Approximate Shape fitting via Linearization

Pankaj K. Agarwal* Sariel Har-Peled† Kasturi R. Varadarajan‡

01/04/2002 10:42

Abstract

Shape fitting is an important optimization problem in computational geometry. We present a general technique for solving such problems: For a given a point set P in $\mathbb{R}^d$, this technique can be used to ε-approximate: (i) the min-width annulus and shell that contains P, (ii) minimum width cylindrical shell containing P, (iii) diameter, width, minimum volume bounding box of P, and (iv) all the previous measures for the case the points are moving. The running time of the resulting algorithms is $O(n + 1/\varepsilon^c)$, where c is a constant that depends on the problem at hand.

Our new general technique enable us to solve those problems in a plug and play manner, and for several of those problems our results are considerably faster than what was previously known. An interesting consequence of this technique for the measures we consider is the existence of a small subset of P that approximates P with respect to any given measure.

1 Introduction

Shape fitting is a fundamental optimization problem in computational geometry. Given a point-set, one would like to find the shape (for example, hyperplane) that best fits the point set. Shape fitting has numerous usages in graphics (shape simplification, collision detection), learning, data-mining, metrology, compression, and geometric optimization.

In this paper, we consider the version of this problem where the shape fitting the input is defined by a (small) constant number of parameters, and the quality of a given shape is measured by the maximum "distance" of an input point from the shape. Such problems fall into the realm of geometric optimization, and numerous variants have been solved. However, no unified theory has evolved, and solutions are usually based on case by case analysis. If the shape is convex, then it is implied by previous work [Gru93, BH01] that using convex shape approximation techniques one can first approximate the given input by a small point-set, and solve the optimization problem on this sampled subset. Since

*Department of Computer Science; Duke University; pankaj@cs.uiowa.edu

†Department of Computer Science, DCL 2111; University of Illinois; 1304 West Springfield Ave.; Urbana, IL 61801; USA; sariel@uiuc.edu; http://www.uiuc.edu/~sariel/

‡ Department of Computer Science, University of Iowa, kvaradar@cs.uiowa.edu

convex shape approximation is (relatively) well understood, this results in fast and efficient approximation algorithms. We state this more formally in Section 5.1.

However, if we are interested in a shape which is not convex (for example, we would like to cover a point-set by a minimum width annulus), then the problem becomes much harder. Although such problems have attracted a lot of research (see below for references), most of the results rely on some additional assumptions on the input.

An interesting feature of most of those problems is that they can be linearized. Namely, they can be rewritten as an optimization problem over a set of linear constraints, restricted to a specific subset of space. Once the problem is linearized one can solve the problem using, essentially, brute-force approach. This technique usually yields an algorithm computing the exact solution in $O(n^c)$ time, where c is an appropriate constant that depends on the problem at hand. This is of course unsatisfying in practice, as value of n might be huge. However, we show in this paper that that once the problem is stated as a linearized optimization problem, it can be approximated efficiently. Linearization is quite powerful and can be applied to numerous problems, and thus our approximation technique is quite broad, as it can be applied to optimization problems of this type. A consequence of our technique is that for any of the shapes that we consider, there is a small subset of points that approximately resembles the original point set as far as that shape is concerned. The size of the subset depends only on the approximation factor and not on the size of the original point set. Such a result was not known for nonconvex shapes despite the fact that these have received considerable attention. In the following, we give more details about the problems to which we apply our general technique.

Cylindrical shell problem. Given a line ℓ in $\mathbb{R}^d$ and two real numbers $0 \leq r \leq R$, the *cylindrical shell* $\Sigma(\ell, r, R)$ is the closed region lying between the two co-axial cylinders of radii r and R with ℓ as their axis, i.e.,

$$\Sigma(\ell, r, R) = \{p \in \mathbb{R}^d \mid r \leq d(\ell, p) \leq R\},$$

where $d(\ell, p)$ is the Euclidean distance between the point p and line ℓ. The *width* of $\Sigma(\ell, r, R)$ is $R - r$.

In the *approximate cylindrical shell* problem, we are given a set P of n points and a parameter $\varepsilon > 0$, and we want to compute a cylindrical shell containing P whose width is at most $(1 + \varepsilon)$ times the width of the minimum-width cylindrical shell containing P.

This problem is motivated by applications in computational metrology, see [AAS00]. Agarwal et al. [AAS00] present an algorithm that computes the exact minimum-width cylindrical shell for a set of n points in $\mathbb{R}^3$ in $O(n^5)$ time. They also present an algorithm that runs in roughly $O(n^2)$ time and computes a shell whose width is at most 26 times the optimal.

For this problem, we give an $(1+\varepsilon)$-approximation algorithm that runs in $O(n + 1/\varepsilon^{cd^2})$ time in $\mathbb{R}^d$, where c is a constant. This is a significant improvement over the algorithm of Agarwal et al. [AAS00].

Spherical shell or annulus problem. Given a point $\mathbf{x}$ in $\mathbb{R}^d$ and two real numbers $0 \leq r \leq R$, the *spherical shell* $\sigma(\mathbf{x}, r, R)$ is the closed region lying between the two concentric spheres of radii r and R with $\mathbf{x}$ as their center, i.e.,

$$\sigma(\ell, r, R) = \{p \in \mathbb{R}^d \mid r \le d(\mathbf{x}, p) \le R\},$$

where $d(\mathbf{x}, p)$ is the Euclidean distance between the points p and line $\mathbf{x}$. The *width* of $\sigma(\mathbf{x}, r, R)$ is $R - r$.

In the *approximate annulus* problem, we are given a set P of n points and a parameter $\varepsilon > 0$, and we want to compute an annulus containing P whose width is at most $(1+\varepsilon)$ times the width of the minimum-width annulus containing P. This problem is also motivated by applications in computational metrology, see [AAS97, AS96, AST94, AGSS89, EFNN89, EGS86, GLR97, LL91, MSY97, PS85, Riv79, RLW91, RZ92, SY95, SJ99, YC97, Cha00, AAHS00] and has been well-studied. The best known exact algorithm runs in roughly $O(n^{3/2})$ time in $\mathbb{R}^2$, and in roughly $O(n^{3-1/19})$ time in $\mathbb{R}^3$.

For the approximate annulus problem, we obtain an algorithm that runs in $O(n+1/\varepsilon^{3d})$ time in $\mathbb{R}^d$. This is an improvement over the algorithm of Chan [Cha00] which runs in roughly $O(n + 1/\varepsilon^{d^2/4})$ time. The running time of our algorithm can be improved further; we mention the annulus result primarily to point out that very good running times can be obtained as a consequence of our very general technique.

Approximating measures of moving points. With the rapid advances in positioning systems, e.g., GPS, ad-hoc networks, and wireless communication, it is becoming increasingly feasible to track and record the changing position of continuously moving objects. These developments have raised a wide range of challenging geometric problems involving moving objects, including efficient data structures for answering proximity queries, for clustering, and for maintaining connectivity information (see [AEGH98, AGMV97, BGH97, AH01, HS01]).

Let $P = \{p_1, \ldots, p_n\}$ be a set of n points moving in $\mathbb{R}^d$. For a given time t, let $p_i(t) = (x_i^1(t), \ldots, x_i^d(t))$ denote the position of point p_i at time t. We will use $P(t)$ to denote the set P at time t. We say that the motion of P is *algebraic* of degree k if all $x_i^j(t)$ are polynomials of degree at most k. For simplicity, we will assume for most of this paper that $k = 1$, in which case the motion is *linear* and the points move along a straight line. Thus $p_i(t) = \mathbf{a}_i + t\mathbf{b}_i$, where $\mathbf{a}_i, \mathbf{b}_i \in \mathbb{R}^d$.

Let $\mu(S)$ denote any of the following measures of the point set S such as diameter, width, minimum-enclosing ball, minimum-volume bounding box of arbitrary orientation, minimum-width enclosing annulus. For any $\varepsilon > 0$, we say that a subset $Q \subseteq P$ ε-approximates P with respect to measure μ if at any time t, $(1-\varepsilon)\mu(P(t)) \le \mu(Q(t))$. We consider the question of upper-bounding the size of an ε-approximation of any set of linearly (algebraically) moving points. In particular, we are interested in a bound that depends only on ε and not an n. We also consider the algorithmic problem of computing such a small ε-approximation for a given P. Currently, the geometry of moving points is not well understood. The results in this paper for moving points demonstrates the flexibility of our approach, and also demonstrate that moving point-set has tractable approximate geometric behavior. This should have further applications in developing KDS (kinetic data-structures) for moving points.

