Approximate Shape fitting via Linearization

Sariel Har-Peled* Kasturi R. Varadarajan†

27/04/2001 09:56

Abstract

Shape fitting is a fundamental optimization problem in computer science. In this paper, we present a general and unified technique for solving a certain family of such problems. 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 without resorting to a careful and painful case by case analysis, as was previously done for those problems. Furthermore, for several of those problems our results are considerably simpler and faster than what was previously known. In particular, for the minimum width cylindrical shell problem, our solution is the first algorithm whose running time is subquadratic in n. (In fact we get running time linear in n.)

1 Introduction

Given a set of points P in $\mathbb{R}^d$, the *shape fitting* is the problem of finding the best shape (for example, hyperplane) that best fits the point-set. Shape fitting is a fundamental optimization problem and has numerous usages in graphics (shape simplification, collusion detection), learning, datamining [FL95], databases [AWY$^+$99] (projective clustering), metrology, compression, and geometric optimization.

A restricted variant of this problem is where the shape to be fitted to the input is defined by a (small) constant number of parameters. Such problems fall into the realm of geometric optimization, and numerous variants were solved. However, no unified theory has evolved, and solutions are usually based on case by case analysis. In this paper, we present a general technique for approximate shape fitting for such variants.

We handle both the classical static case, and also apply it to the more general (and considerably harder) case when the points are moving (i.e., kinetic settings [BGH97]). For example, imagine that one would like to maintain the minimum volume bounding box that contains a set of moving points. Our technique enable us to compute such a bounding box (which is also moving) that contains the point-set, and the volume of this box, at any point in time, is close to the optimal volume of a bounding box that contains the point-set at this point in time. Such a moving bounding box is useful, for example, in implementing an efficient data-structure for collision detection of moving objects.

*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

If the shape we are trying to fit is convex (i.e., find a cylinder that contains the point-set), then it is implied by previous work [BH99] 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 convex shape approximation is (relatively) well understood [Gru93], this results in fast and efficient approximation algorithms. We state this more formally in Section 3.1, where we present a unified approach for this type of problems.

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. Despite the assumptions, some of those works are quite complicated, involved and use non-trivial insights into the problem at hand.

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 [AM94], and thus our approximation technique is quite broad, as it can be applied to optimization problems of this type. In particular, we show that it suffices to approximate the shape for a small subset of the original set of points. Such a result was not known for nonconvex shapes despite the substantial research done on those problems. In the following, we give more details about the applications of this technique that we present in this paper.

Cylindrical shell problem. Given a line ℓ in $\mathbb{R}^d$ and two real numbers $0 \le r \le R$, 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.,

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

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. Since computing the optimal shell is so expensive, they look at the approximate version and 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 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 $0 \le r \le R$, 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, AAHS99] 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 $O(n^{3-\frac{1}{19}})$ 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.

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 $\nu(S)$ 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 ν if at any time t, $(1-\varepsilon)\nu(P(t)) \leq \nu(Q(t))$. 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.

As mentioned above, these problems are motivated by applications in wireless communication, ad-hoc networks, collision detection, clustering, etc. (see [AEGH98, AGMV97, BGH97, AH01, HS01]). 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.

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. What we essentially show is that the convex hull of Q approximates the convex hull of P at all times. This means that Q is an ε-approximation of P with respect to any "convex" measure. Such results were previously known only for diameter and minimum-radius enclosing ball, see [AH01]. These results generalize to algebraic motion and to "non-convex" measures like minimum-width spherical/cylindrical shell.

The paper is organized as follows: In Section 2, we introduce the basic tools, and prove some technical results. In Section 3, we present the applications of the new technique. We conclude in Section 4.

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 $\mathfrak{L}_{\mathcal{F}} : \mathbb{R}^k \to \mathbb{R}$

$$\mathfrak{L}_{\mathcal{F}}(\mathbf{x}) = \min_{f \in \mathcal{F}} f(\mathbf{x}).$$

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 ε-approximation of $\mathcal{F}$ if for each $\mathbf{x} \in \mathbb{R}^k$,

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

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 $\mathbf{x} \in \Delta$,

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

2.1 Linear functions

Theorem 2.1 ([AH01]) *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 $\mathcal{H}$.*

If we look at the dual, the above theorem says that one can approximate an n-point set in $\mathbb{R}^{k+1}$ by an $O(1/\varepsilon^k)$-point subset such that for any line ℓ in $\mathbb{R}^{k+1}$, the interval on ℓ spanned by the projection of the original set on ℓ is approximated by the interval corresponding to this subset; see Section 3.1 for a formal statement. The theorem is thus a consequence of a well-known convex-hull approximation technique [Gru93, BH99]. We improve this result further, as follows:

Theorem 2.2 *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: Follows by using the algorithm of Gärtner [Gär95] together with Dudley's approximation technique [Dud74] on the set generated by Theorem 2.1. The (straightforward but tedious) details are omitted and will appear in the full-version. ■

Theorem 2.3 *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 $O(1/\varepsilon^k)$ simplices, so that for each simplex $\Delta \in \mathcal{C}$ there are two associated linear functions, $h'_\Delta, h''_\Delta \in \mathcal{H}$, that ε-approximate $\mathcal{H}$ inside Δ.*

Proof: Omitted. Will appear in the full-version. ■

Remark 2.4 Theorem 2.3 considerably improves over the result of Agarwal and Har-Peled [AH01]. In particular, the algorithm of [AH01] takes $O((n/\varepsilon^k)\log(1/\varepsilon))$ time, and outputs a decomposition of size $O(1/\varepsilon^k \log(1/\varepsilon))$.

