) \mid y \in Y\}$.

For a flat \mathcal{F} and a point y , we let $\text{dist}(\mathcal{F}, y)$ denote the distance $\|y \text{proj}(\mathcal{F}, y)\|$. We sometimes refer to $\text{dist}(\mathcal{F}, y)$ as the distance of y from \mathcal{F} . For a flat \mathcal{F} and point set $P \subseteq \mathbb{R}^d$, the *radius* of \mathcal{F} with respect to P is $\mathcal{RD}(\mathcal{F}, P) = \max_{p \in P} \text{dist}(\mathcal{F}, p)$.

For $k \geq 0$, let $R_k^{\text{opt}}(P)$ denote the radius of a k -flat which minimizes the radius with respect to P . Formally, $R_k^{\text{opt}}(P) = \min_{\mathcal{F} \in \text{FLT}_k} \mathcal{RD}(\mathcal{F}, P)$, where FLT_k is the set of all k -flats in \mathbb{R}^d .

Given a k -flat \mathcal{F} , a point p on \mathcal{F} , and positive real numbers $0 < \beta \leq \alpha$, an (α, β) -net around p on \mathcal{F} is a set S of points on \mathcal{F} such that for any point q on \mathcal{F} such that $\|qp\| \leq \alpha$, there is a point $s \in S$ such that $\|qs\| \leq \beta$. It is well known that there exists an (α, β) -net S of size $\exp(O(k \log(\alpha/\beta)))$ [Mat], and it is easy to see that a larger (α, β) -net of size $\exp(O(k \log(k\alpha/\beta)))$ can be computed in $\exp(O(k \log(k\alpha/\beta)))$ time.

2.1. Distance Function

Definition 2.3. Given a j -flat \mathcal{G} , the *distance function* $d_{\mathcal{G}}(x)$, for $x \in \mathbb{R}^d$, is the distance of x from its projection onto \mathcal{G} ; namely, $d_{\mathcal{G}}(x) = \|x \text{proj}(\mathcal{G}, x)\|$.

Lemma 2.4. *Let $d_{\mathcal{G}}(\cdot)$ be the distance function to a flat \mathcal{G} . Let $x, y \in \mathbb{R}^d$ be any two points.*

- (i) *For any $0 \leq \beta \leq 1$, $d_{\mathcal{G}}(\beta x + (1 - \beta)y) \leq \beta d_{\mathcal{G}}(x) + (1 - \beta)d_{\mathcal{G}}(y)$. That is, $d_{\mathcal{G}}(\cdot)$ is convex.*
- (ii) *Let ℓ be the line through x and y , and w and z be any two points on ℓ . Then*

$$d_{\mathcal{G}}(z) \leq d_{\mathcal{G}}(w) + 2 \frac{\|wz\|}{\|xy\|} \max(d_{\mathcal{G}}(x), d_{\mathcal{G}}(y)).$$

Proof. Suppose that the line ℓ through x and y is parameterized by $\{x + tv \mid t \in \mathbb{R}\}$, where $v \in \mathbb{R}^d$ is the unit vector $(1/\|xy\|)(y - x)$. Let $d(t)$ denote the distance of the point $\ell(t) = x + tv$ on ℓ from \mathcal{G} , that is, $d(t) = d_{\mathcal{G}}(\ell(t))$. Then $d(t)$ is of the form $d(t) = \sqrt{at^2 + bt + c}$, where a, b, c are appropriate constants that depend on ℓ (specifically x, y) and \mathcal{G} . Moreover, the constants a, b , and c are such that $at^2 + bt + c$ is non-negative for every $t \in \mathbb{R}$. This implies that $a \geq 0$. Since

$$at^2 + bt + c = a \left(t + \frac{b}{2a} \right)^2 + \frac{4ac - b^2}{4a},$$

this also implies that $4ac - b^2 \geq 0$. Let $\alpha = -b/2a$, and let $d = (4ac - b^2)/4a$. Thus $d(t) = \sqrt{a(t - \alpha)^2 + d}$, where $d \geq 0$.

(i) This is well known but we give the proof for completeness. If $d = 0$, $d(t)$ is a linear function in t and is therefore convex. So we assume that $d > 0$ and thus $4ac - b^2 > 0$ and $d(t) > 0$ for all t . Observe that

$$\begin{aligned} d'(t) &= \frac{1}{2}(2at + b)(at^2 + bt + c)^{-1/2} = (at + b/2)(at^2 + bt + c)^{-1/2} \quad \text{and} \\ d''(t) &= a(at^2 + bt + c)^{-1/2} - (at + b/2)^2(at^2 + bt + c)^{-3/2}. \end{aligned}$$

Now

$$f(t) = d''(t)(at^2 + bt + c)^{3/2} = a(at^2 + bt + c) - (at + b/2)^2 = (4ac - b^2)/4 > 0$$

for all t . Since $at^2 + bt + c > 0$ for all t , we conclude that $d''(t) > 0$ for all t . This implies that $d(t)$ is convex [Man, page 89].

(ii) Consider the function $g(t) = \sqrt{a(t - \alpha)^2} = \sqrt{a}|t - \alpha|$. It is easy to verify that $g(t) \leq d(t)$ for all $t \in \mathbb{R}$ and $|d'(t)| \leq |g'(t)|$ for all $t \in \mathbb{R} \setminus \{\alpha\}$.

Let $U = \max(d_{\mathcal{G}}(x), d_{\mathcal{G}}(y))$. Because of the convexity established in part (i), we have $d(t) \leq U$ for any $t \in [t_x, t_y]$, and thus, $g(t) \leq U$, for all $t \in [t_x, t_y]$. Note that $|g'(t)| = \sqrt{a}$, for all $t \in \mathbb{R} \setminus \{\alpha\}$, and the value of \sqrt{a} is maximized when α is the middle point of $[t_x, t_y]$. Otherwise, since $g(\cdot)$ is a symmetric function around α , we could move α to the middle of the segment $[t_x, t_y]$ and increase the value of a , and get a new function which complies with the above conditions. Furthermore, in such a case, a is maximized when $g(t_x) = g(t_y) = U$. In this case, $a = (2U/(t_y - t_x))^2$, and $|g'(t)| = 2U/(t_y - t_x)$, for all $t \in \mathbb{R} \setminus \{\alpha\}$. This implies $|d'(t)| \leq |g'(t)| \leq 2U/(t_y - t_x)$, for all $t \in \mathbb{R}$.




Fig. 1. The triangles bac and bda are similar.

Suppose that $\ell(t_w) = w$ and $\ell(t_z) = z$. Because of the bound on $d'(t)$, we have

$$d(t_z) \leq d(t_w) + (\max_t |d'(t)|) |t_z - t_w| \leq d(t_w) + \frac{2U}{t_y - t_x} \cdot |t_z - t_w|,$$

which completes the proof. \square

2.2. Triangles

We now state a simple lemma that will play an important role in this paper. See Fig. 1.

Lemma 2.5. *Let a, b , and c be points in \mathbb{R}^d such that $\|ab\| \leq r$ and $\|bc\| \geq (1 + \varepsilon)r$, where $r \geq 0$ and $0 < \varepsilon \leq 1$. Then there is a point d on the segment bc such that $\|ad\| \leq (1 - \varepsilon/2)\|ac\|$.*

Proof. Let $\rho = \|ab\|/\|bc\|$, and let d be the point placed on segment bc at distance $\rho\|ab\|$ from b . It is easy to see that $\triangle bac$ is similar to $\triangle bda$, with a scaling factor of ρ . Thus

$$\|ad\| = \rho\|ac\| \leq \frac{1}{1 + \varepsilon} \|ac\| \leq \left(1 - \frac{\varepsilon}{2}\right) \|ac\|. \quad \square$$

2.3. Rotations

Let \mathcal{F} be a k -dimensional flat in \mathbb{R}^d , let \mathcal{G} be a j -dimensional flat that lies in \mathcal{F} , where $j < k$, and let p be any point in \mathbb{R}^d . We define $\text{Rot}(\mathcal{F}, \mathcal{G}, p)$ to be the *rotation of \mathcal{F} around \mathcal{G} such that it passes through p* . Formally, if $p \in \mathcal{F}$, $\text{Rot}(\mathcal{F}, \mathcal{G}, p)$ is \mathcal{F} itself. Otherwise, let $p' = \text{proj}(\mathcal{F}, p)$ and $p'' = \text{proj}(\mathcal{G}, p') = \text{proj}(\mathcal{G}, p)$. If $p' = p''$, we let \mathcal{H} be any $(k - 1)$ -flat that contains \mathcal{G} ; otherwise, we let \mathcal{H} be the $(k - 1)$ -flat contained in \mathcal{F} that passes through p'' and is orthogonal to the vector $\overrightarrow{p''p'}$. It is easy to see that in either case \mathcal{H} contains \mathcal{G} . We define $\text{Rot}(\mathcal{F}, \mathcal{G}, p)$ to be the k -flat that contains \mathcal{H} and p . See Fig. 2.