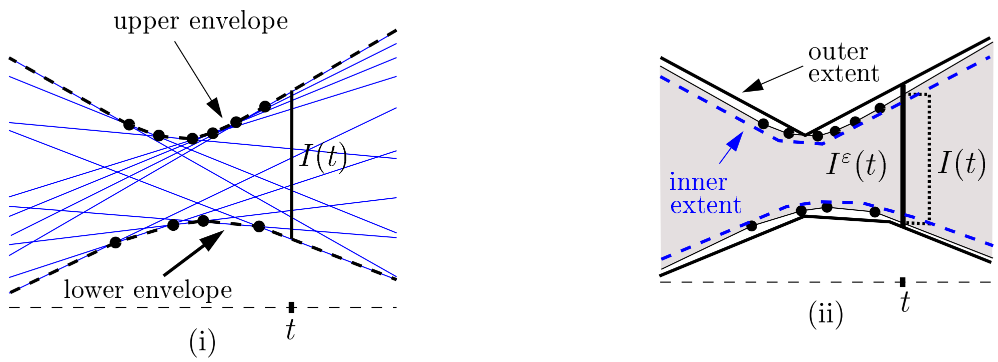

Figure 1: (i) The extent of linear functions, is no more than the vertical segment connecting the lower envelope to the upper envelope. The black dots mark where the description of the extent $I(t)$ changes. (ii) The approximate extent.

We show that for any set P of points in $\mathbb{R}^d$ with linear motion, one can compute, in $O(n + 1/\varepsilon^{2d})$ time an ε-approximation $Q \subseteq P$ of size $O(1/\varepsilon^{2d})$ with respect to all of the following measures: diameter, minimum-radius enclosing ball, width, minimum-volume bounding box of arbitrary orientation, projection width. We show that Q is an ε-approximation of P with respect to any "convex" measure. These results generalize to algebraic motion and to "non-convex" measures like minimum-width spherical/cylindrical shell. Our scheme can also allow insertions and deletions from the set P at a modest cost. These results must be contrasted with schemes for maintaining the exact extant measures. For example, a result of Agarwal et al. [AGHV97] shows that the pair of points realizing the diameter of n points in the plane with linear motion can change $\Omega(n^2)$ times.

In Section 2, we introduce the notion of the *extent* of a set $\mathcal{F}$ of functions. Informally, the extent of $\mathcal{F}$ is the difference between the upper and lower envelopes of $\mathcal{F}$. In Section 3, we show that if $\mathcal{F}$ is a set of linear functions there is a small subset $\mathcal{G} \subseteq \mathcal{F}$ such that the extent of $\mathcal{G}$ approximates the extent of $\mathcal{F}$. In Section 4, we show that this property is also true for polynomials and related functions. In Section 5, we apply these ideas to the problems mentioned above.

2 Preliminaries

Let $\mathcal{F} = \{f_1, \ldots, f_n\}$ be a set of k-variate functions defined over $\mathbf{x} = (x_1, \ldots, x_k) \in \mathbb{R}^k$. The *lower envelope* of $\mathcal{F}$ is the graph of the function $\mathfrak{L}_{\mathcal{F}} : \mathbb{R}^k \to \mathbb{R}$, namely $\mathfrak{L}_{\mathcal{F}}(\mathbf{x}) = \min_{f \in \mathcal{F}} f(\mathbf{x})$. Similarly, the *upper envelope* of $\mathcal{F}$ is the graph of the function $\mathcal{U}_{\mathcal{F}}(\mathbf{x}) = \max_{f \in \mathcal{F}} f(\mathbf{x})$.

The *extent* $I_{\mathcal{F}} : \mathbb{R}^k \to \mathbb{R}$ of $\mathcal{F}$ is defined as

$$I_{\mathcal{F}}(\mathbf{x}) = \mathcal{U}_{\mathcal{F}}(\mathbf{x}) - \mathfrak{L}_{\mathcal{F}}(\mathbf{x}).$$

Let $\varepsilon > 0$ be a parameter. We say that a subset $\mathcal{G} \subseteq \mathcal{F}$ is an (inner) ε-approximation of $\mathcal{F}$ if for each $\mathbf{x} \in \mathbb{R}^k$, we have

$$(1-\varepsilon)I_{\mathcal{F}}(\mathbf{x}) \leq I_{\mathcal{G}}(\mathbf{x}) \leq I_{\mathcal{F}}(\mathbf{x}).$$

One can also define a similar notion of outer extent, see Figure 1. Let $\Delta \subseteq \mathbb{R}^k$ be any region in $\mathbb{R}^k$. We say that a subset $\mathcal{G} \subseteq \mathcal{F}$ is an ε-approximation of $\mathcal{F}$ within Δ if for each $\mathbf{x} \in \Delta$,

$$(1-\varepsilon)I_{\mathcal{F}}(\mathbf{x}) \leq I_{\mathcal{G}}(\mathbf{x}) \leq I_{\mathcal{F}}(\mathbf{x}).$$

A similar concept to the extent can be defined for a set of points: For any set P of points in $\mathbb{R}^d$ and any $u \in \mathbb{R}^d$, we let $\omega(u, P)$ denote the quantity

$$\max_{p \in P} \langle u, p \rangle - \min_{p \in P} \langle u, p \rangle.$$

Notice that when u is a unit vector, $\omega(u, P)$ is simply the *width* of P in the direction u. (This is also known as the *u-breadth* of P, see [GK92].) Let $\varepsilon > 0$ be a parameter. We say that a subset $Q \subseteq P$ is an ε-approximation of P if for each $u \in \mathbb{R}^d$,

$$(1-\varepsilon)\omega(u, P) \leq \omega(u, Q).$$

Let $\Delta \subseteq \mathbb{R}^d$ be any set of vectors. We say that a subset $Q \subseteq P$ is an ε-approximation of P within Δ if for each $u \in \Delta$,

$$(1-\varepsilon)\omega(u, P) \leq \omega(u, Q).$$

A *partition* or *decomposition* of $\mathbb{R}^d$ is a collection $\mathcal{C} = \{\Delta_1, \ldots, \Delta_\ell\}$ of closed subsets (or *regions*) of $\mathbb{R}^d$ such that (a) the union of the regions is $\mathbb{R}^d$, and (b) the interiors of any two regions are disjoint. A set $\mathcal{J}$ of hyperplanes naturally induces a decomposition $\mathcal{A}(\mathcal{J})$ of $\mathbb{R}^d$; the regions of this decomposition are the full-dimensional faces of the arrangement of $\mathcal{J}$.

Definition 2.1 We define a *simple region* in $\mathbb{R}^d$ to be any subset of $\mathbb{R}^d$ that is defined by an intersection of $2d+2$ half-spaces. We define a *simple cone* in $\mathbb{R}^d$ to be any subset of $\mathbb{R}^d$ defined by an intersection of $2d$ half-spaces whose bounding hyperplanes contain the origin.

Lemma 2.2 *For any $\delta > 0$, there is a set $\mathcal{J}$ of $O(1/\delta)$ hyperplanes in $\mathbb{R}^d$ such that (1) each region in the decomposition $\mathcal{A}(\mathcal{J})$ induced by $\mathcal{J}$ is a simple cone, and (b) the angle between any two vectors in a single region is at most δ. Furthermore, the number of regions in $\mathcal{A}(\mathcal{J})$ is $O(1/\delta^{d-1})$.*

Proof: Let $e_i = (0, ..., 0, 1, 0, \ldots, 0)$ be the unit vector having 1 in its i cooridinate, for $i = 1, \ldots, k$, and let $h(i, j)$ be the hyperplane passing through the origin and through $e_1, \ldots e_{i-1}, e_{i+1}, \ldots, e_k$ and has angle $-\delta \cdot j/(k+1)$ with e_i. Finally, let

$$\mathcal{J} = \left\{ h(i,j) \;\middle|\; i = 1, \ldots, k, j = 0, \ldots, \lceil \pi(k+1)/\delta \rceil \right\}.$$

Clearly, each region in the decomposition $\mathcal{A}(\mathcal{J})$ is a simple cone. The angular diameter of each region is smaller than δ by the triangle inequality. ■

Extent of Linear Functions and Duality. Let $\mathcal{H} = \{h_1, \ldots, h_n\}$ be a family of k-variate linear functions and $\varepsilon > 0$ a parameter. Consider the *duality* transformation that maps the k-variate function $h = a_1x_1 + a_2x_2 + \cdots + a_kx_k + a_{k+1}$ to the point $h^* = (a_1, a_2, \ldots, a_k, a_{k+1})$ in $\mathbb{R}^{k+1}$. Let $\mathcal{H}^*$ denote the point set in $\mathbb{R}^{k+1}$ dual to the set $\mathcal{H}$ of k-variate linear functions. The following proposition is immediate.

Proposition 2.3 *Let $\mathcal{H} = \{h_1, \ldots, h_n\}$ be a family of k-variate linear functions and $\varepsilon > 0$ a parameter. If a subset $\mathcal{K}^* \subseteq \mathcal{H}^*$ is an ε-approximation of $\mathcal{H}^*$, then the corresponding set $\mathcal{K}$ is an ε-approximation of $\mathcal{H}$.*

3 Approximating the Extent

3.1 Approximating Fat Polytopes

In this section, we present algorithms for approximating a fat poytope efficiently.