2.2 Polynomials

Let $\mathcal{F} = \{f_1, \dots, f_n\}$ 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 $x_1, \dots, x_k$ denote the variables over which the polynomials in $\mathcal{F}$ are defined. We map each monomial over $x_1, \dots, x_k$ that occurs in $\mathcal{F}$ to a distinct variable u_i. Let $u_1, \dots, u_{k'}$ be the resulting variables. Each polynomial $f_i \in \mathcal{F}$ now becomes a linear function h_i over $\mathbb{R}^{k'}$. Let $\mathcal{H} = \{h_1, \dots, h_n\}$ 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 $\mathcal{F}$.

Theorem 2.5 *Given a family of k-variate polynomials $\mathcal{F} = \{f_1, \dots, f_n\}$ 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 2.3 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 $x_1, \dots, x_k$) and the degree of these polynomials is also bounded by a constant.

Theorem 2.6 *Given a family of k-variate polynomials $\mathcal{F} = \{f_1, \dots, f_n\}$ 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 $\mathcal{F}$.*

2.3 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 2.8 below).

Theorem 2.7 *Let $\mathcal{F} = \{(f_1)^{1/2}, \dots, (f_n)^{1/2}\}$ be a family of k-variate functions (over $\mathbf{x} = (x_1, \dots, x_k) \in \mathbb{R}^k$), 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 $f_1, \dots, f_n$.*

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 $\{f_1, \dots, f_n\}$. Using the algorithm of Theorem 2.5, 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 2.8 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 2.8 *Let $0 \le a \le A \le B \le b$, and $0 < \varepsilon \le 1$ be given parameters, so that: (i) $I = B - A$, (ii) $A - a \le (\delta/2)I$, and (iii) $b - B \le (\delta/2)I$, where $\delta = \varepsilon^2/32$. Then, $\sqrt{A} - \sqrt{a} \le (\varepsilon/2)U$, and $\sqrt{b} - \sqrt{B} \le (\varepsilon/2)U$, where $U = \sqrt{B} - \sqrt{A}$.*

Proof: We have

$$\begin{aligned}\sqrt{b} - \sqrt{B} &= \frac{b - B}{\sqrt{b} + \sqrt{B}} \le \frac{(\delta/2)(B - A)}{2\sqrt{B}} \le \frac{\delta}{4} \cdot \frac{(\sqrt{B} - \sqrt{A})(\sqrt{B} + \sqrt{A})}{\sqrt{B}} \\ &\le \frac{\delta}{2}(\sqrt{B} - \sqrt{A}) \le \frac{\varepsilon}{2}U.\end{aligned}$$

If $A < B/2$, then $U = \sqrt{B} - \sqrt{A} \ge (1 - \sqrt{1/2})\sqrt{B} \ge (1 - \sqrt{1/2})\sqrt{I}$. However, this implies that

$$\sqrt{A} - \sqrt{a} \le \sqrt{A - a} \le \sqrt{I\frac{\delta}{2}} \le \frac{U}{(1 - 1/\sqrt{2})}\sqrt{\frac{\delta}{2}} \le \varepsilon U \cdot \sqrt{\frac{1}{2 \cdot 32}} \cdot \frac{1}{1 - 1/\sqrt{2}} \le \frac{\varepsilon}{2}U.$$

If $A > B/2$ then $\sqrt{a} \ge \sqrt{A - \delta I} \ge \sqrt{B/2 - \delta(B/2)} = \sqrt{(1 - \delta)/2}\sqrt{B} \ge \sqrt{B}/3$, since δ is sufficiently small. So we have:

$$\begin{aligned}\sqrt{A} - \sqrt{a} &= \frac{A - a}{\sqrt{A} + \sqrt{a}} \le \frac{(\delta/2)(B - A)}{2\sqrt{B}/3} \le \delta \cdot \frac{(\sqrt{B} - \sqrt{A})(\sqrt{B} + \sqrt{A})}{\sqrt{B}} \\ &\le 2\delta(\sqrt{B} - \sqrt{A}) = 2\delta U \le \frac{\varepsilon}{2}U.\end{aligned}$$

■

Remark 2.9 It seems that the approach of Lemma 2.8 and Theorem 2.7 should work for any $p \in (0, 1)$. Namely, given a $p \in (0, 1)$, and a set of functions $\mathcal{F} = \{(f_1)^p, \ldots, (f_n)^p\}$, then approximating the extent of $f_1, \ldots, f_n$ up to an appropriate constant $\mu(\varepsilon, p)$, should yield a ε-approximation to the extent of $\mathcal{F}$. We leave the question of determining exact bounds on $\mu(\varepsilon, p)$ as an open question for further research.