Fig. 2. (i) $\ell' = \text{Rot}(\ell, q, p)$ for a line ℓ , a point q on ℓ , and a point p . (ii) In the proof of Lemma 2.6, triangles $pp'p''$ and $rr'r''$ are similar.

As we rotate \mathcal{F} around \mathcal{G} , the distance from an arbitrary point q to it changes. The rate of this change is a function of the distance of q from \mathcal{G} , and is roughly proportional to the ratio between this distance and the distance of p (the “target” point) from the “axis of rotation” \mathcal{G} . The following lemma states this fact more formally.

Lemma 2.6. *Let \mathcal{F} be a k -flat in \mathbb{R}^d , let \mathcal{G} be a j -flat that lies in \mathcal{F} , where $j < k$, and let p be any point in \mathbb{R}^d not on \mathcal{G} . For any $q \in \mathbb{R}^d$, we have*

$$\text{dist}(\text{Rot}(\mathcal{F}, \mathcal{G}, p), q) \leq \text{dist}(\mathcal{F}, q) + \frac{\text{dist}(\mathcal{G}, q)}{\text{dist}(\mathcal{G}, p)} \text{dist}(\mathcal{F}, p).$$

Proof. The claim follows readily when $p \in \mathcal{F}$, since in that case $\text{Rot}(\mathcal{F}, \mathcal{G}, p) = \mathcal{F}$. Otherwise, let $p' = \text{proj}(\mathcal{F}, p)$ and $p'' = \text{proj}(\mathcal{G}, p') = \text{proj}(\mathcal{G}, p)$. Let \mathcal{H} be the $(k-1)$ -flat in the definition of $\text{Rot}(\mathcal{F}, \mathcal{G}, p)$; that is, \mathcal{H} is the intersection of $\text{Rot}(\mathcal{F}, \mathcal{G}, p)$ with \mathcal{F} . It is easy to see that $p'' = \text{proj}(\mathcal{H}, p') = \text{proj}(\mathcal{H}, p)$.

Let $r = \text{proj}(\mathcal{F}, q)$. If r lies on \mathcal{H} , then the lemma follows because $\text{dist}(\text{Rot}(\mathcal{F}, \mathcal{G}, p), q) \leq \text{dist}(\mathcal{H}, q) = \text{dist}(\mathcal{F}, q)$. If r does not lie on \mathcal{H} , let $r' = \text{proj}(\text{Rot}(\mathcal{F}, \mathcal{G}, p), r)$ and $r'' = \text{proj}(\mathcal{H}, r) = \text{proj}(\mathcal{H}, r')$. We claim that $\triangle pp'p''$ and $\triangle rr'r''$ are similar. To see that, consider the projection T , which projects \mathbb{R}^d to the $d-k+1$ subspace perpendicular to \mathcal{H} (namely, \mathcal{H} is mapped to a point by T). Clearly, the lengths of the edges of the two triangles we are interested in are preserved under this projection, and they both lie on a common 2-flat. The triangles are clearly similar, as the situation is now as depicted in Fig. 2. It follows that

$$\|rr''\| = \frac{\|rr''\|}{\|pp''\|} \|pp''\| = \frac{\text{dist}(\mathcal{H}, r)}{\text{dist}(\mathcal{G}, p)} \text{dist}(\mathcal{F}, p) \leq \frac{\text{dist}(\mathcal{G}, r)}{\text{dist}(\mathcal{G}, p)} \text{dist}(\mathcal{F}, p).$$

We conclude that

$$\text{dist}(\text{Rot}(\mathcal{F}, \mathcal{G}, p), q) \leq \|qr\| + \|rr''\| = \text{dist}(\mathcal{F}, q) + \frac{\text{dist}(\mathcal{G}, q)}{\text{dist}(\mathcal{G}, p)} \text{dist}(\mathcal{F}, p). \quad \square$$

3. Minimum Radius Cylinder

In this section we present an efficient algorithm that given a set $P \subseteq \mathbb{R}^d$ of n points and a parameter $0 < \varepsilon \leq 1$ computes a line ℓ such that $\mathcal{RD}(\ell, P) \leq (1 + \varepsilon)R_1^{\text{opt}}(P)$. It is convenient first to describe the algorithm assuming the presence of an appropriate oracle (to be described below). At any time during its execution, the algorithm can present the oracle with a set of possible choices, and the oracle simply picks the right one.

The Algorithm. Let ℓ_{opt} denote an optimal line for the input point set P , that is, $\mathcal{RD}(\ell_{\text{opt}}, P) = R_1^{\text{opt}}(P)$. Let $R^{\text{opt}} = R_1^{\text{opt}}(P)$. For any $p \in P$, let $p' = \text{proj}(\ell_{\text{opt}}, p)$ denote the projection of p onto ℓ_{opt} . Let $P' = \{p' \mid p \in P\}$. Let $\mathcal{I} = \mathcal{CH}(P')$ denote the convex hull of P' . Clearly, \mathcal{I} is the line segment joining u' and v' , for some $u, v \in P$. The length L of the segment \mathcal{I} is clearly at most $\text{diam}(P)$. We need the following lemma, whose proof is presented at the end of this section.

Lemma 3.1. *Let P be a set of n points in \mathbb{R}^d , let ℓ_{opt} be a line such that $\mathcal{RD}(\ell_{\text{opt}}, P) = R_1^{\text{opt}}(P)$, let \mathcal{I} be the convex hull of $\text{proj}(\ell_{\text{opt}}, P)$, let $r \geq 0$ be a number such that $R_1^{\text{opt}}(P) \leq r \leq 4R_1^{\text{opt}}(P)$, let h be any 2-flat, and let $0 < \beta \leq 1$ be a parameter. We can compute, in $O(nd/\beta^2)$ time, a family of $O(1/\beta^2)$ lines on h such that at least one line ℓ in the family has the property that for any point $x \in \mathcal{I}$, $\text{dist}(\ell, x) \leq \text{dist}(\text{proj}(h, \ell_{\text{opt}}), x) + \beta R_1^{\text{opt}}(P)$.*

Let p_Δ be any point of P , and let q_Δ be the farthest point of P from p . Let $\ell_0 = \text{affine}(p_\Delta, q_\Delta)$. It is easy to verify that $\mathcal{RD}(\ell_0, P) \leq 4R^{\text{opt}}$; see Lemma 5.2 for a proof of a more general result. We compute a sequence of lines ℓ_0, \dots, ℓ_ν , where $\nu = c(1/\varepsilon^3) \log(1/\varepsilon)$ and c is a sufficiently large constant to be determined below. We describe below how the line ℓ_i is computed from line ℓ_{i-1} in the i th iteration.

In the i th iteration, let p_i be the point of P farthest away from ℓ_{i-1} . Let h_i be the 2-flat containing ℓ_{i-1} and p_i and let $\widehat{\ell}_i$ denote the projection of ℓ_{opt} onto h_i . Using the algorithm of Lemma 3.1 (setting $r = \mathcal{RD}(\ell_0, P)$), we compute a family of $O(1/\delta^2)$ lines on h_i such that at least one line ℓ in the family has the property that for any $x \in \mathcal{I}$, $\text{dist}(\ell, x) \leq \text{dist}(\widehat{\ell}_i, x) + \delta R^{\text{opt}}$, where $\delta = \varepsilon/4\nu$. We present the oracle for this family of lines, and the oracle picks a line with this property. It can do this by specifying $O(\log 1/\delta)$ bits. We let ℓ_i be the line chosen by the oracle.

At the end of the ν th iteration, we return the best line from the sequence ℓ_0, \dots, ℓ_ν . That is, we return the line ℓ from the sequence that minimizes $\mathcal{RD}(\ell, P)$. We argue below that $\mathcal{RD}(\ell, P) \leq (1 + \varepsilon)R^{\text{opt}}$ for such a line ℓ . Let us assume the contrary, that is, $\mathcal{RD}(\ell_i, P) > (1 + \varepsilon)R^{\text{opt}}$, for each $0 \leq i \leq \nu$. We will derive a contradiction.

Proof of Correctness. For $0 \leq i \leq \nu$, let $d_i(x) = \|x \text{proj}(\ell_i, x)\|$ denote the distance between a point $x \in \mathbb{R}^d$ and ℓ_i . For each $1 \leq i \leq \nu$ and any $x \in \mathcal{I}$, we have

$$d_i(x) \leq \text{dist}(\widehat{\ell}_i, x) + \delta R^{\text{opt}} \leq d_{i-1}(x) + \delta R^{\text{opt}}. \quad (1)$$

Intuitively, d_1, d_2, \dots , are almost monotonically decreasing functions converging to the zero function on the points in \mathcal{I} .