Definition 3.1 A set of points $P \in \mathbb{R}^d$ is *c-fat* if P is contained inside a unit cube and $\mathcal{CH}(P)$ contains a hypercube of length c.

Proposition 3.2 *Let P be a point set that is c_d-fat, where c_d is a constant that depends only on d. For $\varepsilon > 0$, let P' be a point set with the property that for any $p \in P$, there is a $p' \in P'$ such that $d(p, p') \leq \varepsilon c_d$. Then for any $u \in \mathbb{R}^d$, $(1 - 2\varepsilon)\omega(u, P) \leq \omega(u, P')$.*

Proof: It suffices to prove the proposition when u is a unit vector, since the general result follows by scaling. Let u be any unit vector. Let $p, q \in P$ be two points such that $\omega(u, P) = \omega(u, \{p, q\})$, and let $p', q' \in P'$ be two points such that $d(p, p'), d(q, q') \leq \varepsilon c_d$. Since u is a unit vector, $\omega(u, P) \geq c_d$.

Let $w = p - q$ and $w' = p' - q'$, so that $\omega(u, \{p, q\}) = |\langle w, u\rangle|$ and $\omega(u, \{p', q'\}) = |\langle w', u\rangle|$.

$$
\begin{aligned}
\omega(u, P) - \omega(u, P') &\leq \omega(u, \{p, q\}) - \omega(u, \{p', q'\}) = |\langle w, u\rangle| - |\langle w', u\rangle| \\
&\leq |\langle w - w', u\rangle| \leq ||w - w'|| * ||u|| = ||w - w'|| \\
&= ||(p - p') + (q' - q)|| \leq ||p - p'|| + ||q' - q|| \leq 2\varepsilon c_d \\
&\leq 2\varepsilon\omega(u, P).
\end{aligned}
$$

■

Lemma 3.3 *Let P be a point set that is c_d-fat, where c_d is a constant that depends only on d. For any $\varepsilon > 0$, we can compute, in $O(n + 1/\varepsilon^{d-1})$ time, a subset $Q \subseteq P$ of $O(1/\varepsilon^{d-1})$ points that ε-approximates P.*

Proof: We consider the grid of size $\frac{\varepsilon}{4\sqrt{d}}c_d$. In each column of grid cells parallel to the x_d axis, we compute the lowest and highest cell containing some point from P. We let Q denote the set obtained by picking from each such cell any point from P contained in the cell.

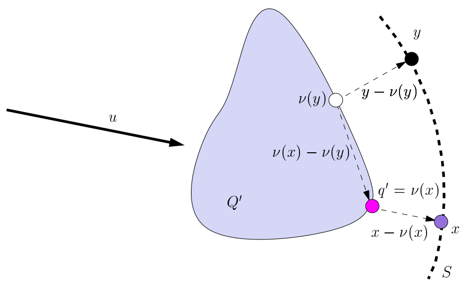

Figure 2: Illustation of the proof of Lemma 3.4.

We now argue that Q ε-approximates P. We conceptually move each point in P to the center of the grid cell containing the point. Let P' denote the new point set. Since each point moved by at most $\frac{\varepsilon}{4}c_d$, Proposition 3.2 implies that $(1-\varepsilon/2)\omega(u,P) \leq \omega(u,P')$ for any $u \in \mathbb{R}^d$. From each line parallel to the x_d axis, we discard any point from P' that is not either the highest or lowest point on P'. Let P'' be the resulting set of points; it is clear that for any $u \in \mathbb{R}^d$, $\omega(u,P') = \omega(u,P'')$. We observe that for any $p'' \in P''$, there is a $q \in Q$ such that $d(p'',q) \leq \varepsilon/4$. Applying Proposition 3.2 once again, we see that for any $u \in \mathbb{R}^d$,

$$\begin{aligned} \omega(u,Q) &\geq (1-\varepsilon/2)\omega(u,P'') = (1-\varepsilon/2)\omega(u,P') \\ &\geq (1-\varepsilon/2)^2\omega(u,P) \geq (1-\varepsilon)\omega(u,P). \end{aligned}$$ ■

Lemma 3.4 *Let P be a point set that is c_d-fat, where c_d is a constant that depends only on d. For any $\varepsilon > 0$, we can compute, in $O(n + 1/\varepsilon^{3(d-1)/2})$ time, a subset $Q \subseteq P$ of $O(1/\varepsilon^{(d-1)/2})$ points that ε-approximates P.*

Proof: The algorithm is based on a scheme of Dudley [Dud74]. For the sake of completeness, we describe it here. Let S be the sphere of radius $\sqrt{d}+1$ centered at the center of the unit cube containing P. Notice that the distance between any point on the sphere and any point within the unit cube is at least 1. Using Lemma 3.3, we compute a set $Q' \subseteq P$ of $O(1/\varepsilon^{d-1})$ points that ε-approximates P. Let $\delta = \sqrt{\varepsilon c_d/4}$. We compute a set I of $O(1/\delta^{d-1}) = O(1/\varepsilon^{(d-1)/2})$ points on the sphere S such that for any point x on sphere S, there is a point $y \in I$ such that $||x-y|| \leq \delta$. For each point $y \in I$, we then compute the point $\nu(y)$ on $\mathcal{CH}(Q')$ that is closest to y. Using the algorithm of Gartner [Gär95], this can be done for each y in time proportional to the size of Q'. Gartner's algorithm in fact returns a subset $b(y) \subseteq Q'$ of at most d points such that $\nu(y)$ is in the convex hull of $b(y)$. Let Q denote $\cup_{y\in B} b(y)$. It is easy to verify that Q has $O(1/\varepsilon^{(d-1)/2})$ points and takes $O(1/\varepsilon^{3(d-1)/2})$ time to compute. We now argue that the set Q ε-approximates P.

Let u be any unit vector, and let $q' \in Q'$ be the point that maximizes the $\langle u, q'\rangle$. Let $x \in B$ be the point on S which the ray shot form q' in direction u hits S. It is easy to see that q' is the unique point on $\mathcal{CH}(Q')$ nearest to x; that is, $q' = \nu(x)$. Also, the vector $x - \nu(x)$ has the same direction as u and $||x-\nu(x)|| \geq ||u|| = 1$. Let $y \in Y$ be such that

$||x - y|| \leq (\delta)^{1/2}$. Since $\nu(x)$ is the closest point to x in $\mathcal{CH}(Q')$,

$$\langle x - \nu(x), \nu(x) - \nu(y) \rangle \geq 0,$$

see Figure 2. Similarly,

$$\langle y - \nu(y), \nu(y) - \nu(x) \rangle \geq 0.$$

Therefore, we have

$$\begin{aligned}
\max_{q \in Q'} \langle u, q \rangle - \max_{q \in Q} \langle u, q \rangle &\leq \langle u, q' \rangle - \langle u, \nu(y) \rangle = \langle u, \nu(x) - \nu(y) \rangle \leq \langle x - \nu(x), \nu(x) - \nu(y) \rangle \\
&\leq \langle x - \nu(x) - (y - \nu(y)), \nu(x) - \nu(y) \rangle,
\end{aligned}$$

as $\langle -(y - \nu(y)), \nu(x) - \nu(y) \rangle \geq 0$. Setting $a = x - y$ and $b = \nu(x) - \nu(y)$, we have

$$\begin{aligned}
\max_{q \in Q'} \langle u, q \rangle - \max_{q \in Q} \langle u, q \rangle &\leq \langle a - b, b \rangle = \langle a, b \rangle - ||b||^2 \leq ||a||^2 - \langle a, b \rangle \\
&\leq ||a||^2 = ||x - y||^2 \leq \delta^2 = c_d \varepsilon / 4,
\end{aligned}$$

as $2\langle a, b \rangle \leq ||a||^2 + ||b||^2$. Similarly, we have

$$\min_{q \in Q} \langle u, q \rangle - \min_{q \in Q'} \langle u, q \rangle \leq c_d \varepsilon / 4,$$

which implies that

$$\omega(u, Q') - \omega(u, Q) \leq c_d \frac{\varepsilon}{2} \leq \frac{\varepsilon}{2} \omega(u, P).$$

Since Q' is an $\varepsilon/2$-approximation of P, we have

$$\omega(u, P) - \omega(u, Q') \leq \frac{\varepsilon}{2} \omega(u, P).$$

Adding the two inequalities, we get

$$(1 - \varepsilon) \omega(u, P) \leq \omega(u, Q),$$

implying that Q is an ε-approximation of P. ■

Lemma 3.5 *Let P be a point set that is c_d-fat, where c_d is a constant that depends only on d. For any $\varepsilon > 0$, we can compute, in $O(n + 1/\varepsilon^{3(d-1)/2})$ time, a partition $\mathcal{C}'$ of $\mathbb{R}^d$ with $O(1/\varepsilon^{d-1})$ regions, such that for each for region $\Delta' \in \mathcal{C}'$ there are two points $p'_\Delta, p''_\Delta \in P$ such that $\{p'_\Delta, p''_\Delta\}$ ε-approximates P inside Δ.*

Furthermore, each region in $\mathcal{C}'$ is a simple cone, and $\mathcal{C}' = \mathcal{A}(\mathcal{J})$, the decompostion induced by a set $\mathcal{J}$ of $O(1/\varepsilon)$ hyperplanes passing through the origin.