Theorem 2.10 *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 2.11 In many applications, we can give a better bound for k' by being more careful with the linearization. In the proof of Theorem 2.5, 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 2.6, Theorem 2.7, and Theorem 2.10. Agarwal and Matoušek [AM94] describe an algorithm for finding the minimum number of monomials needed to linearize a family of polynomials.

3 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 our problems to convex-hull approximation. Convex hull approximation has been used before for shape fitting with convex shapes [BH99, 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.

3.1 Convex Shape Fitting in a General Setting

Let P be a set of n points in $\mathbb{R}^d$. A point-set Q is an *ε-approximation* to P, if for any $\mathbf{x} \in \mathbb{R}^d$, we have $(1-\varepsilon)I_P(\mathbf{x}) \subseteq I_Q(\mathbf{x}) \subseteq I_P(\mathbf{x})$, where

$$I_P(\mathbf{x}) = \left[\min_{\mathbf{p}\in P} \mathbf{x}\cdot\mathbf{p}, \max_{\mathbf{p}\in P} \mathbf{x}\cdot\mathbf{p}\right]$$

is an interval on the real line, and $(1-\varepsilon)I_P(\mathbf{x})$ the result of shrinking this interval by a factor of $(1-\varepsilon)$ and recentering it at the center of $I_P(\mathbf{x})$. The following is a restatement of Theorem 2.1 and Theorem 2.2. (These theorems are in fact proved in the dual setting.)

Theorem 3.1 *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 3.2 A function $\nu(\cdot)$ defined over subsets of $\mathbb{R}^d$, is *convex measure*, if: (i) for any $S \subseteq \mathbb{R}^d$, $\nu(S) \geq 0$, and (ii) $(1 - c\varepsilon)\nu(S) \leq \nu(Q) \leq \nu(S)$ if Q ε-approximates S, where c is a constant that depends on ν.

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 3.1 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 [BH99, Cha00]. However, our scheme does not require us to carefully inspect the problem at hand. Furthermore, it holds for any convex-measure.

3.2 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 $p = (p_1, \dots, p_d)$ 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 $\mathbf{x} = (x_1, \dots, x_d) \in \mathbb{R}^d$ and $\mathbf{y} = (y_1, \dots, y_d)$ 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 $x_1, \dots, x_d, y_1, \dots, y_d$. 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 2.7, we can therefore find a subset $Q \subseteq P$ of $O(1/\varepsilon^{O(d^2)})$ points such that for any line ℓ, $(1-\varepsilon)d(\ell, P) \leq d(\ell, Q)$. 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

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

Theorem 3.3 *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 2.10 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 $d(\ell^*, P) \leq \frac{1}{1-\varepsilon} d(\ell^*, Q) \leq \frac{1}{1-\varepsilon} w^*$.

Remark 3.4 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 52 times the optimal in 3d.

Remark 3.5 Theorem 3.3 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].

3.3 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 2.10 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 3.6 *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.

3.4 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, $(1-c\varepsilon)\,\mathrm{diam}(P(t)) \le \mathrm{diam}(Q(t)) \le \mathrm{diam}(P(t))$. 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 most of those *exact* measures efficiently for a moving point-set. For the case of the diameter, a result of Agarwal *et al.* [AGHV97] shows that the diameter of a point set under linear motion in the plane can change quadratic number of times.

Theorem 3.7 *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 2.5. ■

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 from this abstract.

4 Conclusions

In this paper, we had 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 belief 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, considering the substantial amount of research done on specific problems that can be solved using this new algorithm.

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.

References

[AAHS99] P.K. Agarwal, B. Aronov, S. Har-Peled, and M. Sharir. Approximation and exact algorithms for minimum-width annuli and shells. In *Proc. 15th Annu. ACM Sympos. Comput. Geom.*, pages 380–389, 1999.

[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.

[AWY+99] C.C. Aggarwal, J.L. Wolf, P.S. Yu, C. Procopiuc, and J.S. Park. Fast algorithms for projected clustering. In *SIGMOD 99*, pages 61–72, 1999.

[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.

[BH99] G. Barequet and S. Har-Peled. Efficiently approximating the minimum-volume bounding box of a point set in three dimensions. In *Proc. 10th ACM-SIAM Sympos. Discrete Algorithms*, pages 82–91, 1999.

[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.

[FL95] C. Faloutsos and K. I. Lin. FastMap: a fast algorithm for indexing, data-mining and visualization of traditional and multimedia datasets. *SIGMOD Record (ACM Special Interest Group on Management of Data)*, 24(2):163–174, June 1995.

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

[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.