Clearly, $\text{dist}(u, u') \leq R^{\text{opt}}$ (we remind the reader that u' and v' are the two endpoints of \mathcal{I}). Since $\mathcal{RD}(\ell_0, P) \leq 4R^{\text{opt}}$, we have $\text{dist}(\ell_0, u) \leq 4R^{\text{opt}}$. Thus

$$d_0(u') = \text{dist}(\ell_0, u') \leq \text{dist}(\ell_0, u) + \text{dist}(u, u') \leq 5R^{\text{opt}}.$$

By a symmetric argument, we have $d_0(v') \leq 5R^{\text{opt}}$. From Lemma 2.4(i), we conclude that $d_0(x) \leq 5R^{\text{opt}}$ for any $x \in \mathcal{I}$. Using (1), it then follows that for $i = 1, \dots, \nu$ and any $x \in \mathcal{I}$, we have $d_i(x) \leq d_0(x) + i\delta R^{\text{opt}} \leq 6R^{\text{opt}}$.

Spread $\lceil 160/\varepsilon^2 \rceil + 1$ equally spaced points on the segment \mathcal{I} , and let S denote this set. We say that a point $z \in S$ is *hit* in the i th iteration if $d_{i-1}(z) \geq (\varepsilon/2)R^{\text{opt}}$ and $d_i(z) \leq (1 - \varepsilon/5)d_{i-1}(z)$. (Intuitively, every time a point z in S is being hit, the value associated with it, $d_i(z)$, goes down “considerably”.) Suppose z has been hit m times till the j th iteration; from (1) we have

$$d_j(z) \leq (1 - \varepsilon/5)^m d_0(z) + \nu \delta R^{\text{opt}} \leq (1 - \varepsilon/5)^m 5R^{\text{opt}} + \nu \delta R^{\text{opt}}.$$

Thus, for $m = O((1/\varepsilon) \log(1/\varepsilon))$, we have $5(1 - \varepsilon/5)^m \leq \varepsilon/4$ and

$$d_j(z) \leq (\varepsilon/4)R^{\text{opt}} + \nu \cdot (\varepsilon/(4\nu)) \cdot R^{\text{opt}} \leq (\varepsilon/2)R^{\text{opt}}.$$

Thus, after z has been hit $m = O((1/\varepsilon) \log(1/\varepsilon))$ times, it can never be hit again. Note that the bound on m is independent of ν .

We claim that in every iteration of the algorithm, at least one point of S is being hit. Indeed, consider p_i and p'_i , the projection of p_i into ℓ_{opt} . Let c_i be the projection of p'_i onto ℓ_i . We know that $\|p'_i p_i\| \leq R^{\text{opt}}$, and $\|p_i c_i\| > (1 + \varepsilon)R^{\text{opt}}$ (since $\mathcal{RD}(\ell_{i-1}, P) > (1 + \varepsilon)R^{\text{opt}}$). From Lemma 2.5, we conclude that $\text{dist}(p'_i, p_i c_i) \leq (1 - \varepsilon/2)\|p'_i c_i\|$. Thus, $d_i(p'_i) \leq \text{dist}(p'_i, p_i c_i) + \delta R^{\text{opt}} \leq (1 - \varepsilon/2)\|p'_i c_i\| + \delta R^{\text{opt}} = (1 - \varepsilon/2)d_{i-1}(p'_i) + \delta R^{\text{opt}}$.

Namely, $d_i(p'_i)$ is significantly smaller than $d_{i-1}(p'_i)$. Also, we note that $d_{i-1}(p'_i) \geq \varepsilon R^{\text{opt}}$, as otherwise

$$\begin{aligned} \mathcal{RD}(\ell_{i-1}, P) &= \text{dist}(\ell_{i-1}, p_i) \leq \text{dist}(p'_i, p_i) + \text{dist}(\ell_{i-1}, p'_i) \\ &\leq R^{\text{opt}} + \varepsilon R^{\text{opt}} = (1 + \varepsilon)R^{\text{opt}}. \end{aligned}$$

Thus, we have

$$\begin{aligned} d_i(p'_i) &\leq \left(1 - \frac{\varepsilon}{2}\right) d_{i-1}(p'_i) + \delta R^{\text{opt}} \leq \left(1 - \frac{\varepsilon}{2}\right) d_{i-1}(p'_i) + \delta \frac{d_{i-1}(p'_i)}{\varepsilon} \\ &\leq \left(1 - \frac{\varepsilon}{3}\right) d_{i-1}(p'_i), \end{aligned}$$

since $\delta < \varepsilon^2/6$ for $\varepsilon \leq 1$ and sufficiently large c .

Let $y \in S$ be a point such that $\|p'_i y\| \leq \varepsilon^2 L/320$. We argue that y is hit in the i th iteration. From Lemma 2.4(ii), we have

$$\begin{aligned} |d_{i-1}(y) - d_{i-1}(p'_i)| &\leq \frac{\|p'_i y\|}{u'v'} \cdot 2 \max(d_{i-1}(u'), d_{i-1}(v')) \leq \frac{\varepsilon^2 L/320}{L} \cdot 12R^{\text{opt}} \\ &\leq \frac{\varepsilon^2}{16} R^{\text{opt}}. \end{aligned}$$

Consequently, we have $d_{i-1}(y) \geq d_{i-1}(p'_i) - (\varepsilon^2/16)R^{\text{opt}} \geq \varepsilon R^{\text{opt}} - \varepsilon^2/16 \geq (15\varepsilon/16)R^{\text{opt}}$.

By applying Lemma 2.4 again, this time to $d_i(\cdot)$,

$$\begin{aligned}
d_i(y) &\leq d_i(p'_i) + \frac{\varepsilon^2}{16}R^{\text{opt}} \\
&\leq \left(1 - \frac{\varepsilon}{3}\right)d_{i-1}(p'_i) + \frac{\varepsilon^2}{16}R^{\text{opt}} \\
&\leq \left(1 - \frac{\varepsilon}{3}\right)\left(d_{i-1}(y) + \frac{\varepsilon^2}{16}R^{\text{opt}}\right) + \frac{\varepsilon^2}{16}R^{\text{opt}} \\
&\leq \left(1 - \frac{\varepsilon}{3}\right)d_{i-1}(y) + \frac{2\varepsilon^2}{16}R^{\text{opt}} \\
&\leq \left(1 - \frac{\varepsilon}{3}\right)d_{i-1}(y) + \frac{2\varepsilon}{15} \frac{15\varepsilon}{16}R^{\text{opt}} \\
&\leq \left(1 - \frac{\varepsilon}{3}\right)d_{i-1}(y) + \frac{2\varepsilon}{15}d_{i-1}(y) \\
&\leq \left(1 - \frac{\varepsilon}{5}\right)d_{i-1}(y).
\end{aligned}$$

Thus, the point y is hit in the i th iteration.

We choose the constant c large enough so that the number of iterations $\nu = c(1/\varepsilon^3) \log(1/\varepsilon)$ is larger than $m \cdot |S|$. Since a point from S is hit in each of the ν iterations, but each point in S is hit at most m times, we have a contradiction.

Removing the Oracle. The algorithm as we described it uses $O(\nu \log(1/\delta)) = O((1/\varepsilon^3) \log^2(1/\varepsilon))$ bits from the oracle. To remove the dependence on the oracle, we simply try all possible strings of size $O(\nu \log(1/\delta))$, and execute the algorithm on each of these strings. The overall running time of the resulting algorithm is $n \cdot d \cdot C_\varepsilon$, where $C_\varepsilon = \exp(O((1/\varepsilon^3) \log^2(1/\varepsilon)))$. We therefore conclude:

Theorem 3.2. *Let P be a set of n points in \mathbb{R}^d and let $0 < \varepsilon \leq 1$ be a parameter. We can compute a line ℓ , such that $\mathcal{RD}(\ell, P) \leq (1 + \varepsilon)R_1^{\text{opt}}(P)$, in $n \cdot d \cdot C_\varepsilon$ time, where $C_\varepsilon = \exp(O((1/\varepsilon^3) \log^2(1/\varepsilon)))$.*

Proof of Lemma 3.1. We generate a set of $O(1/\beta^2)$ pairs of points (p_1, p_2) as follows. Set $\gamma = \beta/8$. Let p be any point in P , and let p_1 be any point from a $(5r, \gamma r)$ -net around $\text{proj}(h, p)$ on h . For each choice of p_1 , we find the point $q \in P$ whose projection $\text{proj}(h, q)$ is farthest from p_1 , and we choose p_2 from a $(5r, \gamma r)$ -net around $\text{proj}(h, q)$ on h . For each pair (p_1, p_2) that is generated, we include the line through p_1 and p_2 in our family.