Proof: We first use Lemma 3.4 to compute a subset Q of $O(1/\varepsilon^{(d-1)/2})$ points that $\varepsilon/2$-approximates P. We compute a set J of $O(1/\varepsilon)$ hyperplanes, using Lemma 2.2, such that the partition $\mathcal{C}' = \mathcal{A}(\mathcal{J})$ is made of $O(1/\varepsilon^{d-1})$ simple cones such that for any two unit vectors u, v in the same simple cone, $||u - v|| \leq \frac{\varepsilon}{4\sqrt{d}} c_d$.

From each simplicial cone Δ, we pick any unit vector u_Δ, and find the points p'_Δ and p''_Δ that achieve $\max_{q \in Q} \langle u_\Delta, q \rangle$ and $\min_{q \in Q} \langle u_\Delta, q \rangle$, respectively. We can do this by considering each point in Q. We associate the points p'_Δ and p''_Δ with the simplicial cone Δ.

It is easy to see that running time of the above algorithm is as claimed. We now argue that for each simplicial cone $\Delta \in \mathcal{C}'$, the set $\{p'_\Delta, p''_\Delta\}$ $\varepsilon/2$-approximates Q within Δ. This is sufficient because Q $\varepsilon/2$-approximates P. Let $u = u_\Delta$, $p' = p''_\Delta$, $p'' = p''_\Delta$. Let $v \in \Delta$ be another unit vector, and let q' and q'' be the points in Q that achieve $\max_{q \in Q} \langle v, q \rangle$ and $\min_{q \in Q} \langle v, q \rangle$, respectively. Since Q lies in a unit cube, the lengths of the vectors $p' - p''$ and $q' - q''$ is at most $\sqrt{d}$.

$$\begin{aligned}
\langle p' - p'', v \rangle &= \langle p' - p'', u \rangle + \langle p' - p'', v - u \rangle \geq \langle p' - p'', u \rangle - |\langle p' - p', v - u \rangle| \\
&\geq \langle p' - p'', u \rangle - \frac{\varepsilon}{4\sqrt{d}} c_d * \sqrt{d} = \langle p' - p'', u \rangle - \frac{\varepsilon}{4} c_d \\
&\geq (1 - \varepsilon/4) \langle p' - p'', u \rangle
\end{aligned}$$

Similarly, $\langle q' - q'', u \rangle \geq (1 - \varepsilon/4) \langle q' - q'', v \rangle$. By the definition of p' and p'', we know that $\langle q' - q'', u \rangle \leq \langle p' - p'', u \rangle$. Combining these inequalities, we get

$$(1 - \varepsilon/4) \langle q' - q'', v \rangle \leq \langle q' - q'', u \rangle \leq \langle p' - p'', u \rangle \leq (1 - \varepsilon/4)^{-1} \langle p' - p'', v \rangle,$$

which yields $(1 - \varepsilon/2) \langle q' - q'', v \rangle \leq \langle p' - p'', v \rangle$. This completes the proof. ■

3.2 Approximating the extent

Theorem 3.6 *Given a family of k-variate linear functions $\mathcal{H} = \{h_1, \ldots, h_n\}$, and a parameter $\varepsilon > 0$, one can compute, in $O(n + 1/\varepsilon^k)$ time, a subset $\mathcal{K} \subseteq \mathcal{H}$ of $O(1/\varepsilon^k)$ linear functions, such that $\mathcal{K}$ is an ε-approximation for $\mathcal{H}$.*

Proof: By Proposition 2.3, it suffices to compute an ε-approximation for the point set $\mathcal{H}^*$. Using the algorithm of Barequet and Har-Peled [BH01], we compute in $O(n)$ time two boxes B' and B such that (a) B is obtained from B' by scaling by a factor of at most a_d, a constant that depends only on d, (b) $B' \subseteq \mathcal{CH}(\mathcal{H}^*) \subseteq B$. Let λ be the scaling transformation that maps B to a unit cube C. (For simplicity of exposition, we assume that the point-set $\mathcal{H}^*$ is $k+1$-dimensional; the technique works even if $\mathcal{H}^*$ has smaller dimension.)

Notice that $\mathcal{CH}(\lambda(\mathcal{H}^*))$ contains a hyper-cube of volume at least $1/a_d$. We apply the algorithm of Lemma 3.3 to compute, in $O(n + 1/\varepsilon^k)$ time, an ε-approximation $\lambda(\mathcal{K}^*)$ of $\lambda(\mathcal{H}^*)$ with $O(1/\varepsilon^k)$ points. We return K^* as our ε-approximation of $\mathcal{H}^*$.

We now argue the correctness of this scheme. Observe that since the matrix of the transformation λ is a diagonal matrix and therefore symmetric, $\langle u, p \rangle = \langle \lambda^{-1}(u), \lambda(p) \rangle$. From this, it follows that for any $u \in \mathbb{R}^{k+1}$,

$$(1 - \varepsilon)\omega(u, \mathcal{H}^*) \leq \omega(u, \mathcal{K}^*) \Leftrightarrow (1 - \varepsilon)\omega(\lambda^{-1}(u), \lambda(\mathcal{H}^*)) \leq \omega(\lambda^{-1}(u), \lambda(\mathcal{K}^*)).$$

Therefore $\mathcal{K}^*$ is an ε-approximation of $\mathcal{H}^*$ if and only if $\lambda(\mathcal{K}^*)$ is an ε-approximation of $\lambda(\mathcal{H}^*)$. ■

Theorem 3.7 *Given a family of k-variate linear functions $\mathcal{H} = \{h_1, \ldots, h_n\}$, and a parameter $\varepsilon > 0$, one can compute, in $O(n + 1/\varepsilon^{3k/2})$ time, a subset $\mathcal{K} \subseteq \mathcal{H}$ of $O(1/\varepsilon^{k/2})$ linear functions, such that $\mathcal{K}$ is an ε-approximation for $\mathcal{H}$.*

Proof: We proceed as in Theorem 3.6 but use Lemma 3.4 instead of Lemma 3.3. ■

Theorem 3.8 *Given a family of k-variate linear functions $\mathcal{H} = \{h_1, \ldots, h_n\}$, and a parameter $\varepsilon > 0$, one can compute, in $O(n + 1/\varepsilon^{3k/2})$ time, a decomposition $\mathcal{C}$ of $\mathbb{R}^k$ into $O(1/\varepsilon^k)$ regions, so that for each region $\Delta \in \mathcal{C}$ there are two associated linear functions, $h'_\Delta, h''_\Delta \in \mathcal{H}$, that ε-approximate $\mathcal{H}$ inside Δ.*

Furthermore, each region of $\mathcal{C}$ is a simple region, and $\mathcal{C}$ is $\mathcal{A}(\mathcal{J})$, the decompostion induced by a set $\mathcal{J}$ of $O(1/\varepsilon)$ hyperplanes in $\mathbb{R}^k$.

Proof: We proceed as in Theorem 3.6 but use Lemma 3.5. Let $\mathcal{C}'$ be the decomposition of $\mathbb{R}^{k+1}$ produced by Lemma 3.5. Let β be the function that maps the point $(x_1, \ldots, x_k) \in \mathbb{R}^k$ to the point $(x_1, \ldots, x_k, 1) \in \mathbb{R}^{k+1}$. We compute, for each $\Delta' \in \mathcal{C}'$, a corresponding region Δ in $\mathbb{R}^k$; Δ is the pre-image under β of $\lambda(\Delta')$. If $\{\lambda(h_1^*), \lambda(h_2^*)\}$ are the points associated with Δ', we associate with Δ the two linear functions h_1 and h_2. Let $\mathcal{C}$ denote the collection of Δ's thus computed.

Because β is injective and λ is bijective, $\mathcal{C}$ is a decomposition of $\mathbb{R}^k$. Also observe that if $\{\lambda(h_1^*), \lambda(h_2^*)\}$ ε-approximates $\lambda(\mathcal{H}^*)$ inside region Δ', then $\{h_1^*, h_2^*\}$ ε-approximates $\mathcal{H}^*$ inside $\lambda(\Delta')$, and $\{h_1, h_2\}$ ε-approximates $\mathcal{H}$ inside Δ. Since each $\Delta' \in \mathcal{C}'$ is an intersection of at most $2(k+1)$ half-spaces in $\mathbb{R}^{k+1}$, it is clear that each $\Delta \in \mathcal{C}$ is an intersection of at most $2(k+1)$ half-spaces in $\mathbb{R}^k$ and is therefore a simple region. Furthermore, since the decomposition C' is induced by a set of $O(1/\varepsilon)$ hyperplanes in $\mathbb{R}^{k+1}$, C is induced by the corresponding set of hyperplanes in $\mathbb{R}^k$. ■

4 Polynomials and their Roots

Let $\mathcal{F} = \{f_1, \ldots, f_n\}$ be a family of k-variate polynomials, and $\varepsilon > 0$ be a parameter. We can use linearization to show that there is a subset $\mathcal{G} \subseteq \mathcal{F}$ of $O(1/\varepsilon^{k'})$ polynomials that ε-approximate $\mathcal{F}$. Here, k' is the number of distinct monomials that occur in the polynomials in $\mathcal{F}$.

Specifically, let $x_1, \ldots, x_k$ denote the variables over which the polynomials in $\mathcal{F}$ are defined. We map each monomial over $x_1, \ldots, x_k$ that occurs in $\mathcal{F}$ to a distinct variable u_i. Let $u_1, \ldots, u_{k'}$ be the resulting variables. Each polynomial $f_i \in \mathcal{F}$ now becomes a linear function h_i over $\mathbb{R}^{k'}$. Let $\mathcal{H} = \{h_1, \ldots, h_n\}$ be the resulting set of linear functions thus obtained. If $\mathcal{K} \subseteq \mathcal{H}$ is an ε-approximation for $\mathcal{H}$, then clearly the corresponding subset in $\mathcal{F}$ is an ε-approximation for $\mathcal{F}$. The following theorem thus follows from Theorem 3.6 and Theorem 3.7.

Theorem 4.1 *Given a family of k-variate polynomials $\mathcal{F} = \{f_1, \ldots, f_n\}$ and a parameter $\varepsilon > 0$, one can compute, in $O(n + 1/\varepsilon^{k'})$ time, a subset $\mathcal{G} \subseteq \mathcal{F}$ of $O(1/\varepsilon^{k'})$ polynomials, such that $\mathcal{G}$ is an ε-approximation for $\mathcal{F}$. Here k' is the number of different monomials present in the polynomials in $\mathcal{F}$.*

Alternatively, one can compute in $O(n + 1/\varepsilon^{3k'/2})$ time, a subset $\mathcal{G}' \subseteq \mathcal{F}$, such that $\mathcal{G}'$ is an ε-approximation for $\mathcal{F}$, and $|\mathcal{G}'| = O(1/\varepsilon^{k'/2})$.

Before we can state the analog of Theorem 3.8 for polynomials, we need the following definition. We say that a region in $\mathbb{R}^k$ has *constant description complexity* if it is defined by a constant number of polynomial inequalities (in $x_1, \ldots, x_k$) and the degree of these polynomials is also bounded by a constant.

Theorem 4.2 *Given a family of k-variate polynomials $\mathcal{F} = \{f_1, \ldots, f_n\}$ and a parameter $\varepsilon > 0$, one can compute, in $O(n + 1/\varepsilon^{3k'/2})$ time, a decomposition $\mathcal{C}$ of $\mathbb{R}^k$ into $O(1/\varepsilon^{k'})$ regions of constant description complexity, so that for each region $\Delta \in \mathcal{C}$ there are two associated polynomials, $f'_\Delta, f''_\Delta \in \mathcal{F}$, that ε-approximate $\mathcal{F}$ inside Δ. Here k' is the number of different monomials present in the polynomials in $\mathcal{F}$.*

4.1 Roots of Polynomials

We now consider the problem of ε-approximating a family of square-roots of polynomials. Note, that this is considerably harder than handling polynomials because square-roots of polynomials can not be directly linearized. In certain special cases this can be overcome by special consideration of the functions at hand [AAHS00, Cha00]. It turns out, however, that it is enough to $O(\varepsilon^2)$-approximate the functions inside the roots, and take the root of the resulting approximation (see Lemma 4.4 below).

Theorem 4.3 *Let $\mathcal{F} = \{(f_1)^{1/2}, \ldots, (f_n)^{1/2}\}$ be a family of k-variate functions (over $\mathbf{x} = (x_1, \ldots, x_k) \in \mathbb{R}^k$), where each f_i is a polynomial that is non-negative for every $\mathbf{x} \in \mathbb{R}^k$. Given any $\varepsilon > 0$, we can compute, in $O(n + 1/\varepsilon^{2k'})$ time, a subset $\mathcal{G} \subseteq \mathcal{F}$ of size $O(1/\varepsilon^{2k'})$, such that $\mathcal{G}$ is an ε-approximation for $\mathcal{F}$. Here k' is the number of different monomials present in the polynomials in $f_1, \ldots, f_n$.*

Alternatively, one can compute a set $\mathcal{G}' \subseteq \mathcal{F}$, in $O(n+1/\varepsilon^{3k'})$ time, that ε-approximates$\mathcal{F}$, so that $|\mathcal{G}'| = O(1/\varepsilon^{k'})$.

Proof: Let $\mathcal{F}^2$ denote the family $\{f_1, \ldots, f_n\}$. Using the algorithm of Theorem 4.1, we compute a δ-approximation $\mathcal{G}^2 \subseteq \mathcal{F}^2$ of $\mathcal{F}^2$, where $\delta = c\varepsilon^2$ for a sufficiently small constant c. Let $\mathcal{G} \subseteq \mathcal{F}$ denote the family $\{(f_i)^{1/2} | f_i \in \mathcal{G}^2\}$. Using Lemma 4.4 below, we can conclude that $\mathcal{G}$ is an ε-approximation for $\mathcal{F}$. The bounds on the size of $\mathcal{G}$ and the running time are easily verified. ■

Lemma 4.4 *Let $0 \le a \le A \le B \le b$, and $0 < \varepsilon \le 1$ be given parameters, so that $b - a \le (1 + \delta)(B - A)$, where $\delta = \varepsilon^2/16$. Then, $\sqrt{b} - \sqrt{a} \le (1 + \varepsilon/2)(\sqrt{B} - \sqrt{A})$.*

Proof: Clearly,

$$\sqrt{A}+\sqrt{B} \;\leq\; \sqrt{a}+\sqrt{A-a}+\sqrt{b} \leq \sqrt{a}+\sqrt{\delta b}+\sqrt{b} \leq (1+\sqrt{\delta})(\sqrt{a}+\sqrt{b}).$$

Namely, $\frac{\sqrt{A}+\sqrt{B}}{1+\sqrt{\delta}} \leq \sqrt{a}+\sqrt{b}$. On the other hand,

$$\begin{aligned}
\sqrt{b}-\sqrt{a} &= \frac{b-a}{\sqrt{b}+\sqrt{a}} \leq \frac{(1+\delta)(B-A)}{\sqrt{b}+\sqrt{a}} \leq (1+\delta)(1+\sqrt{\delta})\frac{B-A}{\sqrt{A}+\sqrt{B}} \\
&= (1+\varepsilon^2/16)(1+\varepsilon/4)(\sqrt{B}-\sqrt{A}) \\
&\leq (1+\varepsilon/2)(\sqrt{B}-\sqrt{A}).
\end{aligned}$$

■

Theorem 4.5 *Let $\mathcal{F} = \{(f_1)^{1/2}, \ldots, (f_n)^{1/2}\}$ be a family of k-variate functions (over $\mathbf{x} = (x_1, \ldots, x_k) \in \mathbb{R}^k$), where each f_i is a polynomial that is non-negative for every $\mathbf{x} \in \mathbb{R}^k$. Given any $\varepsilon > 0$, we can compute, in $O(n + 1/\varepsilon^{3k'})$ time, a decomposition $\mathcal{C}$ of $\mathbb{R}^k$ into $O(1/\varepsilon^{2k'})$ regions of constant description complexity, so that for each region $\Delta \in \mathcal{C}$ there are two associated functions, $f'_\Delta, f''_\Delta \in \mathcal{F}$, that ε-approximate $\mathcal{F}$ inside Δ. Here k' is the number of different monomials present in the polynomials in $f_1, \ldots, f_n$.*

Remark 4.6 In many applications, we can give a better bound for k' by being more careful with the linearization. In the proof of Theorem 4.1, any monomial that has the same coefficient in each polynomial f_i can be omitted during linearization (that is, it need not be assigned a variable u_i). Such a monomial does not affect the extent of $\mathcal{F}$. Thus k' is the number of monomials that do not have the same coefficients in all the polynomials in $\mathcal{F}$. This carries over to the statements in Theorem 4.2, Theorem 4.3, and Theorem 4.5. Agarwal and Matousek [AM94] describe an algorithm for finding the minimum number of monomials needed to linearize a family of polynomials, which works by finding a basis for the functions participating in the linearizion.