We now argue that our family has the required line in it. Since projection does not expand distances, we have $\text{dist}(\text{proj}(h, \ell_{\text{opt}}), \text{proj}(h, p)) \leq r$, so there is a choice p_1^* of p_1 such that $\text{dist}(\text{proj}(h, \ell_{\text{opt}}), p_1^*) \leq \gamma r$. Let ℓ' be the translation of $\text{proj}(h, \ell_{\text{opt}})$ through p_1^* . Let $q^* \in P$ be the point whose projection is farthest from p_1^* . Since $\text{dist}(\ell', \text{proj}(h, q^*)) \leq r + \gamma r \leq 2r$, there exists a choice p_2^* for p_2 from the $(5r, \gamma r)$ -net around $\text{proj}(h, q^*)$

such that (a) $\text{dist}(\ell', p_2^*) \leq \gamma r$, and (b) $\|p_1^* p_2^*\| \geq \|p_1^* \text{proj}(h, q^*)\| + r$. Let ℓ be the line through p_1^* and p_2^* . Notice that $\ell = \text{Rot}(\ell', p_1^*, p_2^*)$.

For any $s \in P$, let s' denote $\text{proj}(\ell_{\text{opt}}, s)$. Since projection does not expand distances, we have

$$\|\text{proj}(h, s') p_1^*\| \leq \|\text{proj}(h, s) p_1^*\| + r \leq \|\text{proj}(h, q^*) p_1^*\| + r \leq \|p_1^* p_2^*\|.$$

From Lemma 2.6, we conclude that for any $s \in P$,

$$\begin{aligned} \text{dist}(\ell, \text{proj}(h, s')) &\leq \text{dist}(\ell', \text{proj}(h, s')) + \frac{\|\text{proj}(h, s') p_1^*\|}{\|p_1^* p_2^*\|} \text{dist}(\ell', p_2^*) \\ &\leq \text{dist}(\ell', \text{proj}(h, s')) + \text{dist}(\ell', p_2^*) \leq 2\gamma r. \end{aligned}$$

Let $x \in \mathcal{I}$ be any point. Since $\text{proj}(h, x)$ lies in the convex hull of the set $\{\text{proj}(h, s') \mid s \in P\}$, we conclude using Lemma 2.4(i) that $\text{dist}(\ell, \text{proj}(h, x)) \leq 2\gamma r \leq \beta R_1^{\text{opt}}(P)$, which implies that ℓ has the properties claimed in the lemma. \square

4. A Lower Bound

Megiddo [Me] shows that the problem of determining whether there is a line that intersects a given set of balls is NP-hard. The balls in his construction have the same radius, which implies that the problem of computing the radius of the min-enclosing cylinder (i.e., R_1^{opt}) of a set of points is also NP-hard. We show below that his construction yields the following hardness of approximation result: unless $P = NP$, there is no algorithm that, given a set of n points in \mathbb{R}^d and an $0 < \varepsilon \leq 1$, runs in time polynomial in n, d , and $1/\varepsilon$, and returns a number r such that $R_1^{\text{opt}}(P) \leq r \leq (1 + \varepsilon)R_1^{\text{opt}}(P)$.

Megiddo gives a reduction from 3CNF-satisfiability. Let φ be a 3CNF formula with n variables and m clauses in which, without loss of generality, we assume that each clause consists of three distinct variables. Let x_1, \dots, x_n denote the literals and E_1, \dots, E_m denote the clauses in φ . Let $d = n + 1$. Let e_i denote the point in \mathbb{R}^d with 1 in the i th coordinate and 0 elsewhere. Let $r_d = \sqrt{1 - (1/d)}$, and α be chosen so that

$$(12 - 4/d)\alpha^2 = r_d^2.$$

Let P_1 be the set of $2d$ points $\{\pm e_i \mid 1 \leq i \leq d\}$. Let Q denote the set $\{(x_1, \dots, x_d) \mid x_i = \pm 1/\sqrt{d}\}$ of 2^d points. Let L denote the set of lines obtained by considering each point $q \in Q$ and taking the line passing through the origin and q . Though Q consists of 2^d points, L has only 2^{d-1} lines. As shown by Megiddo, $R_1^{\text{opt}}(P_1)$ equals r_d , and this is attained by exactly the lines in L . Megiddo constructs a set of m points $P_2 = \{p^1, \dots, p^m\}$, one for each clause E_j of φ . The point $p^j = (p_1^j, \dots, p_{n+1}^j)$ is constructed as follows. The last coordinate p_{n+1}^j is set to 3α . For $1 \leq i \leq n$, if the variable x_i does not occur in E_j , then $p_i^j = 0$; if the literal x_i occurs in E_j , then $p_i^j = \alpha$; if the literal \bar{x}_i occurs in E_j , then $p_i^j = -\alpha$.¹ Megiddo shows that the following

¹ Since α may be irrational, we have to use instead a good rational approximation that can be represented with a small number of bits. We ignore this technicality here since it can be easily accommodated.

properties hold:

1. For each point $p \in P_2$, $\|p\| \leq 1$.
2. If φ is satisfiable, then there is a line $\ell \in L$ such that $\text{dist}(\ell, p)^2 \leq (12 - 4/d)\alpha^2 = r_d^2$ for each $p \in P_2$. Thus $R_1^{\text{opt}}(P_1 \cup P_2) = r_d$ in this case.
3. If φ is not satisfiable, then for every line $\ell \in L$, there is a point $p \in P_2$ such that $\text{dist}(\ell, p)^2 = 12\alpha^2 > r_d^2$.

We add, for each point p in P_2 , the point $-p$ into the set P_2 . It is easy to check that the properties above hold for the new P_2 . (We need the fact that for a line ℓ through the origin, $\text{dist}(\ell, p) = \text{dist}(\ell, -p)$.) Our goal here is to show that in the case where φ is not satisfiable, $R_1^{\text{opt}}(P_1 \cup P_2)$ is significantly larger than r_d . We need the following lemma.

Lemma 4.1. *Let $x = (x_1, \dots, x_d)$ and $y = (y_1, \dots, y_d)$ be unit vectors such that $|x_i - y_i| \leq 1/14d^2$, for each $1 \leq i \leq d$. Let ℓ_x (resp. ℓ_y) denote the line through the origin and the point x (resp. y). Let p be a point such that $\|p\| \leq 1$. Then $|\text{dist}(\ell_x, p)^2 - \text{dist}(\ell_y, p)^2| \leq 1/7d$.*

Proof. Notice that $\|x - y\| \leq 1/14d^{3/2} \leq 1/14d$.

$$\begin{aligned}
|\text{dist}(\ell_x, p)^2 - \text{dist}(\ell_y, p)^2| &= |\langle p, x \rangle^2 - \langle p, y \rangle^2| \\
&= |(\langle p, x \rangle - \langle p, y \rangle)(\langle p, x \rangle + \langle p, y \rangle)| \\
&\leq 2|\langle p, x \rangle - \langle p, y \rangle| \\
&= 2|\langle p, x - y \rangle| \\
&\leq 2\|p\| * \|x - y\| \\
&\leq 1/7d. \quad \square
\end{aligned}$$

Suppose that φ is unsatisfiable. Let ℓ denote the line that achieves $R_1^{\text{opt}}(P_1 \cup P_2)$. Since $P_1 \cup P_2$ is symmetric, we may assume that ℓ passes through the origin. Suppose ℓ also passes through the point $z = (z_1, \dots, z_d)$, where z is a unit vector.

Case 1. Suppose that $||z_i| - 1/\sqrt{d}| \leq 1/14d^2$ for each $1 \leq i \leq d$. Then there is a point $q = (q_1, \dots, q_d) \in Q$ such that $|q_i - z_i| \leq 1/14d^2$ for each $1 \leq i \leq d$. Let $\ell_q \in L$ be the line through the origin and q . Since φ is unsatisfiable, there exists a point $p \in P_2$ such that $\text{dist}(\ell_q, p)^2 = 12\alpha^2$. From Lemma 4.1,

$$\begin{aligned}
\text{dist}(\ell, p)^2 &\geq \text{dist}(\ell_q, p)^2 - 1/7d \\
&= 12\alpha^2 - 1/7d \\
&= 12/(12 - 4/d)r_d^2 - 1/7d \\
&\geq (1 + 1/3d)r_d^2 - 1/7d \\
&\geq r_d^2 + 1/6d - 1/7d \\
&= r_d^2 + 1/42d.
\end{aligned}$$

Case 2. Suppose that $||z_i| - 1/\sqrt{d}| \geq 1/14d^2$. It is easy to see that there must be a j

such that $z_j^2 \leq 1/d - 1/196d^3$. Now the square of the distance $\text{dist}(\ell, e_j)^2$ of the point e_j from ℓ equals $\|e_j\|^2 - \langle e_j, z \rangle^2 = 1 - z_j^2$, which is at least $r_d^2 + 1/196d^3$.

Thus, $R_1^{\text{opt}}(P_1 \cup P_2)^2 \geq r_d^2 + \min(1/42d, 1/196d^3) = r_d^2 + 1/196d^3 \geq r_d^2(1 + 1/196d^3)$ for sufficiently large d . We conclude that $R_1^{\text{opt}}(P_1 \cup P_2) \geq r_d(1 + 1/792d^3)$ for sufficiently large d .