5 Applications

In this section, we give efficient algorithms for several problems by applying the results on approximating the extent in the previous section. Since these results in turn depend on convex-hull approximation, what we basically do is reduce these problems to convex-hull approximation. Convex hull approximation has been used before for shape fitting with convex shapes [BH01, Cha00], but not for non-convex shapes. Before giving the applications to non-convex shapes, we state the problem of convex shape fitting in a very general setting and describe a general scheme for solving these problems.

5.1 Convex Shape Fitting in a General Setting

Let P be a set of n points in $\mathbb{R}^d$ and let $\varepsilon > 0$ be a parameter. As defined in Section 2, a subset $Q \subseteq P$ is an ε-approximation of P if for each $u \in \mathbb{R}^d$,

$$(1-\varepsilon)\omega(u,P) \leq \omega(u,Q).$$

The following is a restatement of Theorem 3.6 and Theorem 3.8 in the dual setting. (These theorems are in fact proved in the dual setting.)

Theorem 5.1 *Given a set P of n points in $\mathbb{R}^d$, one can compute a set $Q \subseteq P$ of $O(1/\varepsilon^{d-1})$ points that ε-approximates P in $O(n + 1/\varepsilon^{d-1})$ time.*

Alternatively, one can compute a set $Q' \subseteq P$ of size $O(1/\varepsilon^{(d-1)/2})$ that ε-approximates P in $O(n + 1/\varepsilon^{3(d-1)/2})$ time.

Definition 5.2 A function $\mu(\cdot)$ defined over subsets of $\mathbb{R}^d$, is *convex measure*, if: (i) for any $S \subseteq \mathbb{R}^d$, $\mu(S) \geq 0$, and (ii) $(1 - c\varepsilon)\mu(S) \leq \mu(Q) \leq \mu(S)$ if Q ε-approximates S where c is a constant that depends on μ.

Examples of convex measures of point-sets are quite common, and include volume of $\mathcal{CH}(S)$, surface area of $\mathcal{CH}(S)$, width of S, diameter of S, min-radius cylinder containing S, min-radius ball containing S, minimum volume bounding box that contains S, etc.

Theorem 5.1 implies that any convex measure over a point-set $P \subseteq \mathbb{R}^d$ can be approximated by (i) find a small (i.e., of size $O(1/\varepsilon^d)$ or $O(1/\varepsilon^{d/2})$) subset $Q \subseteq P$ that ε-approximates P, and then apply any exact algorithm (or even approximation algorithm) to Q. This would result in an $O(\varepsilon)$-approximation to the given measure, and the running time would be of the type $O(n + 1/\varepsilon^c)$ where c is a constant that depends on the subroutine being used at the second stage of the algorithm.

For all of the measures mentioned, algorithms of the same (or better) running-time are already known [BH01, Cha00]. However, our scheme does not require us to carefully inspect the problem at hand. Furthermore, it holds for any convex-measure.

5.2 Minimum-width Cylindrical Shell

We are given a set P of n points in $\mathbb{R}^d$, and a parameter $\varepsilon > 0$. Let w^* denote the width of the thinnest cylindrical shell containing P. We wish to find a cylindrical shell containing P whose width is at most $(1 + \varepsilon)w^*$. Let $d(\ell, p)$ denote the distance between a point $p \in \mathbb{R}^d$ and a line $\ell \in \mathbb{R}^d$. Let

$$\mu(\ell, P) = \max_{p \in P} d(\ell, p) - \min_{p \in P} d(\ell, p).$$

In other words, $\mu(\ell, P)$ is the width of the smallest cylindrical shell with ℓ as axis that contains P. Thus, $w^* = \min_\ell \mu(\ell, P)$. Our algorithm finds a line ℓ^* such that $\mu(\ell^*, P) \leq (1 + \varepsilon)w^*$.

Let $p = (p_1, \ldots, p_d)$ be any point in P. We will use $\mathbf{p}$ for p when we want to think of p as a vector. Let us parameterized a line ℓ in $\mathbb{R}^d$ by $\{\mathbf{x} + t\mathbf{y} | t \in R\}$, where $\mathbf{x} = (x_1, \ldots, x_d) \in \mathbb{R}^d$ and $\mathbf{y} = (y_1, \ldots, y_d)$ ranges over the unit vectors in $\mathbb{R}^d$. Then $d(\ell, p)$, the distance of the line ℓ from point p, is the norm of the vector

$$((\mathbf{p} - \mathbf{x}) \cdot \mathbf{y})\mathbf{y} + \mathbf{x} - \mathbf{p}.$$

If we fix p, $d(\ell, p)$ is the square root of a polynomial in the variables $x_1, \ldots, x_d, y_1, \ldots, y_d$. The coefficients of this polynomial depend on the coordinates of p. This polynomial is always non-negative (being a sum of squares) and it has $O(d^2)$ monomials.

Using the algorithm of Theorem 4.3, we can therefore find a subset $Q \subseteq P$ of $O(1/\varepsilon^{O(d^2)})$ points such that for any line ℓ, $(1-\varepsilon)\mu(\ell, P) \leq \mu(\ell, Q)$. The running time of this algorithm is $O(n + 1/\varepsilon^{O(d^2)})$.

We then find a line ℓ^* that minimizes $d(\ell, Q)$. If we do this using a brute-force exact algorithm, we take $O(|Q|^{2d+2})$ time [AAS00]. We have

$$\mu(\ell^*, P) \leq \frac{1}{1-\varepsilon}\mu(\ell^*, Q) \leq \frac{1}{1-\varepsilon}w^*.$$

Theorem 5.3 *Given a set P of n points in $\mathbb{R}^d$, and a parameter $\varepsilon > 0$, we can find in $O(n + 1/\varepsilon^{cd^3})$ time a cylindrical shell containing P whose width is at most $(1+\varepsilon)$ times the width of the thinnest cylindrical shell containing P. Here c is a constant independent of d and n.*

We can improve the running time of this algorithm to $O(n + 1/\varepsilon^{cd^2})$ as follows. Instead of computing the set Q, we use the algorithm of Theorem 4.5 to compute a decomposition $\mathcal{C}$ of $\mathbb{R}^{2d}$ into $O(1/\varepsilon^{O(d^2)})$ regions of constant description complexity, so that for each region $\Delta \in \mathcal{C}$, there are two points $p'_\Delta, p''_\Delta \in P$ so that for any $\ell = \{\mathbf{x}, \mathbf{y} | \; ||\mathbf{y}|| = 1\} \in \Delta$,

$$(1-\varepsilon)\mu(\ell, P) \leq \mu(\ell, \{p'_\Delta, p''_\Delta\}).$$

Having computed this decomposition, we find, for each $\Delta \in \mathcal{C}$, the line that minimizes $\mu(\ell, \{p'_\Delta, p''_\Delta\})$. Since this is a constant-sized optimization problem, it can be solved in constant time. Let ℓ^* be the line that achieves

$$\min_{\Delta \in \mathcal{C}} \min_{\ell \in \Delta} \mu(\ell, \{p'_\Delta, p''_\Delta\}).$$

Then it follows that $\mu(\ell^*, P) \leq \frac{1}{1-\varepsilon}\mu(\ell^*, Q) \leq \frac{1}{1-\varepsilon}w^*$.

Remark 5.4 These algorithms improve over the work of Agarwal et al. [AAS00]. They presented an algorithm that runs in roughly $O(n^2)$ time and computes a shell whose width is at most 26 times the optimal in 3d.

Remark 5.5 Theorem 5.3 implies that given a set of n points in $\mathbb{R}^d$, one can find a set $Q \subseteq P$ of size $O(1/\varepsilon^{cd^2})$, such that the cylindrical width of Q is an ε-approximation to the cylindrical width of P. We believe that this result is of independent interest. Previous attempts to prove similar properties were considerably more involved, and ultimately failed in achieving this. See [AAHS00, AAS00].

5.3 Minimum-width annulus

We are given a set P of n points in $\mathbb{R}^d$, and a parameter $\varepsilon > 0$. Let w^* denote the width of the thinnest annulus (or spherical shell) containing P. We wish to find an annulus containing P whose width is at most $(1+\varepsilon)w^*$. Let $d(\mathbf{x}, p)$ denote the distance between two points $\mathbf{x}, p \in \mathbb{R}^d$. Let

$$\mu(\mathbf{x}, P) = \max_{p \in P} d(\mathbf{x}, p) - \min_{p \in P} d(\mathbf{x}, p).$$

In other words, $\mu(\mathbf{x}, P)$ is the width of the smallest annulus with $\mathbf{x}$ as center that contains P. Thus, $w^* = \min_{\mathbf{x} \in \mathbb{R}^d} \mu(\mathbf{x}, P)$. Our algorithm finds a point $\mathbf{x}^*$ such that $\mu(\mathbf{x}^*, P) \le (1+\varepsilon)w^*$.

Let $p = (p_1, \ldots, p_d)$ be a point in P. Let us parameterized the center by $\mathbf{x} = (x_1, \ldots, x_d) \in \mathbb{R}^d$. Then $d(\mathbf{x}, p)$, the distance of the center ℓ from point p, is the square root of

$$\sum_{1 \le i \le d} x_i^2 - 2 \sum_{1 \le i \le d} p_i x_i + \sum_{1 \le i \le d} p_i^2.$$

Thus $d(\mathbf{x}, p)$ is the square root of a polynomial in $x_1, \ldots, x_d$ which is always non-negative. The number of distinct monomials is $2d$, but the coefficients of only d of the monomials depend on the coordinates of p. Our algorithm for computing $\mathbf{x}^*$ is similar to the second algorithm for the cylindrical shell problem. (We apply Theorem 4.5 with $k' = d$.) We obtain the following result, which improves over the result of Chan [Cha00] in its dependence on ε. (Chan's algorithm runs in roughly $O(n + 1/\varepsilon^{d^2/4})$ time.)

Theorem 5.6 *Given a set P of n points in $\mathbb{R}^d$, and a parameter $\varepsilon > 0$, we can find in $O(n + 1/\varepsilon^{3d})$ time a spherical shell containing P whose width is at most $(1+\varepsilon)$ times the width of the thinnest spherical shell containing P.*

By being more careful, we can get an algorithm that runs in $O(n + 1/\varepsilon^{2d} \log 1/\varepsilon)$ time. Perhaps further improvements are possible by combining Chan's techniques with the results from here. The annulus theorem above demonstrates that good running times can be obtained as a consequence of our very general technique.

5.4 Convex Sampling of a Moving Point Set

Let $P = \{p_1, \ldots, p_n\}$ be a set of n points moving in $\mathbb{R}^d$. For a given time t, let $p_i(t) = (x_i^1(t), \ldots, x_i^d(t))$ denote the position of p_i at time t. We use $P(t)$ denote the set P at time t. We assume that the motion of P is linear, i.e., $p_i(t) = \mathbf{a}_i + \mathbf{b}_i t$, where $\mathbf{a}_i, \mathbf{b}_i \in \mathbb{R}^d$.

For a vector $\mathbf{x} \in \mathbb{R}^d$, and a point $p_i(t) \in P$, let $\lambda_i(t, \mathbf{x}) = \langle \mathbf{a}_i + \mathbf{b}_i t, \mathbf{x} \rangle$ denote the dot-product of $p_i(t)$ with $\mathbf{x}$. Let $\mu_P(t, \mathbf{x}) = \max_i \lambda_i(t, \mathbf{x}) - \min_i \lambda_i(t, \mathbf{x})$. We say that a subset $Q \subseteq P$ is a *convex ε-approximation* to P if for any $\mathbf{x} \in \mathbb{R}^d$ and $t \in \mathbb{R}$,

$$(1-\varepsilon)\mu_P(t, \mathbf{x}) \le \mu_Q(t, \mathbf{x}).$$

Note that if Q is a convex ε-approximation to P, then Q is a $c\varepsilon$-approximation of P (c is a constant) with respect to "convex" measures of P. For example, $(1 - c\varepsilon)\,\mathrm{diam}(P(t)) \leq \mathrm{diam}(Q(t)) \leq \mathrm{diam}(P(t))$. Thus Q is a $c\varepsilon$-approximation of P with respect to diameter, minimum radius enclosing ball, min-radius cylindrical shell, width, and the volume of arbitrarily oriented minimum-volume bounding box.

Thus, being able to compute a small convex ε-approximation to P implies that we can maintain those approximate measures for a moving point-set efficiently. Currently, no efficient algorithms are known for maintaining those *exact* measures efficiently for a moving point-set.

Theorem 5.7 *Given a point-set P with n linearly-moving points, and a parameter $\varepsilon > 0$, one can compute, in $O(n + 1/\varepsilon^{2d})$ time, a subset $Q \subseteq P$ of size $O(1/\varepsilon^{2d})$, such that Q is a convex ε-approximation to P.*

Alternatively, one can compute in $O(n + 1/\varepsilon^{3d})$ time, a subset $Q' \subseteq P$ of size $O(1/\varepsilon^{d})$, such that Q is a convex ε-approximation to P.

Proof: Let $\mathbf{x} = (x_1, \ldots, x_d)$. The function $\lambda_i(t, \mathbf{x}) = \langle \mathbf{a}_i + \mathbf{b}_i t, \mathbf{x} \rangle$ is a polynomial in the variables $x_1, \ldots, x_d, t$ whose coefficients depend on $\mathbf{a}_i$ and $\mathbf{b}_i$. This polynomial has $2d$ monomials. The theorem follows from Theorem 4.1. ■

This result generalizes to point sets with algebraic motion of constant degree. For a point set P with algebraic motion, we can also compute small ε-approximations with respect to "non-convex" measures like minimum-width annulus/cylindrical shell. The approach is similar to the ones we have described so far, and we omit the details.

5.5 Extent of Univariate Polynomials

Let $\mathcal{F} = \{f_1(x), \ldots, f_n(x)\}$ be a set of univariate polynomials of degree at most k. Clearly, one can use Theorem 4.1 to compute a subset $\mathcal{G} \subseteq \mathcal{F}$ that ε-approximates the extent of $\mathcal{F}$. This would yield a subset $\mathcal{G}$ of size $O(1/\varepsilon^k)$. However, we can get the following improved bound.

Theorem 5.8 *Let $\mathcal{F} = \{f_1(x), \ldots, f_n(x)\}$ be a set of one variable polynomials of degree at most k, and let $\varepsilon > 0$ be a parameter. One can compute in $O(n + 1/\varepsilon^k)$ time a subset $\mathcal{G} \subseteq \mathcal{F}$ of size $O(1/\varepsilon)$ that ε-approximates $\mathcal{F}$.*

Proof: We use linearization, where $\mu(x) = (x, x^2, \ldots, x^k)$ is the mapping. The set $\mathcal{F}$ turns into a set $\mathcal{H}$ of k-variate linear functions, and by Theorem 3.8 one can compute a decomposition $\mathcal{C}$ of $\mathbb{R}^k$ into $O(1/\varepsilon^k)$ simple regions, such that for each region $\Delta \in \mathcal{C}$ there are two functions that ε-approximate $\mathcal{H}$ within Δ. Clearly, the only regions of $\mathcal{C}$ which are relevent are those that intersect the curve $\gamma = \cup_x \mu(x)$.

We note $\mathcal{C}$ is the decomposition $\mathcal{A}(\mathcal{J})$ induced by a set $\mathcal{J}$ of $O(1/\varepsilon)$ hyperplanes. Each such hyperplane of $\mathcal{J}$ intersects γ a constant number of times, which immedately implies that γ visits at most $O(1/\varepsilon)$ regions of $\mathcal{A}(\mathcal{J})$. Furthermore, by a careful impelmentation of the algorithm of Theorem 3.8, it follows that we can first compute the set of hyperplanes

$\mathcal{J}$, the regions visited by γ, and then compute the two linear functions that ε-approximate $\mathcal{H}$ within each such visited region. This results in a subset $\mathcal{G} \subseteq \mathcal{F}$ of size $O(1/\varepsilon)$ that ε-approximates $\mathcal{F}$. The computation time is $O(n + 1/\varepsilon^k)$. ■

It may be possible to use this technique to improve the bounds in some instances where $\mathcal{F}$ is a collection of multi-variate polynomials. However, the improvements are unlikely to be as good as in the case of univariate polynomials. The case of univariate polynomials arises in the context of maintaining axis-parallel bounding boxes for moving points [AH01].

5.6 Insertions and Deletions

Let P be a set of points with algebraic motion. Let μ be a measure with the property that we can compute, given any $Q \subseteq P$ and any $\delta > 0$, a subset $Q' \subseteq Q$ of $O(1/\delta^k)$ points in $O(|Q| + 1/\delta^k)$ time, such that Q' is an δ-approximation of Q with respect to μ at any time. We outline a simple scheme that allows to maintain ε-approximations of P when points are inserted into or deleted from Q. Let n be an a-priori bound on the total number of insertions and deletions.

Let T be a balanced binary tree with n leaves. The height h of the tree is $O(\log n)$. Each point in the current set P is stored at one leaf of the tree; note that some leaves may be empty. Let $\delta = \varepsilon/2h$. Each internal node of the tree stores a set of $O(1/\delta^k)$ points. The point-set stored at a node v is defined by induction over the height of the nodes as follows: If $L(v)$ and $R(v)$ are the points stored at the left and right children of v, the point set stored at v is a δ-approximation of $L(v) \cup R(v)$. Our ε-approximation for P is an $\varepsilon/2$-approximation P' of the point set stored at the root of T.