Let $k = |P_1 \cup P_2|$. Since $k > d$, we have that $R_1^{\text{opt}}(P_1 \cup P_2) \geq r_d(1 + 1/792k^3)$ when φ is not satisfiable, and $R_1^{\text{opt}}(P_1 \cup P_2) = r_d$ when φ is satisfiable.

Theorem 4.2. *Unless $P = NP$, there is no polynomial time algorithm that, given a point set P in \mathbb{R}^d with k points, returns a number between $R_1^{\text{opt}}(P)$ and $(1 + 1/792k^3)R_1^{\text{opt}}(P)$.*

Corollary 4.3. *Unless $P = NP$, there is no algorithm that, given a set of n points in \mathbb{R}^d and any $0 < \varepsilon \leq 1$, returns a number r such that $R_1^{\text{opt}}(P) \leq r \leq (1 + \varepsilon)R_1^{\text{opt}}(P)$ and has a running time that is bounded by a polynomial in n , d , and $1/\varepsilon$.*

5. Minimum Radius k -Flat

In this section we describe an efficient algorithm that, given a set P of n points in \mathbb{R}^d and a $0 < \varepsilon \leq 1$, returns a k -flat \mathcal{F} such that $\mathcal{RD}(\mathcal{F}, P) \leq (1 + \varepsilon)R_k^{\text{opt}}(P)$. We first extend, in Section 5.1, the basic machinery used for the case of a line to the more general case of a k -flat. In Section 5.2, we use this machinery in a straightforward fashion to achieve an approximation algorithm for the optimal k -flat.

5.1. Preliminaries

Let $\mathcal{F}^{\text{opt}} = \mathcal{F}_k^{\text{opt}}(P)$ denote a k -flat that realizes $R_k^{\text{opt}}(P)$; there can be more than one such flat but we fix an arbitrary one for the rest of this section. For any $p \in P$, let p' denote $\text{proj}(\mathcal{F}^{\text{opt}}, p)$; let $P' = \text{proj}(\mathcal{F}^{\text{opt}}, P)$. Let $\mathcal{E}_k^{\text{opt}}(P)$ denote the minimum-volume ellipsoid on $\text{affine}(P')$ that encloses P' ; if $\text{affine}(P')$ is j -dimensional, we are speaking here of the j -dimensional volume [GLS]. Let $\mathcal{B}_k^{\text{opt}}(P)$ denote the minimum bounding box (on $\text{affine}(P')$), with axes aligned with those of $\mathcal{E}_k^{\text{opt}}(P)$, that encloses $\mathcal{E}_k^{\text{opt}}(P)$. For the case $k = 1$, the set $\mathcal{I} = \mathcal{CH}(P')$ was a line segment on the optimal line. For $k > 1$, the convex hull $\mathcal{CH}(P')$ does not necessarily look so simple, so the enclosing box $\mathcal{B}_k^{\text{opt}}(P)$ will play the role of \mathcal{I} in the proofs here.

For a point $x = (x_1, \dots, x_d) \in \mathbb{R}^d$ and a real number α , let αx denote the point $(\alpha x_1, \dots, \alpha x_d)$. For a set $X \subseteq \mathbb{R}^d$, let αX denote the set $\{\alpha x \mid x \in \mathbb{R}^d\}$. We say that a compact, convex set $X \subseteq \mathbb{R}^d$ is *centrally symmetric*, or just *symmetric*, if there exists a point $x \in X$ such that for any $y \in \mathbb{R}^d$, $x + y \in X$ if and only if $x - y \in X$. We call such a point x the *center* of X . It is easy to see that a compact, convex, symmetric set has a unique center. For a compact, convex, symmetric set X with center x and a given $\alpha \geq 0$, the *concentric scaling* of X by α is the set $\{x + \alpha(y - x) \mid y \in X\}$. The following lemma is the reason we use $\mathcal{B}_k^{\text{opt}}(P)$ and not any bounding box for P' .

Lemma 5.1. *Let \mathcal{G} be any flat, and let $d : \mathbb{R}^d \rightarrow \mathbb{R}$ be the distance function $d_{\mathcal{G}}(\cdot)$.*

- (i) *Let \mathcal{B}_{in} and \mathcal{B}_{out} be compact, convex, symmetric bodies centered at $v \in \mathbb{R}^d$ such that \mathcal{B}_{out} is the concentric scaling of \mathcal{B}_{in} by a factor of $\alpha > 1$. We have*

$$\max_{x \in \mathcal{B}_{\text{out}}} d(x) \leq \alpha \max_{x \in \mathcal{B}_{\text{in}}} d(x).$$

- (ii) *Let P' and $\mathcal{B}_k^{\text{opt}}(P)$ be as defined above. We have $\max_{x \in \mathcal{B}_k^{\text{opt}}(P)} d(x) \leq k^3 \max_{p' \in P'} d(p')$.*
- (iii) *For any flat \mathcal{H} , we have $\max_{x \in \text{proj}(\mathcal{H}, \mathcal{B}_k^{\text{opt}}(P))} d(x) \leq k^3 \max_{p' \in P'} d(\text{proj}(\mathcal{H}, p'))$.*

Proof. (i) Let $z \in \mathcal{B}_{\text{out}}$ be a point such that $d(z) = \max_{x \in \mathcal{B}_{\text{out}}} d(x)$. If $z = v$, the claim is immediate. Otherwise, let ℓ be the line through v and z , and let its intersection with \mathcal{B}_{in} be the segment yw . We assume that $\|wz\| \leq \|yz\|$. It is easy to check that $\|wz\|/\|yw\| \leq (\alpha - 1)/2$. From Lemma 2.4, we see that

$$d(z) \leq d(w) + (\alpha - 1) \max(d(y), d(w)) \leq \alpha \max(d(y), d(w)) \leq \alpha \max_{x \in \mathcal{B}_{\text{in}}} d(x).$$

(ii) Let $\mathcal{E}_k^{\text{opt}}(P)$ denote the min-enclosing ellipsoid of P' , and let \mathcal{E} be the concentric scaling of $\mathcal{E}_k^{\text{opt}}(P)$ by a factor of $1/k^2$. It follows from John's theorem [GLS, Section 4.6] that $\mathcal{E} \subseteq \mathcal{CH}(P') \subseteq \mathcal{E}_k^{\text{opt}}(P)$ (in fact, John's theorem is slightly stronger, but this does not matter for our discussion). Recall that $\mathcal{B}_k^{\text{opt}}(P)$ is the bounding box of $\mathcal{E}_k^{\text{opt}}(P)$ which is aligned with its main axes. Let \mathcal{B}_{in} be the concentric scaling of $\mathcal{B}_k^{\text{opt}}(P)$ by a factor of $1/k^3$. It is easy to argue that $\mathcal{B}_{\text{in}} \subseteq \mathcal{E}$ and so $\mathcal{B}_{\text{in}} \subseteq \mathcal{CH}(P')$. By part (i) above, we have

$$\max_{x \in \mathcal{B}_k^{\text{opt}}(P)} d(x) \leq k^3 \max_{x \in \mathcal{B}_{\text{in}}} d(x) \leq k^3 \max_{x \in \mathcal{CH}(P')} d(x) \leq k^3 \max_{p' \in P'} d(p'),$$

where the last inequality follows from the convexity established in Lemma 2.4(i).

(iii) We note that $\text{proj}(\mathcal{H}, \mathcal{B}_k^{\text{opt}}(P))$ is a concentric scaling by a factor of k^3 of $\text{proj}(\mathcal{H}, \mathcal{B}_{\text{in}})$ and that $\text{proj}(\mathcal{H}, \mathcal{B}_{\text{in}}) \subseteq \mathcal{CH}(\text{proj}(\mathcal{H}, P')) \subseteq \text{proj}(\mathcal{H}, \mathcal{B}_k^{\text{opt}}(P))$. The claim follows by an argument similar to that in (ii). \square

The statement of the following lemma is folklore but we include a proof here for completeness. The algorithm can be viewed as a version of the Gram–Schmidt orthogonalization carried out for k stages.

Lemma 5.2. *Given a set P of n points in \mathbb{R}^d , and a parameter k , we can compute in $O(ndk)$ time a k -flat \mathcal{F} , such that $\mathcal{RD}(\mathcal{F}, P) \leq 2^{k+1} R_k^{\text{opt}}(P)$.*

Proof. We may assume that $\text{affine}(P)$ has dimension at least $k + 1$; otherwise, we simply compute and return a k -flat that contains P . Let p_1 be an arbitrary point of P , and let $p_i \in P$ be the point such that $\mathcal{RD}(\mathcal{G}_{i-1}, P) = \text{dist}(\mathcal{G}_{i-1}, p_i)$, for $i = 2, \dots, k + 1$, where $\mathcal{G}_{i-1} = \text{affine}(p_1, \dots, p_{i-1})$. We set \mathcal{G}_{k+1} to be the flat \mathcal{F} that we return. Let \mathcal{F}_i be the k -flat that minimizes $\mathcal{RD}(\mathcal{F}, P)$ over all k -flats \mathcal{F} that contain \mathcal{G}_i . We claim that $\mathcal{RD}(\mathcal{F}_i, P) \leq 2^i R_k^{\text{opt}}(P)$. For $i = 1$, the claim is trivial, as we can just translate the optimal k -flat so that it passes through p_1 . We have $\mathcal{RD}(\mathcal{F}_1, P) \leq 2R_k^{\text{opt}}(P)$.