It is straightforward to verify that a node at height i stores an $i\delta$-approximation of the subset of P stored in the leaves of the subtree rooted at the node. Thus the root of T stores an $\varepsilon/2$-approximation of P, and P' is an ε-approximation of P. When a point is inserted into P, it is assigned to some empty leaf; when a point is deleted from P, the corresponding leaf becomes empty. In either case, we only have to recompute the point sets stored at the $h + 1$ nodes on a single leaf-to-root path in T. We do this re-computation starting from the leaf and proceeding to the root. At each node v on this path, we have to compute a δ-approximation of $L(v) \cup R(v)$, which takes $O(1/\delta^k)$ time. The overall time for updating the structure after an insertion or deletion is therefore $O(\log n(\log n/\varepsilon)^k)$.

Theorem 5.9 *Let P be a set of points which we update (via an insertion or deletion) a total of n times. Let μ be a measure with the property that we can compute, given any $Q \subseteq P$ and any $\delta > 0$, a subset $Q' \subseteq Q$ of $O(1/\delta^k)$ points in $O(|Q| + 1/\delta^k)$ time, such that Q' is an δ-approximation of Q with respect to μ at any time. For a given $\varepsilon > 0$, we can maintain a set of $O(1/\varepsilon^k)$ points that ε-approximates P at a cost of $O(\log n(\log n/\varepsilon)^k)$ per insertion/deletion.*

6 Conclusions

In this paper, we have presented a general and unified technique for shape approximation. The new technique relies on the observation that some of the shape approximation/optimization problems can be linearized. Once linearized, using duality and convex-shape approximation techniques one can derive an efficient approximation algorithms. Our applications include near-linear time approximation algorithms for min-width cylindrical shell containing a point-set, min-width annulus, general technique for approximating convex measures, and general technique for approximating convex measures of moving point-set. We believe that there are a lot of other applications of our technique.

To some extent, our algorithm is the ultimate approximation algorithm for such problems: It has linear dependency on n, and a polynomial dependency on $1/\varepsilon$. The existence of such a general (and fast) approximation algorithm is quite surprising.

A possible direction for future research is to investigate how practical is this technique, improving it, and improving existing algorithms for various special problems. In particular, it seems that faster algorithms should exist for the problems of approximating the diameter and width of a point-set.

Acknowledgments

The authors thank Imre Barany, Jeff Erickson, Cecilia Magdalena Procopiuc, and Micha Sharir for helpful discussions.

References

[AAHS00] P.K. Agarwal, B. Aronov, S. Har-Peled, and M. Sharir. Approximation and exact algorithms for minimum-width annuli and shells. *Discrete Comput. Geom.*, 24(4):687–705, 2000.

[AAS97] P.K. Agarwal, B. Aronov, and M. Sharir. Computing envelopes in four dimensions with applications. *SIAM J. Comput.*, 26:1714–1732, 1997.

[AAS00] P.K. Agarwal, B. Aronov, and M. Sharir. Exact and approximation algorithms for minimum-width cylindrical shells. In *Proc. 11th ACM-SIAM Sympos. Discrete Algorithms*, pages 510–517, 2000.

[AEGH98] P.K. Agarwal, D. Eppstein, L.J. Guibas, and M. Henzinger. Parametric and kinetic minimum spanning trees. In *Proc. 39th Annu. IEEE Sympos. Found. Comput. Sci.*, pages 596–605, 1998.

[AGHV97] P.K. Agarwal, L.J. Guibas, J. Hershberger, and E. Veach. Maintaining the extent of a moving point set. In *Proc. 5th Workshop Algorithms Data Struct.*, volume 1272 of *Lecture Notes Comput. Sci.*, pages 31–44. Springer-Verlag, 1997.

[AGMV97] P.K. Agarwal, L.J. Guibas, T.M. Murali, and J.S. Vitter. Cylindrical static and kinetic binary space partitions. In *Proc. 13th Annu. ACM Sympos. Comput. Geom.*, pages 39–48, 1997.

[AGSS89] A. Aggarwal, L.J. Guibas, J. Saxe, and P.W. Shor. A linear-time algorithm for computing the Voronoi diagram of a convex polygon. *Discrete Comput. Geom.*, 4(6):591–604, 1989.

[AH01] P.K. Agarwal and S. Har-Peled. Maintaining the approximate extent measures of moving points. In *Proc. 12th ACM-SIAM Sympos. Discrete Algorithms*, pages 148–157, 2001.

[AM94] P.K. Agarwal and J. Matoušek. On range searching with semialgebraic sets. *Discrete Comput. Geom.*, 11:393–418, 1994.

[AS96] P.K. Agarwal and M. Sharir. Efficient randomized algorithms for some geometric optimization problems. *Discrete Comput. Geom.*, 16:317–337, 1996.

[AST94] P.K. Agarwal, M. Sharir, and S. Toledo. Applications of parametric searching in geometric optimization. *J. Algorithms*, 17:292–318, 1994.

[BGH97] J. Basch, L.J. Guibas, and J. Hershberger. Data structures for mobile data. In *Proc. 8th ACM-SIAM Sympos. Discrete Algorithms*, pages 747–756, 1997.

[BH01] 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.

[Cha00] T.M. Chan. Approximating the diameter, width, smallest enclosing cylinder and minimum-width annulus. In *Proc. 16th Annu. ACM Sympos. Comput. Geom.*, pages 300–309, 2000.

[Dud74] R.M. Dudley. Metric entropy of some classes of sets with differentiable boundaries. *J. Approx. Theory*, 10(3):227–236, 1974.

[EFNN89] H. Ebara, N. Fukuyama, H. Nakano, and Y. Nakanishi. Roundness algorithms using the Voronoi diagrams. In *Abstracts 1st Canad. Conf. Comput. Geom.*, page 41, 1989.

[EGS86] H. Edelsbrunner, L.J. Guibas, and J. Stolfi. Optimal point location in a monotone subdivision. *SIAM J. Comput.*, 15(2):317–340, 1986.

[Gär95] B. Gärtner. A subexponential algorithm for abstract optimization problems. *SIAM J. Comput.*, 24:1018–1035, 1995.

[GK92] P. Gritzmann and V. Klee. Inner and outer j-radii of convex bodies in finite-dimensional normed spaces. *Discrete Comput. Geom.*, 7:255–280, 1992.

[GLR97] J. García-Lopez and P. Ramos. Fitting a set of points by a circle. In *Proc. 13th Annu. ACM Sympos. Comput. Geom.*, pages 139–146, 1997.

[Gru93] P. M. Gruber. Aspects of approximation of convex bodies. In Peter M. Gruber and J. M. Wills, editors, *Handbook of Convex Geometry*, volume A, chapter 1.10, pages 319–345. North-Holland, Amsterdam, Netherlands, 1993.

[HS01] J. Hershberger and S. Suri. Simplified kinetic connectivity for rectangles and hypercubes. In *Proc. 12th ACM-SIAM Sympos. Discrete Algorithms*, pages 158–167, 2001.

[LL91] V.B. Le and D.T. Lee. Out-of-roundness problem revisited. *IEEE Trans. Pattern Anal. Mach. Intell.*, PAMI-13(3):217–223, 1991.

[MSY97] K. Mehlhorn, T.C. Shermer, and C.K. Yap. A complete roundness classification procedure. In *Proc. 13th Annu. ACM Sympos. Comput. Geom.*, pages 129–138, 1997.

[PS85] F.P. Preparata and M.I. Shamos. *Computational Geometry: An Introduction.* Springer-Verlag, New York, NY, 1985.

[Riv79] T.J. Rivlin. Approximating by circles. *Computing*, 21:93–104, 1979.

[RLW91] U. Roy, C.R. Liu, and T.C. Woo. Review of dimensioning and tolerancing: Representation and processing. *Comput. Aided Design*, 23(7):466–483, 1991.

[RZ92] U. Roy and X. Zhang. Establishment of a pair of concentric circles with the minimum radial separation for assessing roundness error. *Comput. Aided Design*, 24(3):161–168, 1992.

[SJ99] M. Smid and R. Janardan. On the width and roundness of a set of points in the plane. *Int. J. Comput. Geom. Appls.*, 9(1):97–108, 1999.

[SY95] T.C. Shermer and C.K. Yap. Probing for near centers and relative roundness. In *Proc. ASME Workshop on Tolerancing and Metrology*, 1995.

[YC97] C.K. Yap and E.-C. Chang. Issues in the metrology of geometric tolerancing. In J.-P. Laumond and M. H. Overmars, editors, *Robotics Motion and Manipulation*, pages 393–400. A.K. Peters, 1997.