Assume the claim is correct for \mathcal{F}_{i-1} . Let \mathcal{F}'_i be the rotation $\text{Rot}(\mathcal{F}_{i-1}, \mathcal{G}_{i-1}, p_i)$ of \mathcal{F}_{i-1} around \mathcal{G}_{i-1} , so it passes through p_i . From Lemma 2.6, it follows that for any $q \in P$,

$$\begin{aligned} \text{dist}(\mathcal{F}'_i, q) &\leq \text{dist}(\mathcal{F}_{i-1}, q) + \frac{\text{dist}(\mathcal{G}_{i-1}, q)}{\text{dist}(\mathcal{G}_{i-1}, p)} \text{dist}(\mathcal{F}_{i-1}, p) \\ &\leq \text{dist}(\mathcal{F}_{i-1}, q) + \text{dist}(\mathcal{F}_{i-1}, p) \leq 2\mathcal{RD}(\mathcal{F}_{i-1}, P). \end{aligned}$$

This implies that $\mathcal{RD}(\mathcal{F}'_i, P) \leq 2\mathcal{RD}(\mathcal{F}_{i-1}, P)$, which implies that $\mathcal{RD}(\mathcal{F}_i, P) \leq 2\mathcal{RD}(\mathcal{F}_{i-1}, P)$. Thus $\mathcal{RD}(\mathcal{F}_{k+1}, P) \leq 2^{k+1} R_k^{\text{opt}}(P)$. However, $\mathcal{F}_{k+1} = \mathcal{G}_{k+1}$, since \mathcal{G}_{k+1} is itself a k -flat. It is clear that \mathcal{G}_{k+1} can be computed in $O(ndk)$ time. \square

The following is a generalization of Lemma 3.1 for $k > 1$.

Lemma 5.3. *Let P , \mathcal{F}^{opt} , $\mathcal{B}_k^{\text{opt}}(P)$ be as defined above. Given P , a $(k+1)$ -flat \mathcal{H} , a parameter $0 < \beta \leq 1$, and a number r such that $R_k^{\text{opt}}(P) \leq r \leq 2^{k+1} R_k^{\text{opt}}(P)$, we can compute in $nd \exp(O(k^3 \log(1/\beta)))$ time a family of $\exp(O(k^3 \log(1/\beta)))$ k -flats on \mathcal{H} such that at least one k -flat $\widehat{\mathcal{F}}$ in the family has the property that for any $x \in \mathcal{B}_k^{\text{opt}}(P)$,*

$$\text{dist}(\widehat{\mathcal{F}}, x) \leq \text{dist}(\mathcal{K}, x) + \beta R_k^{\text{opt}}(P),$$

where $\mathcal{K} = \text{proj}(\mathcal{H}, \mathcal{F}^{\text{opt}})$.

Proof. Let $\gamma = \beta / (2^{k+1} k^3 (k+1))$. We specify a set of sequences v_1, \dots, v_{k+1} of points on \mathcal{H} as follows. Let p_1 be any point in P , and let v_1 be chosen from a $(5r, \gamma r)$ -net around $\text{proj}(\mathcal{H}, p_1)$ on \mathcal{H} . For $2 \leq i \leq k+1$, we choose v_i given the choice of v_1, \dots, v_{i-1} as follows. Let $\mathcal{G}_{i-1} = \text{affine}(v_1, \dots, v_{i-1})$. Let p_i be the point in P whose projection onto \mathcal{H} is farthest from \mathcal{G}_{i-1} . We choose v_i from a $(5r, \gamma r)$ -net around $\text{proj}(\mathcal{H}, p_i)$ on \mathcal{H} .

For each choice of v_1, \dots, v_{k+1} , we add the k -flat $\text{affine}(v_1, \dots, v_{k+1})$ to our family. The claims in the lemma about the size of the family and the running time are readily verified.

We now argue that there is a sequence v_1^*, \dots, v_{k+1}^* such that $\widehat{\mathcal{F}} = \text{affine}(v_1^*, \dots, v_{k+1}^*)$ has the properties claimed in the lemma. We assume for simplicity that $\mathcal{K} = \text{proj}(\mathcal{H}, \mathcal{F}^{\text{opt}})$ is a k -flat; it is straightforward to extend this argument when it is a lower-dimensional flat. For any $p \in P$, let p' denote $\text{proj}(\mathcal{F}^{\text{opt}}, p)$. Let v_1^* be the point in the $(5r, \gamma r)$ -net around $\text{proj}(\mathcal{H}, p_1)$ that is closest to \mathcal{K} , and let \mathcal{F}_1 be the translation of \mathcal{K} through v_1^* . Since $\text{dist}(\mathcal{K}, \text{proj}(\mathcal{H}, p_1)) \leq \text{dist}(\mathcal{F}^{\text{opt}}, p_1) \leq r$, we have $\text{dist}(\mathcal{K}, v_1^*) \leq \gamma r$. It follows that for any $p \in P$, $\text{dist}(\mathcal{F}_1, \text{proj}(\mathcal{H}, p')) \leq \gamma r$.

Let $2 \leq i \leq k+1$ and suppose we have constructed a sequence v_1^*, \dots, v_{i-1}^* and a k -flat \mathcal{F}_{i-1} in \mathcal{H} containing v_1^*, \dots, v_{i-1}^* such that for any $p \in P$, $\text{dist}(\mathcal{F}_{i-1}, \text{proj}(\mathcal{H}, p')) \leq (i-1)\gamma r$. This implies that

$$\begin{aligned} \text{dist}(\mathcal{F}_{i-1}, \text{proj}(\mathcal{H}, p)) &\leq \text{dist}(\mathcal{F}_{i-1}, \text{proj}(\mathcal{H}, p')) + \|\text{proj}(\mathcal{H}, p) - \text{proj}(\mathcal{H}, p')\| \\ &\leq (i-1)\gamma r + r \leq 2r. \end{aligned}$$

Let $p_i^* \in P$ be the point whose projection $\text{proj}(\mathcal{H}, p_i^*)$ is farthest from $\mathcal{G}_{i-1}^* = \text{affine}(v_1^*, \dots, v_{i-1}^*)$. We choose v_i^* to be any point from the $(5r, \gamma r)$ -net around

$\text{proj}(\mathcal{H}, p_i^*)$ such that (1) $\text{dist}(\mathcal{F}_{i-1}, v_i^*) \leq \gamma r$ and (2) $\text{dist}(\mathcal{G}_i^*, v_i^*) \geq \text{dist}(\mathcal{G}_i^*, \text{proj}(\mathcal{H}, p_i^*)) + r$. It is not hard to see that such a v_i^* exists since $\text{dist}(\mathcal{F}_{i-1}, \text{proj}(\mathcal{H}, p_i^*)) \leq 2r$. Let $\mathcal{F}_i = \text{Rot}(\mathcal{F}_{i-1}, \mathcal{G}_i^*, v_i^*)$.

For any $p \in P$, we have

$$\begin{aligned} \text{dist}(\mathcal{G}_i^*, \text{proj}(\mathcal{H}, p')) &\leq \text{dist}(\mathcal{G}_i^*, \text{proj}(\mathcal{H}, p)) + r \\ &\leq \text{dist}(\mathcal{G}_i^*, \text{proj}(\mathcal{H}, p_i^*)) + r \\ &\leq \text{dist}(\mathcal{G}_i^*, v_i^*). \end{aligned}$$

From Lemma 2.6, we conclude that for any $p \in P$,

$$\begin{aligned} \text{dist}(\mathcal{F}_i, \text{proj}(\mathcal{H}, p')) &\leq \text{dist}(\mathcal{F}_{i-1}, \text{proj}(\mathcal{H}, p')) + \frac{\text{dist}(\mathcal{G}_i^*, \text{proj}(\mathcal{H}, p'))}{\text{dist}(\mathcal{G}_i^*, v_i^*)} \text{dist}(\mathcal{F}_{i-1}, v_i^*) \\ &\leq \text{dist}(\mathcal{F}_{i-1}, \text{proj}(\mathcal{H}, p')) + \text{dist}(\mathcal{F}_{i-1}, v_i^*) \leq i\gamma r. \end{aligned}$$

We conclude that for any $p \in P$, $\text{dist}(\widehat{\mathcal{F}}, \text{proj}(\mathcal{H}, p')) \leq (k+1)\gamma r$, where $\widehat{\mathcal{F}} = \mathcal{F}_{k+1} = \text{affine}(v_1^*, \dots, v_{k+1}^*)$. Applying Lemma 5.1(iii), we have that for any $x \in \mathcal{B}_k^{\text{opt}}(P)$,

$$\text{dist}(\widehat{\mathcal{F}}, \text{proj}(\mathcal{H}, x)) \leq k^3 \max_{p \in P} \text{dist}(\widehat{\mathcal{F}}, \text{proj}(\mathcal{H}, p')) \leq k^3(k+1)\gamma r \leq \beta R_k^{\text{opt}}(P),$$

which implies the lemma. \square

Definition 5.4. A sequence of distance functions d_0, \dots, d_v , where $d_i : \mathbb{R}^d \rightarrow \mathbb{R}$ is \mathcal{U} -bounded, if $\forall x \in \mathcal{B}_k^{\text{opt}}(P)$, we have $d_i(x) \leq \mathcal{U}$, for $i = 0, \dots, v$.

Lemma 5.5. Let d_0, \dots, d_v be a sequence of \mathcal{U} -bounded distance functions, let $0 < \varepsilon \leq 1$, and let $\mathcal{C}_\varepsilon = 3\varepsilon^2 R_k^{\text{opt}}(P)/(200\mathcal{U})$. Let $x \in \mathcal{B}_k^{\text{opt}}(P)$ be a point and let i be an integer, such that $d_i(x) \leq (1 - \varepsilon/3)d_{i-1}(x)$ and $d_{i-1}(x) \geq \varepsilon R_k^{\text{opt}}(P)$. Let \mathcal{B}_ε be the translate of $\mathcal{C}_\varepsilon \mathcal{B}_k^{\text{opt}}(P)$ centered at x . Then for any $z \in \mathcal{B}_\varepsilon$, we have (i) $d_{i-1}(z) \geq (\varepsilon/2)R_k^{\text{opt}}(P)$ and (ii) $d_i(z) \leq (1 - \varepsilon/5)d_{i-1}(z)$.

Proof. Consider any $z \in \mathcal{B}_\varepsilon$. Assume that $z \neq x$, for otherwise the claim is immediate. Let ℓ be the line through x and z and let the segment wy denote the intersection of ℓ with $2\mathcal{B}^{\text{opt}}$, the concentric scaling by 2 of $\mathcal{B}_k^{\text{opt}}(P)$. Observe that $\|xz\|/\|wy\| \leq \mathcal{C}_\varepsilon$. Indeed, consider the translate $\widehat{\mathcal{B}}$ of $\mathcal{B}_k^{\text{opt}}(P)$ centered at x , and observe that $\widehat{\mathcal{B}} \subseteq 2\mathcal{B}^{\text{opt}}$, by construction. Thus, $\|wy\| \geq \|\widehat{\mathcal{B}} \cap \ell\| \geq \|xz\|/\mathcal{C}_\varepsilon$.

From Lemma 5.1, we see that $d_{i-1}(w), d_{i-1}(y) \leq 2\mathcal{U}$. We obtain from Lemma 2.4 that

$$|d_{i-1}(z) - d_{i-1}(x)| \leq \frac{\|xz\|}{\|wy\|} 2 \max(d_{i-1}(w), d_{i-1}(y)) \leq 4\mathcal{C}_\varepsilon \mathcal{U}.$$

Consequently,

$$d_{i-1}(z) \geq d_{i-1}(x) - 4\mathcal{C}_\varepsilon \mathcal{U} = \varepsilon R_k^{\text{opt}}(P) - (12\varepsilon^2/200)R_k^{\text{opt}}(P) \geq (9\varepsilon/10)R_k^{\text{opt}}(P),$$

which implies (i). By a similar application of Lemma 2.4 to $d_i(\cdot)$, we get

$$\begin{aligned}
d_i(z) &\leq d_i(x) + 4C_\varepsilon \mathcal{U} \\
&\leq (1 - \varepsilon/3)d_{i-1}(x) + 4C_\varepsilon \mathcal{U} \\
&\leq (1 - \varepsilon/3)(d_{i-1}(z) + 4C_\varepsilon \mathcal{U}) + 4C_\varepsilon \mathcal{U} \\
&\leq (1 - \varepsilon/3)d_{i-1}(z) + 8C_\varepsilon \mathcal{U} \\
&\leq (1 - \varepsilon/3)d_{i-1}(z) + (2\varepsilon/15) \cdot (9\varepsilon R_k^{\text{opt}}(P)/10) \\
&\leq (1 - \varepsilon/3)d_{i-1}(z) + (2\varepsilon/15)d_{i-1}(z) \\
&\leq (1 - \varepsilon/5)d_{i-1}(z). \quad \square
\end{aligned}$$

5.2. The Algorithm

We compute, in $O(ndk)$ time, a k -flat \mathcal{F}_0 such that $\mathcal{RD}(\mathcal{F}_0, P) \leq 2^{k+1}R_k^{\text{opt}}(P)$, using the algorithm of Lemma 5.2. We compute a sequence of k -flats $\mathcal{F}_0, \dots, \mathcal{F}_\nu$, where

$$\nu = \frac{\exp(ck^2)}{\varepsilon^{2k+1}} \log \frac{1}{\varepsilon},$$

and c is a sufficiently large constant to be determined below. We describe below how the flat \mathcal{F}_i is computed from \mathcal{F}_{i-1} in the i th iteration.

In the i th iteration we compute p_i , the point of P farthest away from \mathcal{F}_{i-1} . Let $\widehat{\mathcal{F}}_i$ denote the projection of \mathcal{F}^{opt} onto the $(k+1)$ -flat \mathcal{H}_i containing \mathcal{F}_{i-1} and p_i . Using the algorithm of Lemma 5.3 (setting $r = \mathcal{RD}(\mathcal{F}_0, P)$), we compute a family of k -flats on \mathcal{H}_i such that at least one flat \mathcal{F} in the family has the property that for any $x \in \mathcal{B}_k^{\text{opt}}(P)$, $\text{dist}(\mathcal{F}, x) \leq \text{dist}(\widehat{\mathcal{F}}_i, x) + \delta R_k^{\text{opt}}(P)$, where $\delta = \varepsilon/4\nu$. We present the oracle with this family, and the oracle picks the flat from the family with the above property. It can do this by specifying $O(k^3 \log 1/\delta)$ bits. Let \mathcal{F}_i be the k -flat chosen by the oracle.

At the end of the ν th iteration, we return the best k -flat from the sequence $\mathcal{F}_0, \dots, \mathcal{F}_\nu$. That is, we return the flat \mathcal{F} from the sequence that minimizes $\mathcal{RD}(\mathcal{F}, P)$. We argue below that $\mathcal{RD}(\mathcal{F}, P) \leq (1 + \varepsilon)R_k^{\text{opt}}(P)$ for such a flat \mathcal{F} . We assume the contrary, that is, $\mathcal{RD}(\mathcal{F}_i, P) > (1 + \varepsilon)R_k^{\text{opt}}(P)$, for each $0 \leq i \leq \nu$. We derive a contradiction.

Proof of Correctness. Let $d_i(x)$ denote the distance $\text{dist}(\mathcal{F}_i, x)$ of a point $x \in \mathbb{R}^d$ from \mathcal{F}_i . For each $1 \leq i \leq \nu$ and any $x \in \mathcal{B}_k^{\text{opt}}(P)$, we have

$$d_i(x) \leq \text{dist}(\widehat{\mathcal{F}}_i, x) + \delta R^{\text{opt}} \leq d_{i-1}(x) + \delta R^{\text{opt}}.$$

In particular, this implies that $d_i(x) \leq d_0(x) + i\delta R^{\text{opt}} \leq d_0(x) + R^{\text{opt}}$. On the other hand, we have that for any $x \in \mathcal{B}_k^{\text{opt}}(P)$,

$$\begin{aligned}
d_0(x) &\leq k^3 \max_{p \in P} \text{dist}(\mathcal{F}_0, p') \\
&\leq k^3 \max_{p \in P} \text{dist}(\mathcal{F}_0, p) + \|pp'\| \\
&\leq k^3(2^{k+1} + 1)R^{\text{opt}}.
\end{aligned}$$

It follows that for $i = 1, \dots, \nu$ and $x \in \mathcal{B}_k^{\text{opt}}(P)$, $d_j(x) \leq (k^3(2^{k+1} + 1) + 1)R^{\text{opt}} \leq k^3 2^{k+2} R^{\text{opt}}$. Thus, d_1, \dots, d_ν is a $(k^3 2^{k+2} R_k^{\text{opt}}(P))$ -bounded sequence (see Definition 5.4).

We partition $\mathcal{B}_k^{\text{opt}}(P)$ into a grid \mathcal{G} , where each grid cell is a copy of $\mathcal{C}_\varepsilon \mathcal{B}_k^{\text{opt}}(P)$, where $\mathcal{C}_\varepsilon = 3\varepsilon^2/(200 \cdot 2^{k+2}k^3)$. Let S be the set of vertices of \mathcal{G} . Clearly, $|S| = O((200 \cdot 2^{k+2}k^3)/\varepsilon^2)^k = e^{O(k^2)}/\varepsilon^{2k}$. The set S has the property that any translate of $\mathcal{C}_\varepsilon \mathcal{B}_k^{\text{opt}}(P)$ centered at a point in \mathcal{B}^{opt} intersects some point in S . We say that a point $z \in S$ is *hit* in the i th iteration if $d_{i-1}(z) \geq (\varepsilon/2)R_k^{\text{opt}}(P)$ and $d_i(z) \leq (1 - \varepsilon/5)d_{i-1}(z)$.

Suppose that a point $z \in S$ has been hit m times until the j th iteration; we have

$$d_j(z) \leq (1 - \varepsilon/5)^m d_0(z) + \nu \delta R^{\text{opt}} \leq (1 - \varepsilon/5)^m k^3 2^{k+2} R^{\text{opt}} + \nu \delta R^{\text{opt}}.$$

Thus, for $m = O((k/\varepsilon) \log(1/\varepsilon))$, we have

$$d_j(z) \leq (\varepsilon/4)R^{\text{opt}} + \nu \cdot (\varepsilon/(4\nu)) \cdot R^{\text{opt}} \leq (\varepsilon/2)R^{\text{opt}}.$$

Thus, after z is hit $m = O((k/\varepsilon) \log(1/\varepsilon))$ times, it can never be hit again. Note that the bound on m is independent of ν .

We now argue that for each $1 \leq i \leq \nu$, some point of S is hit in the i th iteration. Using an argument identical to the one for lines, we conclude that $d_i(p'_i) \leq (1 - \varepsilon/3)d_{i-1}(p'_i)$ and $d_{i-1}(p'_i) \geq \varepsilon R^{\text{opt}}$. (Recall that $p'_i = \text{proj}(\mathcal{F}^{\text{opt}}, p_i)$.) Let \mathcal{B} denote the translate of $\mathcal{C}_\varepsilon \mathcal{B}_k^{\text{opt}}(P)$ centered at p'_i . As noted above, some point $z \in S$ lies in \mathcal{B} . Now Lemma 5.5 tells us that z is hit in the i th iteration.

We choose c large enough so that the number of iterations ν is larger than $m \cdot |S|$. Since a point from S is hit in each of the ν iterations, but each point in S is hit at most m times, we have a contradiction.

Removing the Oracle. The algorithm as we described it uses $O(k^3 \log(4\nu/\varepsilon))$ bits from the oracle in each iteration, and therefore $M = O(\nu k^3 \log(4\nu/\varepsilon)) = (e^{O(k^2)}/\varepsilon^{2k+1}) \log^2(1/\varepsilon)$ bits overall. To remove the dependence on the oracle, we simply try all possible strings of size M , and execute the algorithm on each of these strings. The overall running time of the resulting algorithm is $nd \cdot D_\varepsilon$, where $D_\varepsilon = \exp((e^{O(k^2)}/\varepsilon^{2k+1}) \log^2(1/\varepsilon))$. We therefore conclude:

Theorem 5.6. *Given a set P of n points in \mathbb{R}^d and a parameter $0 < \varepsilon \leq 1$, we can compute, in $n \cdot d \cdot \exp((e^{O(k^2)}/\varepsilon^{2k+1}) \log^2(1/\varepsilon))$ time, a k -flat \mathcal{F} such that $\mathcal{RD}(\mathcal{F}, P) \leq (1 + \varepsilon)R_k^{\text{opt}}(P)$.*

Remark 5.7. In estimating the running time, we have not explicitly kept track of the number of bits needed to represent the numbers involved and have assumed that all simple arithmetic operations can be performed in unit time. We note however that the maximum number of bits needed to represent any such number depends only on ε , k , and the maximum number of bits needed to represent any input number. In particular, the number of bits is independent of n and d . A careful scrutiny tells us that the bound on the running time claimed in Theorem 5.6 remains valid even if we take into account the bit complexity of the arithmetic operations.

6. Conclusions

We have presented a linear-time approximation algorithm for fitting low-dimensional flats in high dimensions. This is a substantial improvement over previous algorithms. Currently, the bottleneck in extending this approach to fitting a small number of flats seems to be that we do not know how to do this in low dimensions. Also, unlike previous algorithms, it cannot be used directly for fitting a k -flat in the presence of outliers. These are open questions for further research. In particular, there is a huge gap between our result, and the fastest algorithms for those more general problems [HV2]. Is this gap inherent to the problems, or just a by-product of our techniques? Namely, can outliers or multiple flats be handled in high dimensions in near-linear time?

Acknowledgments

The authors thank Ken Clarkson, Piotr Indyk, and Piyush Kumar for helpful discussions concerning the problems studied in this paper. We also thank the anonymous referees for their useful comments.

References

- [BC] M. Bădoiu and K. L. Clarkson. Optimal core-sets for balls. In *Proc. 14th ACM–SIAM Sympos. Discrete Algorithms*, pages 801–802, 2003.
- [BGKvL] H. L. Bodlaender, P. Gritzmann, V. Klee, and J. van Leeuwen. The computational complexity of norm-maximization. *Combinatorica*, 10:203–225, 1990.
- [BH] G. Barequet and S. Har-Peled. Efficiently approximating the minimum-volume bounding box of a point set in three dimensions. *J. Algorithms*, 38:91–109, 2001.
- [BHI] M. Bădoiu, S. Har-Peled, and P. Indyk. Approximate clustering via core-sets. In *Proc. 34th Annu. ACM Sympos. Theory Comput.*, pages 250–257, 2002.
- [Br] A. Brieden. Geometric optimization problems likely not contained in \mathbb{APX} . *Discrete Comput. Geom.*, 28(2):201–209, 2002.
- [Ed] H. Edelsbrunner. *Algorithms in Combinatorial Geometry*. EATCS Monographs on Theoretical Computer Science. Springer-Verlag, Heidelberg, 1987.
- [FKS] U. Faigle, W. Kern, and M. Streng. Note on the computational complexity of j -radii of polytopes in \mathbb{R}^n . *Math. Programming*, 73:1–5, 1996.
- [GK1] P. Gritzmann and V. Klee. Computational complexity of inner and outer j -radii of polytopes in finite-dimensional normed spaces. *Math. Programming*, 59:163–213, 1993.
- [GK2] P. Gritzmann and V. Klee. On the complexity of some basic problems in computational convexity: I, containment problems. *Discrete Math.*, 136:129–174, 1994.
- [GLS] M. Grötschel, L. Lovász, and A. Schrijver. *Geometric Algorithms and Combinatorial Optimization*, volume 2 of Algorithms and Combinatorics. Springer-Verlag, Berlin, 1st edition, 1988; 2nd edition, 1994.
- [HV1] S. Har-Peled and K. R. Varadarajan. Approximate shape fitting via linearization. In *Proc. 42nd Annu. IEEE Sympos. Found. Comput. Sci.*, pages 66–73, 2001.
- [HV2] S. Har-Peled and K. R. Varadarajan. Projective clustering in high dimensions using core-sets. In *Proc. 18th Annu. ACM Sympos. Comput. Geom.*, pages 312–318, 2002.
- [KMY] P. Kumar, J. S. B. Mitchell, and E. A. Yildirim. Fast smallest enclosing hypersphere computation. In *Proc. 5th Workshop Algorithm Engrg. Exper.*, 2003.
- [Man] O. L. Mangasarian. *Nonlinear Programming*. Classics in Applied Mathematics. SIAM, Philadelphia, PA, 1994.

- [Mat] J. Matoušek. *Lectures on Discrete Geometry*. Springer-Verlag, Berlin, 2002.
- [Me] N. Megiddo. On the complexity of some geometric problems in unbounded dimension. *J. Symbolic Comput.*, 10:327–334, 1990.
- [Ne] Y. Nesterov. Global quadratic optimization via conic relaxation. Technical report, Catholic University of Louvaine, Belgium, 1998.
- [NRT] A. Nemirovski, C. Roos, and T. Terlaky. On maximization of quadratic forms over intersection of ellipsoids with common center. *Math. Programming*, 86(3):463–473, 1999.
- [VVZ] K. R. Varadarajan, S. Venkatesh, and J. Zhang. Approximating the radii of point sets in high dimensions. In *Proc. 43th Annu. IEEE Sympos. Found. Comput. Sci.*, 2002.
- [YZ] Y. Ye and J. Zhang. An improved algorithm for approximating the radii of point sets. In *Proc. APPROX*, 2003.

Received May 27, 2003, and in revised form February 23, 2004. Online publication June 28, 2004.